

Incidence bounds via extremal graph theory

Aleksa Milojević*, Benny Sudakov*, István Tomon†

Abstract

The study of counting point-hyperplane incidences in the d -dimensional space was initiated in the 1990's by Chazelle and became one of the central problems in discrete geometry. It has interesting connections to many other topics, such as additive combinatorics and theoretical computer science. Assuming a standard non-degeneracy condition, i.e., that no s points are contained in the intersection of s hyperplanes, the currently best known upper bound on the number of incidences of m points and n hyperplanes in \mathbb{R}^d is

$$O_{d,s}((mn)^{1-1/(d+1)} + m + n).$$

This bound by Apfelbaum and Sharir is based on geometrical space partitioning techniques, which apply only over the real numbers.

In this paper, we propose a novel combinatorial approach to study such incidence problems over arbitrary fields. Perhaps surprisingly, this approach matches the best known bounds for point-hyperplane incidences in \mathbb{R}^d for many interesting values of m, n, d , e.g. when $m = n$ and d is odd. Moreover, in finite fields our bounds are sharp as a function of m and n in every dimension. We also study the size of the largest complete bipartite graph in point-hyperplane incidence graphs with a given number of edges and obtain optimal bounds as well.

Additionally, we study point-variety incidences and unit-distance problem in finite fields, and give tight bounds for both problems under a similar non-degeneracy assumption. We also resolve Zarankiewicz type problems for algebraic graphs. Our proofs use tools such as induced Turán problems, VC-dimension theory, evasive sets and Hilbert polynomials. Also, we extend the celebrated result of Rónyai, Babai and Ganapathy on the number of zero-patterns of polynomials to the context of varieties, which might be of independent interest.

1 Introduction

The Szemerédi-Trotter theorem [52] is a fundamental result in combinatorial geometry, giving a sharp upper bound on the number of incidences in point-line configurations. It states that m points and n lines on the real plane determine at most $O((mn)^{2/3} + m + n)$ incidences, and as Erdős showed (see e.g. [17]), this bound is the best possible. This deep result found numerous applications and inspired a large number of generalizations and extensions over the last decades. For example, the problem of counting point-line incidences in \mathbb{R}^3 under certain non-degeneracy conditions has been studied by Guth and Katz [27] in their solution to the Erdős distinct distances problem. The Szemerédi-Trotter theorem has also been applied by Elekes [18] to derive sum-product estimates over the reals. For these and other applications of incidence theorems, we refer the reader to the survey [13].

Extending incidence results to finite fields is more tricky, since standard space partitioning techniques do not apply anymore. The importance of point-line incidence bounds in \mathbb{F}_p^2 stems from their close connection to sum-product estimates, as shown in the pioneering work of Bourgain, Katz and Tao [5]. Following this work, point-line incidence bounds over finite fields have been extensively studied, see e.g. [31, 37, 49].

In this paper, we consider several natural incidence questions in higher dimensions. The first question we address is to find the maximum number of incidences between m points and n hyperplanes in \mathbb{F}^d , for an arbitrary field \mathbb{F} . In dimension 3 and above, it is possible that all n hyperplanes

*ETH Zurich, e-mail: {aleksa.milojevic, benjamin.sudakov}@math.ethz.ch. Research supported in part by SNSF grant 200021_196965.

†Umeå University, e-mail: istvan.tomon@umu.se.

contain all m points: take m points on a line, and n hyperplanes containing this line. In order to avoid such trivialities, it is standard to impose the further condition that the *incidence graph* is $K_{s,s}$ -free, where $K_{s,s}$ denotes the complete bipartite graph with vertex classes of size s . Here, we think of s as a large constant, which may depend on d , but no other parameters. In geometric terms, this condition says that there do not exist s hyperplanes whose intersection contains more than s points.

Given a set of points \mathcal{P} and set of geometric objects (e.g. hyperplanes) \mathcal{H} , the *incidence graph* of $(\mathcal{P}, \mathcal{H})$ is the bipartite graph $G(\mathcal{P}, \mathcal{H})$ with vertex classes \mathcal{P} and \mathcal{H} , and $\{x, H\} \subset \mathcal{P} \cup \mathcal{H}$ is an edge if $x \in H$. We denote by $I(\mathcal{P}, \mathcal{H})$ the number of edges of the incidence graph, i.e. $I(\mathcal{P}, \mathcal{H})$ is the number of incidences between \mathcal{P} and \mathcal{H} .

Beyond point-hyperplane incidences, we obtain sharp bounds for several other well-studied incidence problems. We find the maximum number of incidences between points and d -dimensional varieties of bounded degree over arbitrary fields assuming the incidence graph is $K_{s,s}$ -free. Then, we determine the maximum number of unit distances between n points in \mathbb{F}^d , assuming the unit distance graph is $K_{s,s}$ -free. Finally, we conclude with Zarankiewicz's problem for so called algebraic graphs.

1.1 Point-hyperplane incidences

The problem of bounding the number of edges in $K_{s,s}$ -free point-hyperplane incidence graphs has a long history, with the initial motivation coming from computational geometry. How compactly can one record the incidences in a configuration of points and hyperplanes? Instead of writing down whether each of the mn pairs $(x, H) \in \mathcal{P} \times \mathcal{H}$ forms an incidence, one can compress the incidence graph by representing it as the union of complete bipartite subgraphs. This cannot be done very efficiently in general, for example if $G = G_{n,1/2}$ is the Erdős-Rényi random graph, one needs a list of $\Omega(\frac{n^2}{\log n})$ vertices to encode any partition. However, for incidence graphs Chazelle [9] developed a space partitioning technique to show that this can be improved to $O_{d,s}(n^{2-\frac{2}{d+1}} \log n)$ when $m = n$. Another motivation comes from the counting version of Hopcroft's problem, which asks, given m points and n hyperplanes in \mathbb{R}^d , how fast can one determine the number of incidences. The work of Chazelle [9] achieved the first subquadratic algorithm for this problem and Erickson [21] gave a lower bound for the running time of algorithms classified as partitioning algorithms. He showed that particularly hard instances for determining the number of incidences are configurations $(\mathcal{P}, \mathcal{H})$ without a compact representation, i.e. a configuration without a large complete bipartite graph in the incidence graph.

Improving the works of Chazelle [9] and Brass and Knauer [6], Apfelbaum and Sharir [2] proved that if \mathcal{P} is a set of m points and \mathcal{H} is a set of n hyperplanes in \mathbb{R}^d , whose incidence graph contains no $K_{s,s}$, then

$$I(\mathcal{P}, \mathcal{H}) \leq O_{d,s} \left((mn)^{1-\frac{1}{d+1}} + m + n \right). \quad (1)$$

This upper bound has not been improved in the past twenty years, however, matching lower bounds are only known for $d = 2$, which coincides with the Szemerédi-Trotter theorem. The common feature in all of the proofs of (1) and related results is that they rely on certain space partitioning results such as cuttings [9, 10] or polynomial partitioning [27, 47]. These techniques are highly geometric and have no analogues over finite fields.

The currently best known lower bound for $d \geq 3$ was recently achieved by Sudakov and Tomon [51], improving on the constructions of Brass and Knauer [6] and Balko, Cibulka, and Valtr [3]: if s is sufficiently large with respect to d , then there exist a set of m points \mathcal{P} and a set of n hyperplanes \mathcal{H} in \mathbb{R}^d whose incidence graph is $K_{s,s}$ -free and

$$I(\mathcal{P}, \mathcal{H}) \geq \begin{cases} \Omega_d \left((mn)^{1-(2d+3)/(d+2)(d+3)} \right) & \text{if } d \text{ is odd,} \\ \Omega_d \left((mn)^{1-(2d^2+d-2)/(d+2)(d^2+2d-2)} \right) & \text{if } d \text{ is even.} \end{cases}$$

A natural variant of this problem asks about the maximum number of incidences between points and hyperplanes in a d -dimensional space \mathbb{F}^d for a general field \mathbb{F} . If $d = 2$ and p is a prime, taking

every point of \mathbb{F}_p^2 and $n < p^2$ arbitrary lines, we get $m = p^2$ points, and $np = \Theta(nm^{1/2})$ incidences, beating the Szemerédi-Trotter bound as long as $n \gg p$. On the other hand, a simple application of the Kővári-Sós-Turán theorem [32] shows that one cannot have more than $O(nm^{1/2} + m)$ incidences over any field, by noting that the incidence graph of points and lines is always $K_{2,2}$ -free. The work of Bourgain, Katz, and Tao [5] shows that better upper bounds can be obtained over \mathbb{F}_p assuming $m, n \leq p^{2-\delta}$.

However, point-hyperplane incidence bounds in higher dimensions turn out to be more elusive. One such result in three dimensions was obtained by Rudnev [43] (see also de Zeeuw [57] for a shorter proof). He proved that in \mathbb{F}_p^3 , if no s planes contain the same line, then there are at most $O(m\sqrt{n} + ms)$ incidences, assuming $n \leq p^2$. This point-plane incidence bound in \mathbb{F}_p^3 was later used to show various improved sum-product estimates over finite fields in [1, 37, 41]. For $d > 3$, barely anything is known about incidences of points and hyperplanes in \mathbb{F}_p^d .

In this paper, we propose a novel approach to derive bounds on the maximal number of incidences between m points and n hyperplanes in a d -dimensional vector space \mathbb{F}^d , assuming the incidence graph is $K_{s,s}$ -free. We combine new graph-theoretic techniques with probabilistic arguments and prove bounds which are sharp for the whole range of parameters m and n with a suitable choice of field. Surprisingly, for many interesting pairs of values (m, n) , our upper bound matches the upper bound (1). This is fairly unexpected, as all proofs of (1) rely on highly geometric techniques, while we employ only combinatorial ideas. To give a snippet of our most general result, we prove the following in the special case $m = n$.

Theorem 1.1. *Let d, s be positive integers and let \mathbb{F} be a field. If \mathcal{P} is a set of n points and \mathcal{H} is a set of n hyperplanes in \mathbb{F}^d such that no s points lie on the intersection of s hyperplanes, then*

$$I(\mathcal{P}, \mathcal{H}) \leq O\left((s + d^3)n^{2-1/\lceil \frac{d+1}{2} \rceil}\right).$$

Note that this bound matches the best known bound (1) when d is odd. Before stating the general bound, which depends on the relationship of m and n , let us introduce a piece of notation. Namely, we define the quantities $\alpha_t = \frac{t}{d+2-t}$ for $t \in \{2, \dots, d\}$ and $\beta_t = \frac{t}{d+1-t}$ for $t \in \{1, \dots, d\}$. Now we can state our main theorem in its full generality.

Theorem 1.2. *Let d, s, m, n be positive integers and $\alpha > 0$ such that $n = m^\alpha$, and let \mathbb{F} be a field. Let \mathcal{P} be a set of m points and \mathcal{H} be a set of n hyperplanes in \mathbb{F}^d such that the incidence graph of $(\mathcal{P}, \mathcal{H})$ is $K_{s,s}$ -free. Then*

$$I(\mathcal{P}, \mathcal{H}) \leq \begin{cases} O_{s,d}(m) & \text{if } \alpha \in (0, \beta_1], \\ O_{s,d}(m^{1-\frac{1}{d+2-t}}n) & \text{if } \alpha \in [\beta_{t-1}, \alpha_t] \text{ for some } t \in \{2, \dots, d\}, \\ O_{s,d}(mn^{1-\frac{1}{t}}) & \text{if } \alpha \in [\alpha_t, \beta_t] \text{ for some } t \in \{2, \dots, d\}, \\ O_{s,d}(n) & \text{if } \alpha \in [\beta_d, \infty). \end{cases}$$

Moreover, these bounds are tight, i.e. for every d , there exist $s = s(d)$ and $c = c(d)$ such that the following hold. For every $\alpha > 0$ and sufficiently large integer m , there exist a field $\mathbb{F} = \mathbb{F}(d, m, \alpha)$, a set of m points \mathcal{P} and a set of $n = \lfloor m^\alpha \rfloor$ hyperplanes \mathcal{H} in \mathbb{F}^d such that the incidence graph of $(\mathcal{P}, \mathcal{H})$ is $K_{s,s}$ -free and

$$I(\mathcal{P}, \mathcal{H}) \geq \begin{cases} m & \text{if } \alpha \in (0, \beta_1], \\ cm^{1-\frac{1}{d+2-t}}n & \text{if } \alpha \in [\beta_{t-1}, \alpha_t] \text{ for some } t \in \{2, \dots, d\}, \\ cmn^{1-\frac{1}{t}} & \text{if } \alpha \in [\alpha_t, \beta_t] \text{ for some } t \in \{2, \dots, d\}, \\ n & \text{if } \alpha \in [\beta_d, \infty). \end{cases}$$

Let us compare our result with (1) in the non-trivial regime $\alpha \in [1/d, d]$. Interestingly, in case $\alpha = \beta_t = \frac{t}{d+1-t}$ for some integer $t \in \{1, \dots, d\}$, Theorem 1.2 shows that the number of incidences is $O_{s,d}((mn)^{1-1/(d+1)})$, which exactly matches (1). The other extreme is when $\alpha = \alpha_t = \frac{t}{d+2-t}$ for

some $t \in \{2, \dots, d\}$, in which case we get $O_{s,d}((mn)^{1-1/(d+2)})$. For other values of α , our upper bound is $O_{s,d}((mn)^{1-\gamma})$ for some $\frac{1}{d+2} \leq \gamma \leq \frac{1}{d+1}$.

The second half of the theorem shows that these upper bounds cannot be improved unless further assumptions on the field \mathbb{F} are made. To show this, we consider optimal constructions of so called subspace evasive sets [16, 51]. We follow the ideas of [3, 6, 51] to transform large subspace evasive sets into point-hyperplane configuration with many incidences and no large bipartite subgraphs.

We also study another well-known problem about finding large complete bipartite graphs in dense incidence graphs. The motivation behind studying this problem is the intuitive understanding that a large number of incidences between points and hyperplanes is always explained by the existence of large complete bipartite subgraphs in the incidence graph. To make things more precise, given a set of points \mathcal{P} and set of hyperplanes \mathcal{H} , we define $\text{rs}(\mathcal{P}, \mathcal{H})$ as the maximum number of edges in a complete bipartite subgraph of the incidence graph of $(\mathcal{P}, \mathcal{H})$, i.e. the maximum of $r \cdot s$ over all r and s such that $K_{r,s}$ is a subgraph of the incidence graph. We consider configurations of m points and n hyperplanes with εmn incidences.

In the real d -dimensional space, Apfelbaum and Sharir [2] proved that $\text{rs}(\mathcal{P}, \mathcal{H}) = \Omega_d(\varepsilon^{d-1}mn)$ for $\varepsilon > \Omega(n^{-1/(d-1)})$. Also, if $\varepsilon > \Omega((mn)^{-\frac{1}{d-1}})$, then there exist a set of points \mathcal{P} and set of hyperplanes \mathcal{H} in \mathbb{R}^d with $\text{rs}(\mathcal{P}, \mathcal{H}) \leq O_d(\varepsilon^{\frac{d+1}{2}}mn)$. Note that in case $d = 3$, the lower and upper bounds match for $\varepsilon > \Omega(n^{-1/2})$. Do [12] improved the lower bound in \mathbb{R}^4 and \mathbb{R}^5 for a large range of m, n, ε to match the upper bound up to logarithmic factors. One might be also interested how the minimum of $\text{rs}(\mathcal{P}, \mathcal{H})$ depends on the dimension d as well, which is a question motivated by certain variants of the celebrated log-rank conjecture of Lovász and Saks [35]. In this regime, it follows from Fox, Pach, and Suk [23] that $\text{rs}(\mathcal{P}, \mathcal{H}) > \varepsilon^{d+1}2^{-O(d \log d)}mn$. For constant ε , this is improved by a recent result of Singer and Sudan [46], who obtained an exponential dependence on d in the lower bound $\text{rs}(\mathcal{P}, \mathcal{H}) = \Omega(\varepsilon^{2d}mn/d)$. We refer the interested reader to [46] about the relationship of this problem and the log-rank conjecture.

Here, we extend and improve some of these results by considering arbitrary fields, and obtain tight bounds in all cases. Following Apfelbaum and Sharir [2], we define $\text{rs}_d(n, m, I) = \min \text{rs}(\mathcal{P}, \mathcal{H})$, where the minimum is taken over all fields \mathbb{F} , all sets $\mathcal{P} \subset \mathbb{F}^d$ of at most m points, all sets \mathcal{H} of at most n hyperplanes in \mathbb{F}^d , that satisfy $I(\mathcal{P}, \mathcal{H}) \geq I$. Note that requiring that \mathcal{P} is a set of at most m points instead of exactly m points makes no crucial difference, since one can always add points to the set \mathcal{P} which are incident to no hyperplanes (if \mathbb{F} is a finite field, one might need to pass to an extension of \mathbb{F}). A similar remark holds for hyperplanes.

Theorem 1.3. *Let d and $m \leq n$ be positive integers, $\varepsilon \in (0, 1)$ and $I = \varepsilon mn$. Then*

$$\text{rs}_d(m, n, I), \text{rs}_d(n, m, I) = \begin{cases} \Theta_d(\varepsilon^{d-1}mn) & \text{if } \varepsilon > 100 \max\{m^{-\frac{1}{d-1}}, n^{-\frac{1}{d}}\}, \\ \Theta_d(\varepsilon n) & \text{if } \varepsilon < \frac{1}{4} \max\{m^{-\frac{1}{d-1}}, n^{-\frac{1}{d}}\}. \end{cases}$$

In the regime $\varepsilon \gg \max\{m^{-\frac{1}{d-1}}, n^{-\frac{1}{d}}\}$, the implied constant of the lower bound is only exponential in d . Therefore, our theorem matches the dependence on ε coming from the result of Apfelbaum and Sharir [2], as well as the exponential dependence on the dimension obtained by Sudan and Singer [46].

On the other hand, in case $\varepsilon \ll \max\{m^{-\frac{1}{d-1}}, n^{-\frac{1}{d}}\}$, Theorem 1.3 tells us that the largest complete bipartite subgraph might be just a star. Finally, observe that in case $m = n$, there is a large jump around $\varepsilon \approx n^{-1/d}$, where $\text{rs}_d(n, n, I)$ suddenly jumps from $\approx n^{1-1/d}$ to $\approx n^{1+1/d}$.

1.2 Point-variety incidences

The number of incidences between points and varieties in the real space has been extensively studied (see Section 4.1 for a gentle introduction on varieties). Pach and Sharir [39, 40] proved that if \mathcal{P} is a set of m points and \mathcal{V} is a set of n algebraic curves of degree at most Δ in \mathbb{R}^2 such that $G(\mathcal{P}, \mathcal{V})$ is $K_{s,t}$ -free, then

$$I(\mathcal{P}, \mathcal{V}) = O_{s,t,\Delta} \left(m^{\frac{s}{2s-1}} n^{\frac{2s-2}{2s-1}} + m + n \right).$$

Zahl [54] and Basu and Sombra [4] extended this bound to dimension 3 and 4, respectively. Subsequently, Fox, Pach, Sheffer, Suk, and Zahl [22] established the following general bound in \mathbb{R}^D : if \mathcal{P} is a set of m points and \mathcal{V} is a set of n varieties in \mathbb{R}^D of degree at most Δ , and $G(\mathcal{P}, \mathcal{V})$ is $K_{s,t}$ -free, then

$$I(\mathcal{P}, \mathcal{V}) = O_{D,s,t,\Delta,\varepsilon} \left(m^{\frac{(D-1)s}{Ds-1} + \varepsilon} n^{\frac{D(s-1)}{Ds-1}} + m + n \right), \quad (2)$$

for every $\varepsilon > 0$. However, this bound is only known to be tight in the special case when $s = 2$ and t is sufficiently large, see [45].

Here, we consider a similar problem over arbitrary fields. That is, we consider the problem of bounding the number of incidences between points and varieties in \mathbb{F}^D , where \mathbb{F} is an arbitrary field, under the assumption that the incidence graph is $K_{s,s}$ -free (our proofs and results remain essentially the same if we consider the more general problem of forbidding $K_{s,t}$ for $t \geq s$, so we assume $s = t$ to simplify notation).

Theorem 1.4. *Let \mathcal{P} be a set of m points and let \mathcal{V} be a set of n varieties in \mathbb{F}^D , each of dimension d and degree at most Δ . If the incidence graph $G(\mathcal{P}, \mathcal{V})$ is $K_{s,s}$ -free, then*

$$I(\mathcal{P}, \mathcal{V}) \leq O_{d,\Delta,s}(m^{\frac{d}{d+1}}n + m).$$

Let us compare our bound with (2). As $s \rightarrow \infty$, the bound in (2) approaches the shape $m^{\frac{D-1}{D}}n + m + n$, which matches the one in Theorem 1.4 in case $d = D - 1$. For smaller d , Theorem 1.4 provides even better bounds. Curiously, our bound does not depend on the dimension D of the ambient space, it only depends on the degree and the dimension of the varieties. Moreover, our bounds are essentially tight, as the next theorem shows.

Theorem 1.5. *Let $d < D$ be positive integers, $\alpha > 0$, and m, n be positive integers such that $n = \lfloor m^\alpha \rfloor$, and m, n are sufficiently large with respect to D and α . Then there exist a prime p , a set of m points $\mathcal{P} \subseteq \mathbb{F}_p^D$ and a set \mathcal{V} of n varieties in \mathbb{F}_p^D of dimension d and degree at most $\Delta = \lceil (1 + \alpha)(d + 1) \rceil^2$ such that $G(\mathcal{P}, \mathcal{V})$ does not contain $K_{s,s}$ with $s = \Delta^{1/2}$, and*

$$I(\mathcal{P}, \mathcal{V}) \geq \Omega_{d,\alpha}(m^{\frac{d}{d+1}}n).$$

We provide two proofs of Theorem 1.4. The first one uses the framework of induced Turán problems. We construct a bipartite graph $H_{d,\Delta}$ such that incidence graphs of points and varieties in \mathbb{F}^D of dimension d and degree at most Δ avoid $H_{d,\Delta}$ as an induced subgraph, and every vertex of $H_{d,\Delta}$ in one of the parts has degree at most $d + 1$. Then, we adapt the ideas of [30] to prove bounds on the number of edges of induced $H_{d,\Delta}$ -free and $K_{s,s}$ -free bipartite graphs.

An alternative proof of Theorem 1.4 is obtained by establishing a bound on the VC-dimension of incidence graphs and the rate of growth of their shatter functions. To do this we bound the number of incidence-patterns between points and lower-dimensional varieties, extending the celebrated result of Rónyai, Babai and Ganapathy [42] on the number of zero-patterns of polynomials.

We remark that graphs with bounded VC-dimension avoid induced copies of a certain forbidden bipartite subgraph with similar properties as $H_{d,\Delta}$, but our explicit construction of $H_{d,\Delta}$ is more straightforward. Moreover, in certain situations, like the unit-distance problem for instance, finding specific forbidden bipartite graphs leads to better bounds compared to the ones which follow from VC-dimension arguments.

Finally, in order to prove Theorem 1.5 as well as other constructions in this paper, we use the results from [51], which are based on the random algebraic method originally introduced by [8].

1.3 Unit distances in finite fields

The celebrated Erdős unit distance problem [19, 20] is one of the most notorious open problems in combinatorial geometry, asking to estimate the function $f_d(n)$, the maximum number of unit distances spanned by n points in \mathbb{R}^d . The cases $d = 2, 3$ are the most difficult ones, with a large

gap remaining between the best known lower and upper bounds. For $d = 2$, the state-of-the-art is $n^{1+\Omega(1/\log \log n)} < f_2(n) < O(n^{4/3})$, where the lower bound is due to Erdős [20], and the upper bound is due to Spencer, Szemerédi and Trotter [48]. For $d = 3$, we know that $\Omega(n^{4/3} \log \log n) < f_3(n) < n^{1.498}$ by [20] and [55], respectively. On the other hand, for $d \geq 4$, it is not difficult to construct a set of points achieving $f_d(n) = \Theta_d(n^2)$ [33]. Indeed, if $d \geq 4$, one can take two orthogonal linear subspaces V_1 and V_2 , each of dimension at least 2. Then, placing $n/2$ points on the origin-centered sphere $S_i \subset V_i$ of radius $1/\sqrt{2}$ for $i = 1, 2$, we get at least $n^2/4$ unit distances. However, as it is shown in [22], in some sense this construction is the only way one can get a quadratic number of unit distances. More precisely, if we assume that no s points are contained in a $(d-3)$ -dimensional sphere, then the maximum number of unit distances is at most

$$O_{d,s,\varepsilon}(n^{2-\frac{2}{d+1}+\varepsilon})$$

for every $\varepsilon > 0$ (see also [25] for further strengthening). The key observation is that such unit distance graphs are semi-algebraic of complexity at most 2 containing no copy of $K_{s,s}$. The bounds on the number of edges of such graphs follow from the more general bounds (3) on Zarankiewicz's problem for algebraic graphs, which will be discussed in the next section. In the context of finite fields, this problem also received attention when $d = 3$. In this dimension, there are stronger natural restrictions on the unit-distance graphs, beyond being $K_{s,s}$ -free. Under these conditions, Zahl [56] proved that a set of n points determines at most $O(n^{3/2})$ unit distances assuming -1 is not a square in \mathbb{F} , while Rudnev [44] established the upper bound $O(n^{8/5})$ without this assumption.

Here, we completely resolve the problem of bounding the number of unit distances over arbitrary fields, assuming the unit distance graph is $K_{s,s}$ -free. A *unit sphere* in \mathbb{F}^d with center $a = (a_1, \dots, a_d) \in \mathbb{F}^d$ is defined as the set of points $(x_1, \dots, x_d) \in \mathbb{F}^d$ which satisfy

$$(x_1 - a_1)^2 + \dots + (x_d - a_d)^2 = 1.$$

Furthermore, points a and b are at *unit distance* if b is contained in the unit sphere with center a .

Theorem 1.6. *Let \mathcal{P} be a set of n points in \mathbb{F}^d . If the unit distance graph of \mathcal{P} does not contain $K_{s,s}$, then the number of unit distances spanned by \mathcal{P} is at most $O_{d,s}(n^{2-\frac{1}{\lceil d/2 \rceil + 1}})$.*

Furthermore, this bound is sharp in the following sense.

Theorem 1.7. *Let n, d be positive integers. There exists a constant $s = s(d)$, a finite field \mathbb{F}_q and a set of n points $\mathcal{P} \subseteq \mathbb{F}_q^d$ such that the unit distance graph on \mathcal{P} does not contain $K_{s,s}$ and \mathcal{P} spans $\Omega_d(n^{2-\frac{1}{\lceil d/2 \rceil + 1}})$ unit distances.*

1.4 Algebraic graphs

Zarankiewicz's problem asks for the maximum number of edges in an n -vertex graph without a copy of $K_{s,s}$. Fox, Pach, Sheffer, Suk, and Zahl [22] considered a variant of this problem in which case the host graph is semi-algebraic. We say that a bipartite graph $G = (A, B, E)$ is *semi-algebraic* of description complexity t in $(\mathbb{R}^{d_1}, \mathbb{R}^{d_2})$, if $A \subset \mathbb{R}^{d_1}$, $B \subset \mathbb{R}^{d_2}$, and the edges are defined by the sign-patterns of t polynomials $f_1, \dots, f_t : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \rightarrow \mathbb{R}$, each of degree at most t . In [22] it is proved that if such a graph G is $K_{s,s}$ -free with $|A| = m$ and $|B| = n$, then

$$|E(G)| = O_{d_1, d_2, t, s, \varepsilon} \left(m^{\frac{d_2(d_1-1)}{d_1 d_2 - 1} + \varepsilon} n^{\frac{d_1(d_2-1)}{d_1 d_2 - 1}} + m + n \right)$$

for every $\varepsilon > 0$. In the special case of interest $d = d_1 = d_2$, this gives

$$|E(G)| = O_{d,t,s,\varepsilon} \left((mn)^{\frac{d}{d+1} + \varepsilon} + m + n \right). \quad (3)$$

This bound is only known to be tight for $d = 2$. It is worth highlighting that this upper bound matches (1) up to the ε error term.

We consider the analogous problem for *algebraic graphs*, that is, graphs that are defined with respect to zero-patterns of polynomials over arbitrary fields. Given two sets of points $\mathcal{P} \subseteq \mathbb{F}^{d_1}$, $\mathcal{Q} \subseteq \mathbb{F}^{d_2}$, a collection of t polynomials $f_1, \dots, f_t : \mathbb{F}^{d_1} \times \mathbb{F}^{d_2} \rightarrow \mathbb{F}$, and a boolean formula Φ , we define the bipartite graph G on the vertex set $\mathcal{P} \cup \mathcal{Q}$, where $x \in \mathcal{P}$ and $y \in \mathcal{Q}$ are adjacent if and only if

$$\Phi([f_1(x, y) = 0], \dots, [f_t(x, y) = 0]) = 1.$$

We say that G is an *algebraic graph of description complexity t* if it can be described by t polynomials f_1, \dots, f_t , each of which has degree at most t .

Theorem 1.8. *Let G be an algebraic graph of description complexity at most t on the vertex set $\mathcal{P} \cup \mathcal{Q}$, where $\mathcal{P} \subseteq \mathbb{F}^{d_1}$, $\mathcal{Q} \subseteq \mathbb{F}^{d_2}$ and $|\mathcal{P}| = m$, $|\mathcal{Q}| = n$. If G is $K_{s,s}$ -free, the number of edges in G is at most*

$$O_{d_1, d_2, t, s}(\min\{m^{1-1/d_1}n, mn^{1-1/d_2}\}).$$

Moreover, this bound is tight, for every d_1, d_2 , there exists t such that if m, n are sufficiently large and $s \geq d_1 + d_2$, the following holds. There exist a field \mathbb{F} and a $K_{s,s}$ -free algebraic graph G of description complexity at most t in $(\mathbb{F}^{d_1}, \mathbb{F}^{d_2})$ with parts of size m, n and at least $\Omega(\min\{m^{1-1/d_1}n, mn^{1-1/d_2}\})$ edges.

Paper organization: In Section 2, we give a short illustration of our techniques discussing point-hyperplane incidences in 3-dimensional spaces, and introduce the key lemma connecting combinatorics and geometry. Then, in Section 3, we prove our general results on point-hyperplane incidences, Theorems 1.2 and Theorem 1.3, along with the constructions showing the tightness of these bounds. In Section 4 we discuss point-variety incidences and show two proofs of Theorem 1.4, while postponing the constructions until Section 6. The unit distance problem is discussed in Section 5, where the proofs of Theorems 1.6 and 1.7 are presented. In Section 6 we discuss the Zarankiewicz problem for algebraic graphs and show Theorems 1.5 and 1.8. We conclude our paper with a few remarks.

2 Forbidden induced patterns and illustration in dimension 3

The key idea used to prove our bounds is that incidence graphs of points and hyperplanes in \mathbb{F}^d avoid a simple family \mathcal{F}_d of graphs as induced subgraphs. The family \mathcal{F}_d includes all bipartite graphs on vertex classes $\{a_1, \dots, a_d\}$ and $\{b_1, \dots, b_d\}$, where $a_i b_j$ is an edge for $i \geq j - 1$, and $a_i b_{i+2}$ is a non-edge for $i = 1, \dots, d - 2$.

An alternative and perhaps easier way to think about the family is through the notion of a *pattern*. We define a *pattern* to be an edge labeling of a complete bipartite graph, in which every edge receives one of the labels 0, 1 or *. Every such pattern Π defines a family of graphs \mathcal{F}_Π on the same vertex set, where a graph F is in the family if all edges of the pattern labeled by 0 do not appear in F , all edges labeled by 1 appear in F , while the edges labeled by * may or may not appear. To simplify the terminology, we say that a graph G contains a pattern Π if it contains an induced copy of some graph in the family \mathcal{F}_Π .

The family \mathcal{F}_d discussed above can be constructed in this way from the following pattern Π_d . Namely, Π_d is a pattern on a balanced bipartite graph on vertices $a_1, \dots, a_d, b_1, \dots, b_d$ where $a_i b_j$ is labeled by 1 for $i \geq j - 1$ and $a_i b_{i+2}$ is labeled by 0, while all other edges are labeled by *. See Figure 1 for an illustration.

The following lemma is the main geometric ingredient that we use, which allows us to reduce the incidence bounds to a problem in extremal graph theory.

Lemma 2.1. *Let $d \geq 3$ be an integer and let \mathcal{P} be a set of distinct points and \mathcal{H} be a set of distinct hyperplanes in \mathbb{F}^d . Then, the incidence graph $G(\mathcal{P}, \mathcal{H})$ does not contain the pattern Π_d .*

Proof. Assume, for the sake of contradiction, that $G = G(\mathcal{P}, \mathcal{H})$ contains Π_d . Since the pattern Π_d is symmetric, we may assume that the vertices a_1, \dots, a_d correspond to points $x_1, \dots, x_d \in \mathcal{P}$, and the vertices b_1, \dots, b_d correspond to hyperplanes $H_1, \dots, H_d \in \mathcal{H}$. We prove by induction that




Figure 1: The pattern Π_5 , where black edges are labeled 1, and the red dashed edges are labeled 0.

$\dim \bigcap_{i=1}^k H_i = d - k$ for all $2 \leq k \leq d$. This suffices to derive the contradiction by observing that $x_{d-1}, x_d \in \bigcap_{i=1}^d H_i$, which is at most a 0-dimensional space, i.e. a single point. This contradicts the assumption that x_1, \dots, x_d are distinct.

To show that $\dim \bigcap_{i=1}^k H_i = d - k$, we begin by observing that $H_1 \cap H_2$ must have dimension $d - 2$, since H_1, H_2 are distinct hyperplanes with non-empty intersection. For any $k \geq 3$, the induction hypothesis implies that $\bigcap_{i=1}^{k-1} H_i$ is an affine space of dimension $d - k + 1$, which is not a subspace of H_k since $x_{k-2} \in \bigcap_{i=1}^{k-1} H_i$ and $x_{k-2} \notin H_k$. Thus, H_k must intersect $\bigcap_{i=1}^{k-1} H_i$ in an affine subspace of dimension $\dim \bigcap_{i=1}^{k-1} H_i - 1 = d - k$. This completes the proof. \square

For a linear algebraic perspective on the pattern Π_d , one might consider the $(d-2) \times (d-2)$ matrix whose rows and columns are indexed by a_1, \dots, a_{d-2} and b_3, \dots, b_d , where the entry corresponding to a_i and b_j equals 1 if the edge $a_i b_j$ is labeled 0 and equals 0 if the edge $a_i b_j$ is labeled 1 (entries corresponding to the edges labeled by $*$ may be arbitrary). This defines an upper-triangular matrix with non-zero diagonal entries, a condition which ensures full rank. A similar family of graphs was also studied recently by Sudakov and Tomon [50] to establish Ramsey theoretic properties of algebraically defined graphs. Using the above lemma, we can just focus on the hereditary family of graphs that avoid the pattern Π_d . Then, we use the dependent random choice technique [24] and combinatorial methods to show that if G is a Π_d -free graph with sufficiently many edges, then G must contain a large complete bipartite graph.

Let us now illustrate how Lemma 2.1 can be combined with the technique of dependent random choice to give a very short proof of Theorem 1.1 in dimension three. Note that when $d = 3$, the forbidden pattern Π_3 is quite simple, since containing the pattern Π_3 is equivalent to containing the bipartite graph $K_{3,3}$ minus an edge as an induced subgraph.

Throughout this paper, we use standard graph theoretic notation. Given a graph G and a vertex $v \in V(G)$, we denote by $\deg v$ the degree of v , and by $N(v)$ the neighbourhood of v . Also, if $v_1, \dots, v_t \in V(G)$, we denote by $N(v_1, \dots, v_t)$ the *common neighbourhood* of v_1, \dots, v_t , i.e. $N(v_1, \dots, v_t) = N(v_1) \cap \dots \cap N(v_t)$.

Lemma 2.2. *Let $s \geq 2$ be an integer, let \mathcal{P} be a set of n points and \mathcal{H} be a set of n hyperplanes in \mathbb{F}^3 . If $G(\mathcal{P}, \mathcal{H})$ does not contain $K_{s,s}$, then $I(\mathcal{P}, \mathcal{H}) \leq 2sn^{3/2}$.*

Proof. Throughout the proof, we denote the bipartition of the incidence graph $G = G(\mathcal{P}, \mathcal{H})$ by $A \cup B$, where A corresponds to the set of points and B to the set of hyperplanes. This will help to separate the graph-theoretic arguments from the geometric ones.

We show that if $G = G(\mathcal{P}, \mathcal{H})$ satisfies $e(G) \geq 2sn^{3/2}$ and does not contain $K_{s,s}$, then it must contain the pattern Π_3 . To do this, we use the method of dependent random choice. More precisely, our goal is to find four vertices $v_1, v_2 \in A, u_1, u_2 \in B$ forming a copy of $K_{2,2}$ such that the common neighbourhoods $N(v_1) \cap N(v_2)$ and $N(u_1) \cap N(u_2)$ have size at least s . Since the graph G does not contain $K_{s,s}$, there will be an edge missing between these two common neighbourhoods, which allows us to find the pattern Π_3 in G .

Let us begin by picking a random two vertex subset $\{v_1, v_2\} \subset A$ chosen from the uniform distribution on all $\binom{n}{2}$ pairs, and consider the common neighbourhood $X = N(v_1) \cap N(v_2)$ in B . The expected size of this common neighbourhood can be lower bounded by counting the number of

three element sets $\{v_1, v_2, u\}$ with $v_1, v_2 \in A, u \in B$ and $v_1u, v_2u \in E(G)$ (such sets are sometimes also called as cherries). If the number of such sets is denoted by T , we have

$$\mathbb{E}[|X|] = \frac{T}{\binom{n}{2}} = \frac{1}{\binom{n}{2}} \sum_{u \in B} \binom{\deg u}{2} \geq \frac{2}{n-1} \left(\frac{1}{n} \sum_{u \in B} \deg u \right) \geq \frac{e(G)^2}{4n^3} \geq s^2,$$

where the first inequality holds by the convexity of the function $\binom{x}{2} = \frac{x(x-1)}{2}$ for $x \geq 1$. On the other hand, let us denote by Y the number of two element sets $\{u_1, u_2\} \subset X$ with common neighbourhood $N(u_1) \cap N(u_2)$ of size less than s . Then the expected value of Y can be bounded as follows. For each two element set $\{u_1, u_2\} \subset B$, we have $u_1, u_2 \in X$ if $v_1, v_2 \in N(u_1) \cap N(u_2)$. Hence, if $|N(u_1) \cap N(u_2)| < s$, then the probability of $u_1, u_2 \in X$ is less than $\binom{s}{2} / \binom{n}{2} < \left(\frac{s}{n}\right)^2$, so

$$\mathbb{E}[Y] \leq \sum_{\{u_1, u_2\} \subset B} \left(\frac{s}{n}\right)^2 \leq \binom{n}{2} \left(\frac{s}{n}\right)^2 \leq \frac{s^2}{2}.$$

Since $\mathbb{E}[|X| - Y] \geq s^2/2$, there exists a pair of vertices v_1, v_2 satisfying $|X| - Y \geq s^2/2$. This means that at least one pair $u_1, u_2 \in X$ has a common neighbourhood of size at least s . Also, since $X \geq s^2/2 \geq s$, the vertices v_1, v_2 have common neighbourhood of size at least s . As G does not contain $K_{s,s}$, there must be vertices $v_3 \in N(u_1) \cap N(u_2), u_3 \in N(v_1) \cap N(v_2)$ for which $u_3v_3 \notin E(G)$. But then $v_1, v_2, v_3, u_1, u_2, u_3$ form the pattern Π_3 , which is a contradiction. This completes the proof. \square

3 Point-hyperplane incidences

3.1 Upper bounds for $K_{s,s}$ -free incidence graphs

In this section, we prove the upper bounds in Theorem 1.2. We prepare the proof with the following lemma about finding the pattern Π_d under some simple conditions. We recall that Π_d is the pattern on a balanced bipartite graph on vertices $a_1, \dots, a_d, b_1, \dots, b_d$ where $a_i b_j$ is labeled by 1 for $i \geq j - 1$ and $a_i b_{i+2}$ is labeled by 0, while all other edges are labeled by $*$.

Lemma 3.1. *Let s and $2 \leq t \leq d$ be positive integers, and let $G = (A, B; E)$ be a $K_{s,s}$ -free bipartite graph. Assume that $v_1, \dots, v_t \in A$ satisfy the following conditions:*

- $N(v_1, v_2) \supseteq N(v_1, v_2, v_3) \supseteq \dots \supseteq N(v_1, \dots, v_t)$,
- *there exists a set $X \subseteq N(v_1, \dots, v_t)$ of size at least $d^{d+1}s$ such that every $(d+1-t)$ -tuple of elements in X has a common neighbourhood of size at least s .*

Then G contains the pattern Π_d .

Proof. The first condition implies that there exist vertices u_1, \dots, u_{t-2} which satisfy

$$u_i \in N(v_1, \dots, v_{i+1}) \setminus N(v_{i+2}).$$

To find vertices u_{t-1}, \dots, u_d and v_{t+1}, \dots, v_d , we define a $(d+2-t)$ -uniform hypergraph F on the set of vertices X . Let \prec be an arbitrary total ordering on X , then the edges of the hypergraph F will be those sets $\{u_{t-1}, \dots, u_d\}$ with $u_{t-1} \prec \dots \prec u_d$ for which there exists an index $i \in \{t-1, \dots, d-1\}$ such that $N(u_i, \dots, u_d) = N(u_{i+1}, \dots, u_d)$. We claim that F has at most

$$s(d+1-t)|X|^{d+1-t}$$

edges. To this end, note that there are $d+1-t$ choices for the index i at which the equality $N(u_i, \dots, u_d) = N(u_{i+1}, \dots, u_d)$ can occur. Furthermore, for each of these choices, there are at most $|X|^{d-i}$ ways to choose $\{u_{i+1}, \dots, u_d\}$. Having fixed these vertices, which by assumption satisfy $|N(u_{i+1}, \dots, u_d)| \geq s$, we have less than s vertices u_i for which $N(u_i) \supseteq N(u_{i+1}, \dots, u_d)$, otherwise

we have a $K_{s,s}$. Finally, there are at most $|X|^{i-(t-1)}$ ways to choose vertices u_{t-1}, \dots, u_{i-1} . Thus, the total number of edges of the hypergraph F is at most

$$|E(F)| < (d+1-t)|X|^{d-i} \cdot s \cdot |X|^{i-(t-1)} = s(d+1-t)|X|^{d+1-t},$$

as claimed.

Observe that the inequality $\binom{|X|}{d+2-t} \geq \left(\frac{|X|}{d+2-t}\right)^{d+2-t} \geq s(d+1-t)|X|^{d+1-t}$ is satisfied by our assumption that $|X| \geq sd^{d+1}$. Hence, there exist $u_{t-1} \prec \dots \prec u_d$ in X such that $\{u_{t-1}, \dots, u_d\}$ is not an edge of F , or equivalently, $N(u_i, \dots, u_d) \subsetneq N(u_{i+1}, \dots, u_d)$ for all indices $t-1 \leq i \leq d-1$. Picking vertices

$$v_{i+2} \in N(u_{i+1}, \dots, u_d) \setminus N(u_i, \dots, u_d)$$

for all $t-1 \leq i \leq d-1$, the vertices $v_1, \dots, v_d, u_1, \dots, u_d$ span a subgraph with pattern Π_d . \square

The following is the main graph-theoretical lemma which allows us to prove the upper bounds on the number of incidences of points and hyperplanes. Let us recall that we have defined $\alpha_t = \frac{t}{d+2-t}$ for $t \in \{2, \dots, d\}$ and $\beta_t = \frac{t}{d+1-t}$ for $t \in \{1, \dots, d\}$.

Lemma 3.2. *For integers $d, s \geq 2$, there exists $C = C(d, s) > 0$ such that the following holds for every sufficiently large m . Let $G = (A, B; E)$ be a bipartite graph, $|A| = m$, $|B| = n$ and $n = m^\alpha$. If G contains no $K_{s,s}$ as a subgraph and does not contain the pattern Π_d , then*

$$e(G) \leq \begin{cases} Cm & \text{if } \alpha \in [0, \beta_1], \\ Cmn^{1-\frac{1}{t}} & \text{if } \alpha \in [\alpha_t, \beta_t] \text{ for some } t \in \{2, \dots, d\}, \\ Cm^{1-\frac{1}{d+2-t}}n & \text{if } \alpha \in [\beta_{t-1}, \alpha_t] \text{ for some } t \in \{2, \dots, d\}, \\ Cn & \text{if } \alpha \in [\beta_d, \infty). \end{cases}$$

Moreover, if $\alpha = 1$, i.e. when $m = n$, the constant $C = 8(s + d^3)$ suffices.

Proof. Let us start by addressing the cases when $\alpha \in [\alpha_t, \beta_t]$ for some $t \in \{2, \dots, d\}$ or $\alpha \in [\beta_d, \infty)$. In the latter case, we set $t := d$. Suppose that $e(G) \geq C(mn^{1-\frac{1}{t}} + n)$ for some large constant C , and observe that $mn^{1-\frac{1}{t}} \geq n$ if $\alpha \in [\alpha_t, \beta_t]$, and $n \geq mn^{1-\frac{1}{t}}$ if $\alpha \in [\beta_d, \infty)$. We show that there exists a t -tuple of vertices $v_1, \dots, v_t \in A$ which satisfies the assumptions of Lemma 3.1, which then implies that G contains the pattern Π_d , contradiction. To do this, we use a variant of the dependent random choice method.

We call an ordered t -tuple of vertices (v_1, \dots, v_t) *bad* if $N(v_1, \dots, v_t) \geq s$ and $N(v_1, \dots, v_i) = N(v_1, \dots, v_{i+1})$ for some $i \in \{1, \dots, t-1\}$, and *good* otherwise. We choose a random good t -tuple (v_1, \dots, v_t) with the following sampling procedure. We sample the vertices v_i in order, starting from v_1 which is a uniformly random element of A . At step i , having chosen vertices v_1, \dots, v_{i-1} , we have two cases.

Case 1. $|N(v_1, \dots, v_{i-1})| < s$. In this case, any t -tuple containing v_1, \dots, v_{i-1} is good, and then all subsequent vertices v_j for $j \geq i$ can be sampled uniformly at random from $A \setminus \{v_1, \dots, v_{j-1}\}$.

Case 2. $|N(v_1, \dots, v_{i-1})| \geq s$. Define the set $S_i = \{v \in A : N(v) \supseteq N(v_1, \dots, v_{i-1})\}$. Since $G(P, H)$ does not contain a copy of $K_{s,s}$, there are less than s vertices in the set S_i . Sample v_i uniformly at random from $A \setminus (S_i \cup \{v_1, \dots, v_{i-1}\})$.

Let us denote by X the size of the common neighbourhood of v_1, \dots, v_t , and let Y be the number of $(d+1-t)$ -tuples of vertices $u_t, \dots, u_d \in N(v_1, \dots, v_t)$ which has less than s common neighbours. If we show that $\mathbb{E}[X - Y] \geq d^{d+1}s$, then one can pick a good t -tuple v_1, \dots, v_t satisfying $X - Y \geq d^{d+1}s$ and delete one vertex from every $(d+1-t)$ -tuple in $N(v_1, \dots, v_t)$ which has less than s common neighbours. Then, one is left with the t -tuple v_1, \dots, v_t and a subset of their common neighbourhood which satisfy all conditions of Lemma 3.1.

Thus, the main objective is to show that $\mathbb{E}[X - Y] \geq d^{d+1}s$. By linearity of expectation, we have $\mathbb{E}[X] = \sum_{u \in B} \mathbb{P}[u \in N(v_1, \dots, v_t)]$. Let us fix a vertex $u \in B$ with $\deg u \geq d + s$, and estimate $\mathbb{P}[u \in N(v_1, \dots, v_t)]$. We have

$$\begin{aligned} \mathbb{P}[u \in N(v_1, \dots, v_t)] &= \mathbb{P}[v_1, \dots, v_t \in N(u)] = \prod_{i=1}^t \mathbb{P}[v_i \in N(u) \mid v_1, \dots, v_{i-1} \in N(u)] \\ &\geq \prod_{i=1}^t \frac{\deg u - s - i + 1}{m - i + 1} > \left(\frac{\deg u - s - d}{m} \right)^t. \end{aligned}$$

Let $B_0 \subseteq B$ be the set of vertices u satisfying $\deg u \geq 2(d + s)$. Then, the number of edges incident to B_0 is $\sum_{u \in B_0} \deg u \geq C(mn^{1-1/t} + n) - 2(d + s)n \geq e(G)/2$ for $C \geq 4(d + s)$. Therefore,

$$\begin{aligned} \mathbb{E}[X] &\geq \sum_{u \in B_0} \left(\frac{\deg u - s - d}{m} \right)^t \geq \sum_{u \in B_0} \left(\frac{\deg u}{2m} \right)^t \\ &\geq |B_0| \left(\frac{\sum_{u \in B_0} \deg u}{2|B_0|m} \right)^t \geq n \left(\frac{e(G)}{4mn} \right)^t \geq \left(\frac{e(G)}{4mn^{1-1/t}} \right)^t, \end{aligned}$$

where the third inequality holds by convexity. On the other hand, the probability that a given $(d + 1 - t)$ -tuple of vertices $u_t, \dots, u_d \in B$ which has less than s common neighbours lies in $N(v_1, \dots, v_t)$ can be bounded by

$$\mathbb{P}[u_t, \dots, u_d \in N(v_1, \dots, v_t)] = \mathbb{P}[v_1, \dots, v_t \in N(u_t, \dots, u_d)] \leq \left(\frac{|N(u_t, \dots, u_d)|}{m - s - d} \right)^t \leq (2s)^t m^{-t},$$

assuming m is sufficiently large with respect to s and d . Thus, the expectation of Y can be bounded as

$$\mathbb{E}[Y] \leq \sum_{\substack{u_t, \dots, u_d \in B \\ |N(u_t, \dots, u_d)| < s}} \left(\frac{2s}{m} \right)^t \leq n^{d+1-t} \left(\frac{2s}{m} \right)^t.$$

When $\alpha \leq d = \beta_d$, we have $\alpha \leq \beta_t = \frac{t}{d+1-t}$. Hence, for a sufficiently large value of C ,

$$\mathbb{E}[X - Y] \geq \left(\frac{e(G)}{4mn^{1-1/t}} \right)^t - n^{d+1-t} \left(\frac{2s}{m} \right)^t \geq (C/4)^t - (2s)^t m^{\alpha(d+1-t)-t} \geq (C/4)^t - (2s)^t \geq d^{d+1}s.$$

On the other hand, when $\alpha \geq d = \beta_d$, we have

$$\mathbb{E}[X - Y] \geq \left(\frac{Cn}{4mn^{1-1/t}} \right)^t - m \left(\frac{2s}{m} \right)^t = ((C/4)^d - (2s)^d) \frac{n}{m^d} \geq (C/4)^d - (2s)^d \geq d^{d+1}s,$$

for a sufficiently large value of C . Thus, we indeed have $\mathbb{E}[X - Y] \geq (C/4)^t - (2s)^t \geq d^{d+1}s$ in all cases and therefore this completes the first case of the proof.

Before we consider the rest of the cases, let us justify the assertion that when $\alpha = 1$ one can take $C = 8(s + d^3)$. Namely, if $\alpha = 1$ we have $\alpha_t \leq \alpha \leq \beta_t$ for $t = \lceil \frac{d+1}{2} \rceil$. Then, for the proof to go through, one needs to verify that the inequality $\mathbb{E}[X - Y] \geq (C/4)^t - (2s)^t \geq sd^{d+1}$ holds with the proposed value of C , which is a consequence of a simple calculation.

Observe that the rest of the cases already follow by symmetry. Since $1/\alpha_t = \alpha_{d+2-t}$ and $1/\beta_{t-1} = \beta_{d+2-t}$, we have that if $\alpha \in [\beta_{t-1}, \alpha_t]$ for some $t \in \{2, \dots, d\}$, then $1/\alpha \in [\alpha_{d+2-t}, \beta_{d+2-t}]$. Also, if $\alpha \in [0, \beta_1]$, then $1/\alpha \in [\beta_d, \infty)$. Thus, if we reverse the roles of m and n , together with the observation $m = n^{1/\alpha}$, we find ourselves in the setting of the cases already proved. This completes the proof of the theorem. \square

Now we are ready to prove the upper bounds of Theorem 1.1 and Theorem 1.2, which we restate as follows.

Theorem 3.3. *Let $d, s \geq 2$ be integers, then there exists $C = C(d, s) > 0$ such that the following holds for every $\alpha > 0$ and field \mathbb{F} . Let \mathcal{P} be a set of m points and \mathcal{H} be a set of $n = m^\alpha$ hyperplanes in \mathbb{F}^d such that the incidence graph of $(\mathcal{P}, \mathcal{H})$ is $K_{s,s}$ -free. Then*

$$I(\mathcal{P}, \mathcal{H}) \leq \begin{cases} Cm & \text{if } \alpha \in [0, \beta_1], \\ Cmn^{1-\frac{1}{t}} & \text{if } \alpha \in [\alpha_t, \beta_t] \text{ for some } t \in \{2, \dots, d\}, \\ Cm^{1-\frac{1}{d+2-t}}n & \text{if } \alpha \in [\beta_{t-1}, \alpha_t] \text{ for some } t \in \{2, \dots, d\}, \\ Cn & \text{if } \alpha \in [\beta_d, \infty). \end{cases}$$

Moreover, if $\alpha = 1$, the implied constant can be taken to be $C = 8(s + d^3)$.

Proof. By Lemma 2.1, the graph $G(\mathcal{P}, \mathcal{H})$ does not contain the pattern Π_d . Since the conclusion of Lemma 3.2 does not hold, one of its assumptions must fail. Thus, we conclude that $G(\mathcal{P}, \mathcal{H})$ either contains $K_{s,s}$ as a subgraph or $I(\mathcal{P}, \mathcal{H})$ satisfies the required upper bound. \square

3.2 Large complete bipartite subgraph of dense incidence graphs

The main goal of this section is to prove our lower bounds on $\text{rs}_d(m, n, I)$ and $\text{rs}_d(n, m, I)$ for $m \leq n$, that is, to prove Theorem 1.3. Recall that $\text{rs}_d(n, m, I) = \min \text{rs}(\mathcal{P}, \mathcal{H})$, where the minimum is taken over all fields \mathbb{F} , all sets $\mathcal{P} \subset \mathbb{F}^d$ of at most m points, all sets \mathcal{H} of at most n hyperplanes in \mathbb{F}^d , that satisfy $I(\mathcal{P}, \mathcal{H}) \geq I$. As observed before, the lower bounds are only interesting in the regime $\varepsilon > 100 \max\{m^{-\frac{1}{d-1}}, n^{-\frac{1}{d}}\}$, otherwise taking the maximum degree vertex with its neighborhood gives a complete bipartite graph of the required size. Let us restate our theorem in this case.

Theorem 3.4. *Let $m \leq n$ be integers, $\varepsilon > 0$, and $I = \varepsilon mn$. If $\varepsilon > 100 \max\{m^{-\frac{1}{d-1}}, n^{-\frac{1}{d}}\}$, then*

$$\text{rs}_d(m, n, I), \text{rs}_d(n, m, I) \geq \left(\frac{\varepsilon}{100}\right)^{d-1} mn.$$

Let us prepare the proof with some definitions. For a bipartite graph G , we denote by $\text{rs}(G)$ the number of edges in the largest complete bipartite subgraph of $G = (A, B; E)$. Also, we say that a t -tuple of vertices $(v_1, \dots, v_t) \in A^t$ is a *good t -tuple* if $N(v_1, v_2) \supseteq N(v_1, v_2, v_3) \supseteq \dots \supseteq N(v_1, \dots, v_t)$ and $|N(v_1, \dots, v_t)| \geq 2$. The reason this definition is useful is the following lemma, which allows us to embed the pattern Π_d using good d -tuples. One can think of Lemma 3.5 as a simpler version of Lemma 3.1 from Section 3.1.

Lemma 3.5. *Suppose that G contains a good d -tuple (v_1, \dots, v_d) . Then G contains the pattern Π_d .*

Proof. Picking $u_i \in N(v_1, \dots, v_{i+1}) \setminus N(v_1, \dots, v_{i+2})$ for $i \leq d-2$ and $u_{d-1}, u_d \in N(v_1, \dots, v_d)$, the vertices $v_1, \dots, v_d, u_1, \dots, u_d$ induce a subgraph with pattern Π_d . \square

Let us reformulate Theorem 3.4 in terms of good d -tuples.

Lemma 3.6. *Let $d \geq 2$ be a positive integer and let $G = (A, B; E)$ be a bipartite graph with at least εmn edges, where $|A| = m$, $|B| = n$ and $n \geq m$. If $\varepsilon \geq 100 \max\{m^{-\frac{1}{d-1}}, n^{-\frac{1}{d}}\}$, then either A contains a good d -tuple or*

$$\text{rs}(G) \geq \left(\frac{\varepsilon}{100}\right)^{d-1} mn.$$

We prove this lemma after some preparations. The following lemma shows how to build good t -tuples under the assumption that G does not contain large complete bipartite subgraphs.

Lemma 3.7. *Let $\gamma, \varepsilon > 0$, and let $G = (A, B; E)$ be a bipartite graph such that $|A| = m$, $|B| = n$, and every vertex $u \in B$ has degree at least $\deg u \geq \frac{\varepsilon}{3}n$. If $\text{rs}(G) < \frac{\varepsilon}{6}\gamma \cdot mn$, then for any set $X \subseteq B$ of size at least γn , there exists a vertex $v \in A$ for which $|N(v) \cap X| \geq \frac{\varepsilon}{6}|X|$ and $N(v) \cap X \neq X$.*

Proof. Let $X \subseteq B$ be a subset of size at least γn . The idea of this proof is to double count the edges incident to X . Let V_1 denote the set of vertices $v \in A$ for which $X \subseteq N(v)$ and let V_2 denote the set of vertices $v \in A$ with $|N(v) \cap X| < \frac{\varepsilon}{6}|X|$. To show the lemma, it suffices to find a vertex $v \in A$ belonging to neither V_1 nor V_2 .

Suppose, for the sake of contradiction, that $V_1 \cup V_2 = A$. Then, $e(A, X) = e(V_1, X) + e(V_2, X)$. Since the degree of every vertex in B is at least $\frac{\varepsilon m}{3}$, we have $e(A, X) \geq \frac{\varepsilon}{3}m|X|$. On the other hand, since the graph induced on $V_1 \cup X$ is complete bipartite, we have $e(V_1, X) \leq \text{rs}(G) \leq \frac{\varepsilon}{6}\gamma \cdot mn$. Also, by the definition of V_2 , we have $e(V_2, X) = \sum_{v \in V_2} |N(v) \cap X| < |V_2| \cdot \frac{\varepsilon}{6}|X| \leq \frac{\varepsilon}{6}m|X|$. Hence,

$$\frac{\varepsilon}{6}\gamma \cdot mn + \frac{\varepsilon}{6}m|X| > \frac{\varepsilon}{3}m|X|.$$

Rearranging this inequality gives $|X| < \gamma n$, which is a contradiction to our initial assumption. \square

Now, we explain how to build good $(d-1)$ -tuples in the graph G .

Lemma 3.8. *Let $d \geq 3$ be a positive integer and let $G = (A, B; E)$ be a bipartite graph with εmn edges, where $|A| = m$, $|B| = n$. If $\text{rs}(G) \leq (\frac{\varepsilon}{6})^{d-1}mn$, then G contains a good $(d-1)$ -tuple $(v_1, \dots, v_{d-1}) \in A^{d-1}$ with common neighbourhood of size at least $|N(v_1, \dots, v_{d-1})| \geq 2(\frac{\varepsilon}{6})^{d-1}n$.*

Proof. For $t = 2, \dots, d-1$, we prove that there is a good t -tuple $(v_1, \dots, v_t) \in A^t$ with $|N(v_1, \dots, v_t)| \geq 2(\frac{\varepsilon}{6})^t n$. We prove this by induction on t . For $t = 2$, we simply need to find a pair of vertices $v_1, v_2 \in A$ whose common neighbourhood has size at least $2(\frac{\varepsilon}{6})^2 n$. Since the average degree of vertices in B is at least εm , the average size of the common neighbourhood of two vertices of A is $n \binom{\varepsilon m}{2} / \binom{m}{2} \geq \frac{\varepsilon^2}{2}n$, which is sufficient.

For $t > 2$, the induction hypothesis implies that G contains a good $(t-1)$ -tuple $(v_1, \dots, v_{t-1}) \in A^{t-1}$ with the common neighbourhood of size $|N(v_1, \dots, v_{t-1})| \geq 2(\frac{\varepsilon}{6})^{t-1}n$. Let $\gamma = (\frac{\varepsilon}{6})^{t-1}$, then $\text{rs}(G) \leq (\frac{\varepsilon}{6})^{d-1}mn \leq \frac{\varepsilon}{6}\gamma \cdot mn$ and $|X| \geq \gamma n$ are satisfied, so we can apply Lemma 3.7 to $X = N(v_1, \dots, v_{t-1})$. We get that there exists a vertex $v_t \in A$ for which $|N(v_1, \dots, v_t)| \geq \frac{\varepsilon}{6}|X| \geq (\frac{\varepsilon}{6})^t n$ and $N(v_1, \dots, v_{t-1}) \not\supseteq N(v_1, \dots, v_t)$. We conclude that (v_1, \dots, v_t) is a good t -tuple with sufficiently large common neighbourhood, completing the proof. \square

Now, we are ready to prove Lemma 3.6.

Proof of the Lemma 3.6. Assume that the statement is false, and let G be a counterexample with d minimal, and among those $\min\{m, n\}$ is also minimal. That is, G has at least εmn edges, G contains no good d -tuple and $\text{rs}(G) < (\frac{\varepsilon}{100})^{d-1}mn$. Clearly, we have $\min\{m, n\} \geq 2$. Also, if $d = 2$, finding a good d -tuple corresponds to finding a pair of vertices of A with at least two common neighbours, i.e. finding a cycle of length 4 in G . Since $|E(G)| \geq \varepsilon mn \geq 100 \max\{n^{-1/2}, m^{-1}\}mn \geq mn^{1/2} + n$, a simple application of the classical Kővári-Sós-Turán theorem shows that G must have a cycle of length 4, which suffices. Hence, we may assume that $d \geq 3$.

First, we show that $\min_{v \in A} \deg v \geq (1 - \frac{2}{d})\varepsilon|B|$ and $\min_{v \in B} \deg v \geq (1 - \frac{2}{d})\varepsilon|A|$. Since A and B are symmetric, we only discuss the case when G has a vertex $v \in A$ with degree $\deg v \leq (1 - \frac{2}{d})\varepsilon|B|$. Define the graph $G' = G - \{v\}$, which has parts of size $m' = m - 1$ and $n' = n$. Furthermore, G' has $\varepsilon' m' n'$ edges, where $\varepsilon' m' n' \geq \varepsilon mn - (1 - \frac{2}{d})\varepsilon n$. Thus, as G' does not contain a good d -tuple and by the minimality of G , we conclude

$$\begin{aligned} \text{rs}(G) &\geq \text{rs}(G') \geq \left(\frac{\varepsilon'}{100}\right)^{d-1} m' n' \geq \left(\frac{\varepsilon mn - (1 - \frac{2}{d})\varepsilon n}{100(m-1)n}\right)^{d-1} (m-1)n \\ &= \left(\left(\frac{\varepsilon}{100}\right)^{d-1} mn\right) \cdot \frac{m-1}{m} \left(\frac{m - (1 - \frac{2}{d})}{m-1}\right)^{d-1} \geq \left(\frac{\varepsilon}{100}\right)^{d-1} mn. \end{aligned}$$

The last inequality is true due to simple calculations. We should also verify that G' satisfies the conditions of Lemma 3.6 by showing that $\varepsilon'^{d-1} m' \geq 100^{d-1}$ and $\varepsilon'^d n' \geq 100^d$. But this also follows,




Figure 2: An illustration for the proof of Lemma 3.6

since we have already shown that $\varepsilon'^{d-1}m'n' \geq \varepsilon^{d-1}mn$, which implies both $\varepsilon'^{d-1}m' \geq \varepsilon^{d-1}m \geq 100^{d-1}$ and $\varepsilon'^d n' \geq \varepsilon^d n \geq 100^d$ (since $\varepsilon'^{d-1}m'n' \geq \varepsilon^{d-1}mn$ also implies $\varepsilon' \geq \varepsilon$). We conclude that $\text{rs}(G) \geq (\varepsilon/100)^{d-1}mn$, a contradiction.

In what follows, we assume that G contains no vertex $v \in A$ of degree less than $(1 - \frac{2}{d})\varepsilon n$ and no vertex $v \in B$ of degree less than $(1 - \frac{2}{d})\varepsilon m$. We choose a good $(d-1)$ -tuple (v_1, \dots, v_{d-1}) with the maximal size of the common neighbourhood, which we denote by $X = N(v_1, \dots, v_{d-1})$. By Lemma 3.8, there must be a good $(d-1)$ -tuple with $|X| \geq (\frac{\varepsilon}{6})^{d-1}|B|$. If A does not contain a good d -tuple in G , then any vertex $v \in A$ either has $X \subseteq N(v)$ or $|N(v) \cap X| = 1$. Lemma 3.7 applied with $\gamma = (\frac{\varepsilon}{6})^{d-2}$ shows that $|X| \leq (\frac{\varepsilon}{6})^{d-2}|B|$. We conclude that one can partition A into two sets, V_1 and V_2 , where $V_1 = \{v \in A : X \subseteq N(v)\}$ and $V_2 = \{v \in A : |X \cap N(v)| \leq 1\}$.

Let us consider the edges incident to X . By our assumption on the minimum degree of vertices of X , we have $e(X, A) \geq |X| \cdot (1 - \frac{2}{d})\varepsilon m \geq \frac{\varepsilon}{3}m|X|$. On the other hand,

$$\frac{\varepsilon}{3}m|X| \leq e(X, A) = e(X, V_1) + e(X, V_2) \leq |X||V_1| + |V_2| \leq |X||V_1| + m.$$

Since $|X| \geq (\frac{\varepsilon}{6})^{d-1}n$ and $(\frac{\varepsilon}{100})^d n \geq 1$, we have $|X| \geq \frac{12}{\varepsilon}$. Thus, $\frac{\varepsilon}{12}m|X| \geq m$ and so $\frac{\varepsilon}{4}m|X| \leq |X||V_1|$, showing that $|V_1| \geq \frac{\varepsilon}{4}m$.

Let us now consider the bipartite graph $G' = G[V_1 \cup (B \setminus X)]$. The number of edges of this graph is at least

$$\begin{aligned} e(G') &= e(V_1, B) - e(V_1, X) \geq |V_1| \cdot \left(1 - \frac{2}{d}\right)\varepsilon|B| - |V_1||X| \\ &\geq |V_1| \left(\frac{d-2}{d}\varepsilon|B| - \left(\frac{\varepsilon}{6}\right)^{d-2}|B| \right) \geq \left(1 - \frac{5/2}{d}\right)\varepsilon|B||V_1|, \end{aligned}$$

since $6^{d-2} \geq 2d$ for all $d \geq 3$. Thus, the sizes of the vertex classes of G' are $m' = |V_1| \geq \frac{\varepsilon m}{4}$ and $n' = |B \setminus X| \geq n - (\frac{\varepsilon}{6})^{d-2}n \geq \frac{5}{6}n$. Finally, the above computation shows that $e(G') = \varepsilon'm'n'$, where $\varepsilon' \geq (1 - \frac{5/2}{d})\varepsilon$.

Next, we argue that G' does not have a good $(d-1)$ -tuple. Suppose for contradiction that G' contains a good $(d-1)$ -tuple (u_1, \dots, u_{d-1}) . The vertices (u_1, \dots, u_{d-1}) form a good $(d-1)$ -tuple in the graph G as well, and their common neighbourhood contains both X and at least two vertices from $B \setminus X$ (since this is a good $(d-1)$ -tuple in G'). But this implies u_1, \dots, u_{d-1} is a $(d-1)$ -tuple with a larger common neighbourhood in G than v_1, \dots, v_{d-1} , contradicting the maximality. Thus, we conclude that G' cannot contain a good $(d-1)$ -tuple.

Furthermore, we show that $\varepsilon'^{d-2}m' \geq 100^{d-2}$ and $\varepsilon'^{d-1}n' \geq 100^{d-1}$. This is a consequence of an easy computation since

$$\varepsilon'^{d-2}m' \geq \left(1 - \frac{5/2}{d}\right)^{d-2} \varepsilon^{d-2} \frac{\varepsilon m}{4} \geq \frac{e^{-5/2}}{4} 100^{d-1} \geq 100^{d-2}.$$

In the above calculation, we used that $\varepsilon^{d-1}n \geq 100^{d-1}$ and that $\left(1 - \frac{5/2}{d}\right)^{d-2} \geq e^{-5/2}$, which holds since $\left(1 - \frac{5/2}{d}\right)^{d-2}$ is a decreasing function for $d \geq 3$ with the limit $e^{-5/2}$ when $d \rightarrow \infty$. A similar computation shows that $\varepsilon'^{d-1}n' \geq \left(1 - \frac{5/2}{d}\right)^{d-1} \varepsilon^{d-1} \cdot \frac{5}{6}n \geq 100^d \varepsilon^{-1} \left(1 - \frac{5/2}{3}\right)^2 \cdot \frac{5}{6} \geq 100^{d-1}$, where we have used that $\left(1 - \frac{5/2}{d}\right)^{d-1}$ is an increasing function of d , when $d \geq 3$. As G is a counterexample with d minimal, G' is not a counterexample, so

$$\text{rs}(G') \geq \left(1 - \frac{5/2}{d}\right)^{d-2} \left(\frac{\varepsilon}{100}\right)^{d-2} |V_1| |B \setminus X| \geq e^{-5/2} \frac{\varepsilon^{d-1}}{4 \cdot 100^{d-2}} m \cdot \frac{5}{6} n \geq \frac{\varepsilon^{d-1}}{100^{d-1}} mn.$$

Hence, we have $\text{rs}(G) \geq \text{rs}(G') \geq \left(\frac{\varepsilon}{100}\right)^{d-1} mn$, which is a contradiction. \square

Proof of the lower bounds in Theorem 3.4. Proving the lower bounds on $\text{rs}_d(m, n, I)$ and $\text{rs}_d(n, m, I)$ is completely symmetric, so we just discuss $\text{rs}_d(m, n, I)$. If $\varepsilon > 100 \max\{m^{-\frac{1}{d-1}}, n^{-\frac{1}{d}}\}$, let us consider any sets of m points and n hyperplanes, say \mathcal{P} and \mathcal{H} . The incidence graph $G = G(\mathcal{P}, \mathcal{H})$ avoids the pattern Π_d by Lemma 2.1 and so G contains no good d -tuples. Since G has at least εmn edges, Lemma 3.6 implies $\text{rs}(G) \geq \left(\frac{\varepsilon}{100}\right)^{d-1} mn$. \square

3.3 Constructions $K_{s,s}$ -free incidence graphs

The main goal of this section is to provide examples which show that the bounds given in Theorem 1.2 are tight. The main building block of all constructions in this section are *subspace evasive sets*. Throughout this section, we consider only finite fields \mathbb{F} .

Definition 3.9. A set of points $S \subset \mathbb{F}^d$ is (k, s) -*subspace evasive* if every k -dimensional affine subspace of \mathbb{F}^d contains less than s elements of S .

It is clear that the size of a (k, s) -subspace evasive set is at most $O_s(|\mathbb{F}|^{d-k})$, as one can partition \mathbb{F}^d into $|\mathbb{F}|^{d-k}$ affine subspaces of dimension k . It was proved by Dvir and Lovett [16] that for sufficiently large s , this trivial bound is optimal. Sudakov and Tomon [51] gave an alternative proof using the random algebraic method.

Lemma 3.10 ([16]). *For integers $1 \leq k \leq d$, there exists $s \leq d^d$ such that for every large enough finite field \mathbb{F} , the space \mathbb{F}^d contains a (k, s) -subspace evasive set of size at least $|\mathbb{F}|^{d-k}$.*

To be more precise, Dvir and Lovett [16] construct a (k, s) -subspace evasive set of size at least $\frac{1}{3}|\mathbb{F}|^{d-k}$. However, taking the union of four random translates of such a set results in a $(k, 4s)$ -subspace evasive set of size at least $|\mathbb{F}|^{d-k}$ with high probability. Removing the factor $1/3$ makes our calculations a bit nicer, but other than that has no real effect.

We prove the following equivalent formulation of the upper bounds in Theorem 1.2.

Theorem 3.11. *For every dimension d , there exist $s = s(d)$ and $c = c(d) > 0$ such that the following holds. For every $\alpha > 0$ and sufficiently large integer m , there exists a field \mathbb{F} , a set of m points \mathcal{P} and a set of $n = \lfloor m^\alpha \rfloor$ hyperplanes \mathcal{H} in \mathbb{F}^d such that the incidence graph of $(\mathcal{P}, \mathcal{H})$ is $K_{s,s}$ -free and*

$$I(\mathcal{P}, \mathcal{H}) \geq \begin{cases} m & \text{if } \alpha \in (0, \beta_1], \\ cmn^{1-\frac{1}{t}} & \text{if } \alpha \in [\alpha_t, \beta_t] \text{ for some } t \in \{2, \dots, d\}, \\ cm^{1-\frac{1}{d+2-t}}n & \text{if } \alpha \in [\beta_{t-1}, \alpha_t] \text{ for some } t \in \{2, \dots, d\}, \\ n & \text{if } \alpha \in [\beta_d, \infty). \end{cases}$$

The main building block of the proof of this theorem is the following lemma.

Lemma 3.12. *Let $2 \leq t \leq d$ be integers, then there exists an integer $s = s(d) > 0$ such that for every prime p the following holds. There exists a set of points \mathcal{P} and a set of hyperplanes \mathcal{H} in \mathbb{F}_p^d such that $G(\mathcal{P}, \mathcal{H})$ does not contain $K_{s,s}$, $|\mathcal{P}| \geq p^t$, $|\mathcal{H}| \geq p^{d-t+2}/s$, and*

$$I(\mathcal{P}, \mathcal{H}) = \frac{|\mathcal{P}||\mathcal{H}|}{p} \geq \Omega_d \left((|\mathcal{P}||\mathcal{H}|)^{1 - \frac{1}{d+2}} \right).$$

Proof. By Lemma 3.10, there exist a constant $s = s(d)$ and sets $\mathcal{P}, \mathcal{N}_0 \subset \mathbb{F}_p^d$ such that $|\mathcal{P}| \geq p^t$, $|\mathcal{N}_0| \geq p^{d-t+1}$, \mathcal{P} is $(d-t, s)$ -subspace evasive and \mathcal{N}_0 is $(t-1, s)$ -subspace evasive. In other words, any s points of \mathcal{P} span an affine subspace of dimension at least $d-t+1$, and any s points of \mathcal{N}_0 span an affine subspace of dimension at least t .

In particular, no line through the origin contains more than s points of \mathcal{N}_0 . Thus, by keeping at most one point in each line through the origin, we get a subset \mathcal{N} of \mathcal{N}_0 of size at least $\geq \frac{1}{s}|\mathcal{N}_0|$.

Let \mathcal{H} be the set of all hyperplanes whose normal vectors are from \mathcal{N} , i.e. all hyperplanes of the form $\langle w, x \rangle = b$ for $w \in \mathcal{N}, b \in \mathbb{F}_p$. Since no line through the origin contains more than one point of \mathcal{N} , we observe that every pair $(w, b) \in \mathcal{N} \times \mathbb{F}_p$ yields a different hyperplane. Thus, the size of \mathcal{H} is $|\mathcal{H}| = p|\mathcal{N}| \geq p^{d-t+2}/s$.

To count the incidences between \mathcal{P} and \mathcal{H} , we observe that for every $w \in \mathcal{N}$ and every $x \in \mathcal{P}$, there is a unique $b \in \mathbb{F}_p$ such that $\langle w, x \rangle = b$. This means that out of p hyperplanes of \mathcal{H} with normal vector $w \in \mathcal{N}$, exactly one is incident to x . Hence, the number of incidences between \mathcal{P} and \mathcal{H} is $I(\mathcal{P}, \mathcal{H}) = |\mathcal{P}||\mathcal{N}| = \frac{1}{p}|\mathcal{P}||\mathcal{H}|$ and so $I(\mathcal{P}, \mathcal{H}) \geq \Omega_d \left((|\mathcal{P}||\mathcal{H}|)^{1 - \frac{1}{d+2}} \right)$.

Finally, we argue that the incidence graph $G(\mathcal{P}, \mathcal{H})$ is $K_{s,s}$ -free. Indeed, any s points of \mathcal{P} span an affine subspace of dimension at least $d-t+1$. On the other hand, the intersection of s hyperplanes is nonempty only when no two of them are parallel, meaning that their normal vectors are distinct. If these s hyperplanes intersect in a $(d-t+1)$ -dimensional affine space, this means that the s normal vectors lie in a $(t-1)$ -dimensional linear space orthogonal to it. But this is not possible since \mathcal{N} is $(t-1, s)$ -subspace evasive. Hence, the intersection of s hyperplanes can be at most $(d-t)$ -dimensional. As $d-t < d-t+1$, this shows that $G(\mathcal{P}, \mathcal{H})$ is $K_{s,s}$ -free. \square

Observe that Lemma 3.12 implies the existence of \mathcal{P}, \mathcal{H} with $|\mathcal{P}| = m_0 = \Omega_d(p^t), |\mathcal{H}| = n_0 = \Omega_d(p^{d+2-t})$ and $I(\mathcal{P}, \mathcal{H}) = \Omega_d((m_0 n_0)^{1 - \frac{1}{d+2}})$. This only covers a sparse set of potential values of m and n . However, by using an approximation and a subsampling trick, one can cover all possible values of m and n .

Proof of Theorem 3.11. Let us first discuss the boundary cases, i.e. when $\alpha \in (0, \beta_1]$ or $\alpha \in [\beta_d, \infty)$. In this case, we show that there exist sets \mathcal{P} and \mathcal{H} with $|\mathcal{P}| = m, |\mathcal{H}| = n = m^\alpha, I(\mathcal{P}, \mathcal{H}) \geq \max\{m, n\}$ and such that $G(\mathcal{P}, \mathcal{H})$ does not contain $K_{2,2}$. In the case $\alpha \in [0, \beta_1]$, it suffices to choose $p \geq \max\{m, n\}$ and take \mathcal{H} to be n parallel hyperplanes in \mathbb{F}_p^d and \mathcal{P} to be m points on one of the hyperplanes. On the other hand, when $\alpha \in [\beta_d, \infty)$, let \mathcal{P} be m collinear points and \mathcal{H} be a set of n hyperplanes containing exactly one of these points.

Let us now present the main argument. The idea of the proof is to pick a configuration $(\mathcal{P}_0, \mathcal{H}_0)$ given by Lemma 3.12 and randomly sample a proportion of it to obtain the required values of m, n .

First, consider the case when $\alpha_t \leq \alpha \leq \beta_t$ for some $t \in \{2, \dots, d\}$. Let p be the smallest prime larger than $2(sn)^{1/t}$. By Bertrand's postulate, we have $p \leq 4(sn)^{1/t}$. By applying Lemma 3.12 with $d+2-t$ instead of t , we obtain a set of points \mathcal{P}_0 and a set of hyperplanes \mathcal{H}_0 in \mathbb{F}_p^d with $|\mathcal{P}_0| = m_0 \geq p^{d+2-t}, |\mathcal{H}_0| = n_0 \geq p^t/s$ and $I(\mathcal{P}_0, \mathcal{H}_0) = \frac{1}{p}|\mathcal{P}_0||\mathcal{H}_0|$. Moreover, $G(\mathcal{P}_0, \mathcal{H}_0)$ does not contain $K_{s,s}$ as a subgraph.

Let \mathcal{P} and \mathcal{H} be random subsets of $\mathcal{P}_0, \mathcal{H}_0$ containing m and n elements, respectively. Observe that $n_0 \geq \frac{1}{s}p^t \geq n$ and $m_0 \geq p^{d+2-t} \geq m^{\frac{d+2-t}{t}} \geq n^{1/\alpha} = m$, so it indeed makes sense to consider such subsets of $\mathcal{P}_0, \mathcal{H}_0$. The expected number of incidences between \mathcal{P} and \mathcal{H} is simply $\mathbb{E}[I(\mathcal{P}, \mathcal{H})] =$

$\frac{m}{m_0} \cdot \frac{n}{n_0} I(\mathcal{P}_0, \mathcal{H}_0) \geq \frac{1}{p} mn$. Since $p \leq 4(sn)^{-1/t}$, picking a pair $(\mathcal{P}, \mathcal{H})$ with at least as many incidences as the expectation, we get

$$I(\mathcal{P}, \mathcal{H}) \geq \frac{1}{p} mn \geq \frac{1}{4s^{1/t}} n^{-\frac{1}{t}} \cdot mn.$$

In the case $\beta_{t-1} \leq \alpha \leq \alpha_t$, we proceed similarly. We pick p to be the smallest prime larger than $2s^{\frac{1}{t}} m^{\frac{1}{d+2-t}}$, then $p < 4s^{\frac{1}{t}} m^{\frac{1}{d+2-t}}$. We apply Lemma 3.12 with $d+2-t$ instead of t to find a set of points \mathcal{P}_0 and a set of hyperplanes \mathcal{H}_0 in \mathbb{F}_p^d with $|\mathcal{P}_0| = m_0 \geq p^{d+2-t}$, $|\mathcal{H}_0| = n_0 \geq p^t/s$ and $I(\mathcal{P}_0, \mathcal{H}_0) = \frac{1}{p} |\mathcal{P}_0| |\mathcal{H}_0|$. By an almost identical computation as before one can check that $m_0 \geq m$ and $n_0 \geq m^{\frac{1}{\alpha}} \geq n$. Hence, if $\mathcal{P} \subset \mathcal{P}_0$ is a random m -element subset and $\mathcal{H} \subset \mathcal{H}_0$ is a random n -element subset, they satisfy $\mathbb{E}(I(\mathcal{P}, \mathcal{H})) \geq \frac{1}{p} mn \geq \frac{1}{4s^{1/t}} m^{-\frac{1}{d+2-t}} \cdot mn$. Therefore, we can pick a pair $(\mathcal{P}, \mathcal{H})$ with at least as many incidences as the expectation, completing the proof. \square

3.4 Constructions of dense incidence graphs with no large bipartite subgraphs

In this section, we prove the upper bounds of Theorem 1.3. However, instead of assuming $m \leq n$ and giving bounds on both $\text{rs}_d(m, n, I)$ and $\text{rs}_d(n, m, I)$, as originally stated in Theorem 1.3, we prefer to keep the notation consistent by always denoting the number of hyperplanes by m and the number of points by n . Thus, we will give bounds on $\text{rs}_d(m, n, I)$ which depend on $\max\{m, n\}$ and $\min\{m, n\}$. With this notational change, Theorem 1.3 can be equivalently restated as follows.

Theorem 3.13. *Let d, m, n be positive integers, $\varepsilon \in (0, 1)$ and $I = \varepsilon mn$.*

- *If $\varepsilon \geq \frac{1}{4} \max\{m, n\}^{-1/d}$ and $\varepsilon \geq \frac{1}{4} \min\{m, n\}^{-1/(d-1)}$, then $\text{rs}_d(m, n, I) \leq O_d(\varepsilon^{d-1} mn)$.*
- *If $\varepsilon < \frac{1}{4} \max\{m, n\}^{-1/d}$ or $\varepsilon < \frac{1}{4} \min\{m, n\}^{-1/(d-1)}$ then $\text{rs}_d(m, n, I) \leq O_d(\varepsilon \max\{m, n\})$.*

The implied constant depending on the dimension in the above theorem is at most $2^{O(d \log d)}$. The main geometric construction is contained in the following lemma, which can be thought of as a strengthening of Lemma 3.12.

Lemma 3.14. *Let p be a prime and let d, t, k, ℓ be positive integers satisfying $t \leq d-1$ and $k, \ell \leq p^{d-t}$. Then, there exist an integer $s = O(d^d)$, a set of points \mathcal{P} and a set of hyperplanes \mathcal{H} in \mathbb{F}_p^d satisfying the following properties:*

- (i) $kp^t/2 \leq |\mathcal{P}| \leq kp^t$, $\frac{1}{4}\ell p^t \leq |\mathcal{H}| \leq \ell sp^t$, and every point of \mathcal{P} is contained in exactly $|\mathcal{H}|/p$ hyperplanes of \mathcal{H} . In particular, $I(\mathcal{P}, \mathcal{H}) = \frac{|\mathcal{P}| \cdot |\mathcal{H}|}{p}$.
- (ii) For $a \in \{1, \dots, t-1\}$ and any set $A \subseteq \mathcal{P}$ of size $|A| \geq skp^a$, A is not contained in a $(d-t+a)$ -dimensional affine subspace.
- (iii) For $b \in \{1, \dots, t-2\}$ and any set $B \subseteq \mathcal{H}$ of size $|B| \geq s^2 \ell p^b$, the intersection $\bigcap_{H_i \in B} H_i$ does not contain a $(t-1-b)$ -dimensional affine subspace.

Proof. By Lemma 3.10, there exists $s = O(d^d)$ such that for every prime p there exists a $(d-t, s)$ -subspace evasive set \mathcal{P}_0 containing p^t points of \mathbb{F}_p^d . In other words, any s points of \mathcal{P}_0 span an affine subspace of dimension at least $d-t+1$. Furthermore, there exists a set of points \mathcal{Q}_0 containing p^{t-1} points of \mathbb{F}_p^d such that \mathcal{Q}_0 is $(d-t+1, s)$ -subspace evasive, which means that any s points of \mathcal{Q}_0 span an at least $(d-t+2)$ -dimensional affine subspace.

The set \mathcal{P} is constructed as the union of k random translates of \mathcal{P}_0 . More precisely, we choose k random vectors $u_1, \dots, u_k \in \mathbb{F}_p^d$ (with repetition) and set $\mathcal{P} = \bigcup_{i=1}^k (\mathcal{P}_0 + u_i)$. The set \mathcal{P} constructed in this way is $(d-t, ks)$ -subspace evasive, since any affine space of dimension $d-t$ containing ks elements of \mathcal{P} must contain at least s elements of one of the translates $\mathcal{P}_0 + u_i$. But this is impossible, since \mathcal{P}_0 is a $(d-t, s)$ -subspace evasive set.

Next, we compute the expected size of \mathcal{P} and argue that u_1, \dots, u_k can be chosen such that $|\mathcal{P}| \geq \frac{1}{2} kp^t$. For a given point $x \in \mathbb{F}_p^d$, the probability that x lies in $\mathcal{P}_0 + u_i$ is

$$\mathbb{P}[x \in (\mathcal{P}_0 + u_i)] = \mathbb{P}[u_i \in (x - \mathcal{P}_0)] = \frac{1}{p^d} |\mathcal{P}_0|.$$

Thus, the expected size of \mathcal{P} is

$$\begin{aligned}\mathbb{E}[|\mathcal{P}|] &= \sum_{x \in \mathbb{F}_p^d} \mathbb{P}[x \in \mathcal{P}] = \sum_{x \in \mathbb{F}_p^d} \left(1 - \prod_{i=1}^k \mathbb{P}[x \notin (\mathcal{P}_0 + u_i)]\right) = \sum_{x \in \mathbb{F}_p^d} \left(1 - \left(1 - \frac{|\mathcal{P}_0|}{p^d}\right)^k\right) \\ &\geq p^d \left(1 - \exp\left(-\frac{k|\mathcal{P}_0|}{p^d}\right)\right) \geq p^d \frac{k|\mathcal{P}_0|}{2p^d} = \frac{1}{2}kp^t.\end{aligned}$$

In the above calculation, we have used the inequalities $1 - \frac{x}{2} \geq e^{-x} \geq 1 - x$ for $x \in (0, 1)$. Therefore, we conclude that u_1, \dots, u_k can be chosen such that \mathcal{P} has size at least $\frac{1}{2}kp^t$. On the other hand, we always have $|\mathcal{P}| \leq k|\mathcal{P}_0| = kp^t$.

Let us now explain how to process the set \mathcal{Q}_0 and how to construct from it the set of hyperplanes \mathcal{H} . Let $U = \{x \in \mathbb{F}_p^d : x(d) = 0\}$. Since we may translate the set \mathcal{Q}_0 without altering its properties, we may assume that at most $\frac{|\mathcal{Q}_0|}{p}$ elements of \mathcal{Q}_0 lie in U . Then, we define the set $\mathcal{Q}_1 = \mathcal{Q}_0 \setminus U$, which is still $(d - t + 1, s)$ -subspace evasive and has size $|\mathcal{Q}_1| \geq |\mathcal{Q}_0|/2$.

Pick $s\ell$ uniformly random vectors $v_1, \dots, v_{s\ell} \in U$ and define $\mathcal{Q}_2 = \bigcup_{i=1}^{s\ell} (\mathcal{Q}_1 + v_i)$. In a similar manner as above, we observe that \mathcal{Q}_2 is a $(d - t + 1, s\ell)$ -subspace evasive set. Our goal now is to show that $v_1, \dots, v_{s\ell}$ can be chosen in a way such that at least $\frac{1}{4}\ell p^{t-1}$ lines through the origin contain a point of \mathcal{Q}_2 . If $L \subset \mathbb{F}_p^d$ is a fixed line through the origin, note that L intersects $\mathcal{Q}_1 + v_i$ if and only if the translate of L given by $L - v_i$ intersects \mathcal{Q}_1 . Note that union of all the translates of L covers \mathcal{Q}_1 . Since \mathcal{Q}_1 is a $(d - t + 1, s)$ -subspace evasive set, each translate of L contains at most s points of \mathcal{Q}_1 and so at least $|\mathcal{Q}_1|/s$ translates of L intersect \mathcal{Q}_1 . Thus, we conclude $\mathbb{P}[L \cap (\mathcal{Q}_1 + v_i) \neq \emptyset] \geq \frac{|\mathcal{Q}_1|/s}{p^{d-1}}$. If we denote the number of lines through the origin intersecting \mathcal{Q}_2 by Y , by a similar calculation as above we have

$$\mathbb{E}[Y] = \sum_{L \text{ line through } 0} 1 - \mathbb{P}[L \cap \mathcal{Q}_2 = \emptyset] \geq \sum_{L \text{ line through } 0} 1 - \left(1 - \frac{|\mathcal{Q}_1|}{sp^{d-1}}\right)^{s\ell} \geq p^{d-1} \frac{\ell|\mathcal{Q}_1|}{2p^{d-1}} \geq \frac{1}{4}\ell p^{t-1}.$$

We form the subset $\mathcal{Q} \subseteq \mathcal{Q}_2$ by retaining at most one point on each line through the origin. The above calculation implies that $v_1, \dots, v_{s\ell}$ can be chosen such that $|\mathcal{Q}| \geq \frac{1}{4}\ell p^{t-1}$. On the other hand, the size of \mathcal{Q} is bounded by $|\mathcal{Q}| \leq |\mathcal{Q}_2| \leq s\ell|\mathcal{Q}_0| \leq s\ell p^{t-1}$ and \mathcal{Q} is a $(d - t + 1, s\ell)$ -subspace evasive set.

Finally, we choose \mathcal{H} to be the set of hyperplanes given by the equations $\langle q, x \rangle = b$, for $q \in \mathcal{Q}$ and $b \in \mathbb{F}_p$. We ensured that no two points $q_1, q_2 \in \mathcal{Q}$ lie on the same line through the origin, meaning that the hyperplanes given by $\langle q_1, x \rangle = b_1$, $\langle q_2, x \rangle = b_2$ are distinct for all $q_1, q_2 \in \mathcal{Q}$ and $b_1, b_2 \in \mathbb{F}_p$.

All that is left is to show that the conditions of the theorem hold for the sets \mathcal{P}, \mathcal{H} . We have already shown that $|\mathcal{P}| \in [kp^t/2, kp^t]$. Furthermore, since every pair $(q, b) \in \mathcal{Q} \times \mathbb{F}_p$ yields a different hyperplane of \mathcal{H} , we conclude $|\mathcal{H}| = p|\mathcal{Q}| \in [\frac{1}{4}\ell p^t, s\ell p^t]$. Finally, for every $x \in \mathcal{P}$ and every $q \in \mathcal{Q}$, there is a unique b for which $\langle x, q \rangle = b$, and therefore the number of hyperplanes containing x is precisely $|\mathcal{Q}| = |\mathcal{H}|/p$. This verifies (i).

Now, we prove that for any $a \in \{1, \dots, t-1\}$, no $(d - t + a)$ -dimensional affine subspace contains skp^a points of \mathcal{P} . Suppose this is not the case and there was a set $A \subseteq \mathcal{P}$ of size $|A| \geq skp^a$ which was contained in a $(d - t + a)$ -dimensional affine space V . Note that V can be partitioned into p^a translates of some $(d - t)$ -dimensional affine subspace $V' < V$. But then there exists a translate of V' containing at least ks points of \mathcal{P} . This is a contradiction since \mathcal{P} is a $(d - t, ks)$ -subspace evasive set, thus establishing (ii).

We proceed by a similar argument to prove (iii). Let B be a set of $s^2\ell p^b$ hyperplanes, denoted by $H_1, \dots, H_{s^2\ell p^b}$. Furthermore, we assume that the equation of H_i is given by $\langle q_i, x \rangle = b_i$. Observe that if the intersection of the hyperplanes H_i is nonempty, one must have $q_i \neq q_j$ for all $i, j \in [s^2\ell p^b]$, since otherwise two parallel hyperplanes belong to B . Let us denote the set of all points q_i for $i \in [s^2\ell p^b]$ by $B' \subseteq \mathcal{Q}$. Assume, for contradiction, that $\bigcap_{i=1}^{s^2\ell p^b} H_i$ contains a $(t - 1 - b)$ -dimensional affine subspace. This affine subspace can be written as $V + q$ for some $q \in \mathbb{F}_p^d$ and some $(t - 1 - b)$ -dimensional linear subspace V . Then all points q_i are contained in the orthogonal complement of

V , which we denote by V^\perp . But V^\perp is a $(d - t + 1 + b)$ -dimensional linear space, which contains $s^2\ell p^b$ points of \mathcal{Q} . By partitioning this space into p^b translates of a $(d - t + 1)$ -dimensional affine space, and using that \mathcal{Q} is a $(d - t + 1, \ell s^2)$ -subspace evasive set, one arrives at a contradiction as above. \square

In what follows, we prove Theorem 3.13 in a series of lemmas, addressing different regimes.

Lemma 3.15. *Let $\varepsilon \in (0, 1)$, and let d and m, n be integers which satisfy*

$$\varepsilon \leq \frac{1}{4} \max\{m, n\}^{-\frac{1}{d}}.$$

Then there exists a prime p , a set \mathcal{P} of m points and a set \mathcal{H} of n hyperplanes in \mathbb{F}_p^d such that $I(\mathcal{P}, \mathcal{H}) \geq \varepsilon mn$ and $\text{rs}(\mathcal{P}, \mathcal{H}) \leq (2s)^4 \varepsilon \max\{m, n\}$, where s is given by Lemma 3.14.

Proof. For notational convenience, we assume that $\max\{m, n\} = n$, noting that the same argument applies if $\max\{m, n\} = m$. Let p be a prime satisfying $\varepsilon^{-1}/2 < p \leq \varepsilon^{-1}$, and set $t = \lfloor \log_p 4n \rfloor$ and $\ell = \lceil \frac{4n}{p^t} \rceil \leq p$. Note that $4n \leq \ell p^t \leq 8n$. Since $p > \varepsilon^{-1}/2 \geq 2n^{1/d}$, we have $p^d > 4n$ and so $t \leq d - 1$. Thus, Lemma 3.14 can be applied with $k = \ell$ to get a set of points \mathcal{P}_0 and a set of hyperplanes \mathcal{H}_0 with $m_0 = |\mathcal{P}_0| \geq \ell p^t/2$, $n_0 = |\mathcal{H}_0| \geq \ell p^t/4$ and $I(\mathcal{P}_0, \mathcal{H}_0) = \frac{1}{p} m_0 n_0$, satisfying (i), (ii) and (iii).

Let us now argue that $\text{rs}(\mathcal{P}_0, \mathcal{H}_0) \leq (2s)^4 \varepsilon n$. Let $A \subseteq \mathcal{P}_0, B \subseteq \mathcal{H}_0$ such that every point of A lies on every hyperplane in B . We consider three cases: either $|A| \geq s^2 \ell p^{t-2}$, $|B| \geq s^2 \ell p^{t-2}$ or $\max\{|A|, |B|\} < s^2 \ell p^{t-2}$. In all cases, we show that

$$|A| \cdot |B| \leq s^4 \ell p^{t-1} \leq 2s^4 \varepsilon \ell p^t \leq 16s^4 \varepsilon n.$$

If $|A| \geq s^2 \ell p^{t-2}$, then by (ii) of Lemma 3.14, the points of A span an affine subspace of dimension at least $d - t + t - 1 = d - 1$, which means that at most one of the hyperplanes contains them all. On the other hand, again by (ii) of Lemma 3.14, any hyperplane can contain at most $s \ell p^{t-1}$ points of \mathcal{P}_0 and thus $|A| \leq s \ell p^{t-1}$. In conclusion, we have $|A| \leq s \ell p^{t-1}$ and $|B| = 1$, implying $|A| \cdot |B| \leq s \ell p^{t-1}$ as claimed.

The case $|B| \geq s^2 \ell p^{t-2}$ is very similar. By (iii) of Lemma 3.14, the intersection of hyperplanes of B does not contain a line, meaning that $\bigcap_{H \in B} H$ is a single point. Furthermore, there are exactly $|\mathcal{H}_0|/p$ hyperplanes of \mathcal{H}_0 containing any point of \mathcal{P}_0 . Since $|\mathcal{H}_0| \leq \ell s p^t$, we have $|B| \leq |\mathcal{H}_0|/p \leq s \ell p^{t-1}$, and therefore $|A| \cdot |B| \leq s \ell p^{t-1}$.

Finally, if $\max\{|A|, |B|\} < s^2 \ell p^{t-2}$, we let a and b be the smallest integers for which $s^2 \ell p^a \leq |A| < s^2 \ell p^{a+1}$ and $s^2 \ell p^b \leq |B| < s^2 \ell p^{b+1}$. By property (ii) of Lemma 3.14, the points of A span an at least $(d - t + a + 1)$ -dimensional affine space. Similarly, by property (iii) of Lemma 3.14, the hyperplanes of B intersect in at most a $(t - 2 - b)$ -dimensional space. Since all points of A belong to all hyperplanes of B , it must be that $d - t + a + 1 \leq t - 2 - b$, i.e. $a + b + 2 \leq 2t - d - 1$. Thus,

$$|A| \cdot |B| \leq s^4 \ell^2 p^{a+b+2} \leq s^4 \ell p \cdot p^{2t-d-1} \leq s^4 \ell p^{t-1},$$

where the last inequality holds by noting that $t \leq d - 1$. This completes the third case.

To finish the proof, let \mathcal{P} be a random subset of \mathcal{P}_0 containing m points and \mathcal{H} be a random subset of \mathcal{H}_0 containing n hyperplanes. One should verify that $m \leq |\mathcal{P}_0|$ and $n \leq |\mathcal{H}_0|$, but this is clear since $|\mathcal{P}_0|, |\mathcal{H}_0| \geq \frac{1}{4} \ell p^t \geq n$ and $n \geq m$. We also have $\text{rs}(\mathcal{P}, \mathcal{H}) \leq \text{rs}(\mathcal{P}_0, \mathcal{H}_0) \leq (2s)^4 \varepsilon n$. Furthermore, the expected number of incidences between \mathcal{P} and \mathcal{H} is

$$\mathbb{E}[I(\mathcal{P}, \mathcal{H})] = \frac{n}{n_0} \frac{m}{m_0} I(\mathcal{P}_0, \mathcal{H}_0) = \frac{1}{p} mn \geq \varepsilon mn.$$

Thus, there exist subsets \mathcal{P}, \mathcal{H} of $\mathcal{P}_0, \mathcal{H}_0$ with at least εmn incidences and $\text{rs}(\mathcal{P}, \mathcal{H}) \leq (2s)^4 \varepsilon n$. \square

Lemma 3.16. *Let $\varepsilon \in (0, 1)$, and let d and $m \leq n$ be positive integers which satisfy*

$$\varepsilon \leq \frac{1}{4} m^{-\frac{1}{d-1}}.$$

Then there exists a prime power q , a set \mathcal{P} of m points and a set \mathcal{H} of n hyperplanes in \mathbb{F}_q^d such that $I(\mathcal{P}, \mathcal{H}) \geq \varepsilon mn$ and $\text{rs}(\mathcal{P}, \mathcal{H}) \leq (4s)^4 \varepsilon n$, where s is given by Lemma 3.14.

Proof. The main idea of the proof is to take a symmetric $(d-1)$ -dimensional construction described in Lemma 3.15 and “blow it up” to obtain our desired configuration.

More precisely, we apply Lemma 3.15 to find a set of m points \mathcal{P}_0 and a set of m hyperplanes \mathcal{H}_0 in \mathbb{F}_p^{d-1} for some prime p for which $I(\mathcal{P}_0, \mathcal{H}_0) \geq \varepsilon m^2$ and $\text{rs}(\mathcal{P}_0, \mathcal{H}_0) \leq (2s)^4 \varepsilon m$. The hyperplanes and points may be considered in the larger ambient space \mathbb{F}_q^{d-1} , for any prime power $q = p^k$. Passing to \mathbb{F}_q^{d-1} does not change the incidence graph. In the rest of the proof, we fix q to be a power of p larger than n/m .

Now, for each $x \in \mathbb{F}_q^{d-1}$, let $x' \in \mathbb{F}_q^d$ be the point whose first $d-1$ coordinates are the same as x , and the last coordinate is 0. Set $\mathcal{P} = \{x' | x \in \mathcal{P}_0\}$. Furthermore, we fix a set $S = \{s_1, \dots, s_\ell\} \subseteq \mathbb{F}_q$ of size $\ell = \lceil n/m \rceil$. Then, for every hyperplane $H \in \mathcal{H}_0$, if H is given by the equation $\langle a, x \rangle = a_1 x_1 + \dots + a_{d-1} x_{d-1} = b$, we define hyperplanes $H^{(1)}, \dots, H^{(\ell)}$ in \mathbb{F}_q^d , where $H^{(i)}$ is given by the equations $a_1 x_1 + \dots + a_{d-1} x_{d-1} + s_i x_d = b$. Then, we let $\mathcal{H}_1 = \{H^{(i)} | i \in [\ell], H \in \mathcal{H}_0\}$.

The sizes of the newly constructed sets are $|\mathcal{P}| = m, |\mathcal{H}_1| = m \lceil \frac{n}{m} \rceil > n$. Furthermore, if $x \in \mathcal{P}_0$ is incident to $H \in \mathcal{H}_0$, then x' is incident to all of the hyperplanes $H^{(1)}, \dots, H^{(\ell)}$. Thus, the number of incidences between \mathcal{P} and \mathcal{H}_1 is $I(\mathcal{P}, \mathcal{H}_1) \geq \ell I(\mathcal{P}_0, \mathcal{H}_0) \geq \varepsilon \ell |\mathcal{P}_0| |\mathcal{H}_0| = \varepsilon |\mathcal{P}| |\mathcal{H}_1|$. Finally, the size of the largest complete bipartite graph is bounded by

$$\text{rs}(\mathcal{P}, \mathcal{H}_1) \leq \ell \cdot \text{rs}(\mathcal{P}_0, \mathcal{H}_0) \leq \ell \cdot (2s)^4 \varepsilon m \leq (4s)^4 \varepsilon n.$$

Taking a random m element subset \mathcal{H} of \mathcal{H}_1 , the expectation of $I(\mathcal{P}, \mathcal{H})$ is at least εmn , so there is a choice for \mathcal{H} such that $(\mathcal{P}, \mathcal{H})$ satisfies our desired conditions. \square

Lemma 3.17. *Let $\varepsilon \in (0, 1)$, and let d and $m \geq n$ be positive integers which satisfy*

$$\varepsilon \leq \frac{1}{4} n^{-\frac{1}{d-1}}.$$

Then there exists a prime power q , a set \mathcal{P} of m points and a set \mathcal{H} of n hyperplanes in \mathbb{F}_q^d such that $I(\mathcal{P}, \mathcal{H}) \geq \varepsilon mn$ and $\text{rs}(\mathcal{P}, \mathcal{H}) \leq (4s)^4 \varepsilon m$, where s is given by Lemma 3.14.

Proof. The proof is verbatim the same as the proof of Lemma 3.16, the only difference being that we “blow up” points instead of hyperplanes. More precisely, we start from a set of n points \mathcal{P}_0 and a set of n hyperplanes \mathcal{H}_0 in \mathbb{F}_p^{d-1} for which $I(\mathcal{P}_0, \mathcal{H}_0) \geq \varepsilon n^2$ and $\text{rs}(\mathcal{P}_0, \mathcal{H}_0) \leq (2s)^4 \varepsilon n$ and consider them in \mathbb{F}_q^{d-1} , where $q > m/n$.

Then, we fix a set $S = \{s_1, \dots, s_\ell\} \subseteq \mathbb{F}_q$ of size $\ell = \lceil m/n \rceil$ and for each point $x \in \mathbb{F}_q^{d-1}$ we define $x^{(i)} \in \mathbb{F}_q^d$ be the point whose first $d-1$ coordinates are the same as x , and the last coordinate is s_i . Set $\mathcal{P}_1 = \{x^{(i)} | x \in \mathcal{P}_0, i \in [\ell]\}$. Then, for every hyperplane $H \in \mathcal{H}_0$, if H is given by the equation $\langle a, x \rangle = a_1 x_1 + \dots + a_{d-1} x_{d-1} = b$, we define the hyperplane H' in \mathbb{F}_q^d , given by the equation $a_1 x_1 + \dots + a_{d-1} x_{d-1} + 0 \cdot x_d = b$. Then, we let $\mathcal{H} = \{H' | H \in \mathcal{H}_0\}$.

The sizes of the newly constructed sets are $|\mathcal{H}| = n, |\mathcal{P}_1| = n \lceil \frac{m}{n} \rceil > n$. Thus, by choosing a random subset $\mathcal{P} \subseteq \mathcal{P}_1$ and performing the same calculation as in the proof of Lemma 3.16 one arrives at the desired conclusion. \square

Lemma 3.18. *Let $\varepsilon \in (0, 1)$, and let d and m, n be positive integers which satisfy*

$$4\varepsilon \geq n^{-\frac{1}{d-2}} \quad \text{and} \quad 4\varepsilon \geq m^{-\frac{1}{d-2}}.$$

Then there exists a prime power q , a set \mathcal{P} of m points and a set \mathcal{H} of n hyperplanes in \mathbb{F}_q^d such that $I(\mathcal{P}, \mathcal{H}) \geq \varepsilon mn$ and $\text{rs}(\mathcal{P}, \mathcal{H}) \leq 2^6 s^4 (4\varepsilon)^{d-1} mn$, where s is given by Lemma 3.14.

Proof. The proof of this lemma is somewhat similar to the proofs of Lemmas 3.16 and 3.17, since we use the idea of “blowing-up” a lower-dimensional construction. However, in this case we use a $(d-2)$ -dimensional configuration as a starting point and we “blow-up” both points and hyperplanes.

Let p be a prime between $\varepsilon^{-1}/2$ and ε^{-1} , and set $N = (4\varepsilon)^{-(d-2)}$. We remark that the inequalities $N \leq m, n$ are satisfied. Note that the condition of Lemma 3.15 is also satisfied with $d-2$ instead

of d , and N instead of m and n . Therefore, we can find sets $\mathcal{P}_0, \mathcal{H}_0$ of points and hyperplanes in \mathbb{F}_p^{d-2} , both of size N , and satisfying $I(\mathcal{P}_0, \mathcal{H}_0) \geq \varepsilon N^2$ and $\text{rs}(\mathcal{P}_0, \mathcal{H}_0) \leq (2s)^4 \varepsilon N$.

As in the proof of Lemma 3.16, we consider $\mathcal{P}_0, \mathcal{H}_0$ in \mathbb{F}_q^{d-2} instead of \mathbb{F}_p^{d-2} , where q is the power of p satisfying $q > \max\{\frac{m}{N}, \frac{n}{N}\}$. Furthermore, we fix two sets $S, T \subseteq \mathbb{F}_q$, such that $S = \{s_1, \dots, s_k\}$ and $T = \{t_1, \dots, t_\ell\}$, where $k = \lceil \frac{m}{N} \rceil, \ell = \lceil \frac{n}{N} \rceil$. To every point $x \in \mathcal{P}_0$, we associate k points of \mathbb{F}_q^d as follows. For $i \in [k]$, let $x^{(i)} \in \mathbb{F}_q^d$ be equal to x on the first $d-2$ coordinates, equal to s_i on the coordinate $d-1$, and equal to 0 on the last coordinate. Similarly, for a hyperplane $H \in \mathcal{H}_0$ given by the equation $a_1 x_1 + \dots + a_{d-2} x_{d-2} = b$, we associate the hyperplanes $H^{(j)}$ in \mathbb{F}_q^d for $j \in [\ell]$ defined by the equations $a_1 x_1 + \dots + a_{d-2} x_{d-2} + 0 \cdot x_{d-1} + t_j x_d = b$. Finally, we set $\mathcal{P}_1 = \{x^{(i)} | i \in [k], x \in \mathcal{P}_0\}$ and $\mathcal{H}_1 = \{H^{(j)} | H \in \mathcal{H}_0, j \in [\ell]\}$.

If $x \in H$, then it is easy to check that $x^{(i)} \in H^{(j)}$ for all $i \in [k], j \in [\ell]$, and no other type of incidences emerge in $(\mathcal{P}_1, \mathcal{H}_1)$. Thus, $I(\mathcal{P}_1, \mathcal{H}_1) \geq k\ell I(\mathcal{P}_0, \mathcal{H}_0) \geq \varepsilon k |\mathcal{P}_0| \ell |\mathcal{H}_0| = \varepsilon |\mathcal{P}_1| |\mathcal{H}_1|$. Furthermore,

$$\text{rs}(\mathcal{P}_1, \mathcal{H}_1) \leq k\ell \text{rs}(\mathcal{P}_0, \mathcal{H}_0) \leq \frac{4mn}{N^2} \cdot (2s)^4 \varepsilon N = 2^6 s^4 \varepsilon \frac{mn}{N}.$$

Recalling that $N = (4\varepsilon)^{-(d-2)}$, we obtain $\text{rs}(\mathcal{P}_1, \mathcal{H}_1) \leq 2^6 s^4 (4\varepsilon)^{d-1} mn$. Subsampling a random m element subset \mathcal{P} of \mathcal{P}_1 , and a random n element subset \mathcal{H} of \mathcal{H}_1 , completes the proof. \square

Proof of Theorem 3.13. We claim that it suffices to take $C = 4^{d+4} s^4 = O((4d)^{4d})$. Consider three cases.

Case 1. $\varepsilon \geq \frac{1}{4} \max\{m, n\}^{-1/d}$ and $\varepsilon \geq \frac{1}{4} \min\{m, n\}^{-1/(d-1)}$.

Then $4\varepsilon > m^{-1/(d-2)}$ and $4\varepsilon > n^{-1/(d-2)}$ are also satisfied, so Lemma 3.18 applies. Therefore, we have $\text{rs}_d(m, n, \varepsilon mn) \leq C\varepsilon^{d-1} mn$.

Case 2. $\varepsilon \leq \frac{1}{4} \max\{m, n\}^{-1/d}$.

In this case, Lemma 3.15 shows that $\text{rs}_d(m, n, \varepsilon mn) \leq (2s)^4 \varepsilon \max\{m, n\} \leq C\varepsilon \max\{m, n\}$.

Case 3. $\varepsilon \leq \frac{1}{4} \min\{m, n\}^{-1/(d-1)}$.

We apply either Lemma 3.16 or Lemma 3.17, depending on whether $m \geq n$ or $n \geq m$. This shows that $\text{rs}_d(m, n, \varepsilon mn) \leq (4s)^4 \varepsilon \max\{m, n\} \leq C\varepsilon \max\{m, n\}$. \square

4 Point-variety incidences

4.1 Basics of algebraic geometry

Much of the following material is based on the Appendix of [14] and the excellent book of Cox, Little and O'Shea [11].

Let $\overline{\mathbb{F}}$ be the algebraic closure of the field \mathbb{F} . We denote by $\overline{\mathbb{F}}[x_1, \dots, x_D]$ the set of polynomials in variables x_1, \dots, x_D with coefficients in $\overline{\mathbb{F}}$. An *affine algebraic set* V is a set of common zeros of polynomials $f_1, \dots, f_k \in \overline{\mathbb{F}}[x_1, \dots, x_D]$ in $\overline{\mathbb{F}}^D$, and this set is denoted by $V(f_1, \dots, f_k)$. Formally,

$$V(f_1, \dots, f_k) = \left\{ (a_1, \dots, a_d) \in \overline{\mathbb{F}}^D : f_i(a_1, \dots, a_d) = 0 \text{ for all } i \in [k] \right\}.$$

In a similar way, one can define projective algebraic sets. In this case, the ambient projective space $\mathbb{P}^D(\overline{\mathbb{F}})$ over the field $\overline{\mathbb{F}}$ is defined as the set of equivalence classes of points in $\overline{\mathbb{F}}^{D+1} \setminus \{(0, \dots, 0)\}$, where $(a_0, \dots, a_D) \sim (b_0, \dots, b_D)$ if $a_i = \lambda b_i$ for every i for some $\lambda \in \overline{\mathbb{F}}$. Note that for every homogeneous polynomial $f \in \overline{\mathbb{F}}[x_0, \dots, x_D]$ and every $\lambda \neq 0$, we have $f(a_0, \dots, a_D) = 0$ if and only if $f(\lambda a_0, \dots, \lambda a_D) = 0$. Hence, the set of zeros of a homogeneous polynomial in $\mathbb{P}^D(\overline{\mathbb{F}})$ is well-defined. Thus, we can define a *projective algebraic set* determined by homogeneous polynomials $f_1, \dots, f_k \in \overline{\mathbb{F}}[x_0, \dots, x_D]$ as the common zero set of these polynomials.

One should think of the projective space as the “completion” of the affine space, in which one adds certain points at infinity. Hence, to every affine algebraic set one can uniquely associate a projective variety which extends it. Formally, one can define a homogenization f^h of the degree d polynomial $f \in \overline{\mathbb{F}}[x_1, \dots, x_D]$ by setting $f^h(x_0, \dots, x_D) = x_0^d f(x_1/x_0, \dots, x_D/x_0)$. Then, an affine algebraic set $V(f_1, \dots, f_k) \subseteq \overline{\mathbb{F}}^D$ extends to the projective algebraic set defined by polynomials f_1^h, \dots, f_k^h .

Finite unions and arbitrary intersections of algebraic sets are also algebraic. An algebraic set V is *irreducible* if there are no algebraic sets $V_1, V_2 \subsetneq V$ for which $V_1 \cup V_2 = V$, and an irreducible algebraic set is called a *variety*. The decomposition theorem for algebraic sets states that there is a unique way to express an algebraic set as the finite union of varieties, which are then called the *components* of the algebraic set (see Section 6 of Chapter 4 in [11]).

Given an affine algebraic set $Y \subset \overline{\mathbb{F}}^D$, the *ideal of Y* , denoted by $I(Y)$, is defined as the set of polynomials $f \in \overline{\mathbb{F}}[x_1, \dots, x_D]$ that vanish on Y . Clearly, if $g, h \in \overline{\mathbb{F}}[x_1, \dots, x_D]$ are polynomials, then the restrictions of g and h on Y are identical if and only if $g - h \in I(Y)$. Therefore, the space of polynomial functions on Y is isomorphic to the quotient $\overline{\mathbb{F}}[x_1, \dots, x_D]/I(Y)$. One can analogously define the ideal of a projective algebraic set $Y \subseteq \mathbb{P}^D(\overline{\mathbb{F}})$, denoted by $I(Y)$, as the set of all homogeneous polynomials in $\overline{\mathbb{F}}[x_0, \dots, x_D]$ vanishing on Y .

Next, we define the dimension and the degree of a variety. We give two definitions of the dimension of a variety (for the proof of their equivalence, see Theorem I.7.5 in [28]). Both of these definitions will be useful in our proofs. According to the first definition, the *dimension* of the variety V is the maximum value of r for which there exists a chain of varieties V_0, V_1, \dots, V_r such that $\emptyset = V_0 \subsetneq V_1 \subsetneq \dots \subsetneq V_r \subsetneq V$.

Next, we give a definition the dimension based on the ideal $I(V)$. Fix a projective variety $V \subseteq \mathbb{P}^D(\overline{\mathbb{F}})$ and let $\overline{\mathbb{F}}_t[x_0, x_1, \dots, x_D]$ be the space of homogeneous polynomials of degree t , which is a finite-dimensional vector space over $\overline{\mathbb{F}}$. Further, let $I_t = I(V) \cap \overline{\mathbb{F}}_t[x_0, x_1, \dots, x_D]$ be the space of those homogeneous polynomials in $I(V)$ which have degree t . Since I_t is a subspace of $\overline{\mathbb{F}}_t[x_0, x_1, \dots, x_D]$, one may consider the dimension of the quotient $\dim(\overline{\mathbb{F}}_t[x_0, x_1, \dots, x_D]/I_t)$, which is called the *Hilbert function* of I . By Hilbert’s theorem, the function $t \mapsto \dim(\overline{\mathbb{F}}_t[x_0, \dots, x_D]/I_t)$ is a polynomial for t sufficiently large and it is called the *Hilbert polynomial* of $I(V)$ (see e.g. Proposition 3 on p. 487. in [11]). The *dimension* of the variety V , denoted by $\dim(V)$, is then the degree of the Hilbert polynomial associated to $I(V)$. Furthermore, the *degree* of the variety, denoted by $\deg V$, is defined as $(\dim V)!$ times the leading coefficient of this polynomials. Finally, the dimension and the degree of the affine variety are simply the dimension and the degree of the associated projective variety.

To get some intuition for these concepts, we remark that if $V = V(f)$ is a variety in D -dimensional space defined by a single irreducible polynomial f , then the dimension of V is $D - 1$ and the degree of V is $\deg f$.

To conclude this section, let us mention two important results that we use in our proofs. The first one concerns the intersections of varieties whose degree and dimension are known and holds for affine and projective varieties alike (for a reference, see Example 8.4.6 in [26]).

Theorem 4.1. *Let V_1 and V_2 be algebraic varieties of dimensions $d_1 \leq d_2$ in a projective or affine ambient space of dimension D such that $V_1 \not\subseteq V_2$. If Z_1, \dots, Z_k are the irreducible components of $V_1 \cap V_2$, then $\dim(Z_i) \leq d_1 - 1$ for all i and $\sum_{i=1}^k \deg(Z_i) \leq \deg V_1 \cdot \deg V_2$.*

The second statement we use is a uniform version of Hilbert’s theorem, which can be used to bound the Hilbert function (for a reference, see e.g. Chapter 9 of [38]).

Theorem 4.2. *Let V be a projective variety and let $I = I(V)$ be the associated ideal. Then, for every positive integer t , $\dim(\overline{\mathbb{F}}_t[x_0, x_1, \dots, x_D]/I_t) \leq \deg(V)t^{\dim V} + \dim V$.*

4.2 Proofs of upper bounds via induced Turán problems

In this section, we prove Theorem 1.4 through the following approach. We construct a bipartite graph $H_{d,\Delta}$ with the property that $G(\mathcal{P}, \mathcal{V})$ does not contain an induced copy of $H_{d,\Delta}$. Then, we

show that any $K_{s,s}$ -free bipartite graph with sides of size m, n which does not contain an induced copy of $H_{d,\Delta}$ has at most $O(m^{\frac{d}{d+1}}n)$ edges.

Let us begin by describing the forbidden induced bipartite subgraph $H = H_{d,\Delta}$. Let A and B be the parts of $H_{d,\Delta}$, and $k = 2^{\Delta^d} + 1$. Part B consists of $d + 1$ ‘‘layers’’ of vertices, which we describe as follows. The first two layers contain one vertex each, denoted by v_1 and v_2 , respectively. For $3 \leq \ell \leq d + 1$, the ℓ -th layer has $k^{\ell-2}$ vertices, denoted by $v_\ell^{(i_3, \dots, i_\ell)}$, where $i_3, \dots, i_\ell \in [k]$.

Next, we describe part A . For each $3 \leq \ell \leq d + 1$ and each sequence $(i_3, \dots, i_\ell) \in [k]^{\ell-2}$, we add k vertices $w_{\ell,1}^{(i_3, \dots, i_\ell)}, \dots, w_{\ell,k}^{(i_3, \dots, i_\ell)}$ to A whose neighbours are $v_1, v_2, v_3^{(i_3)}, v_4^{(i_3, i_4)}, \dots, v_\ell^{(i_3, \dots, i_\ell)}$. Note that every vertex of A has degree at most $d + 1$. Furthermore, observe that vertices $v_t^{(i_3, \dots, i_t)}$ and $v_r^{(j_3, \dots, j_r)}$ have a common neighbour if and only if one of the sequences (i_3, \dots, i_t) and (j_3, \dots, j_r) is a prefix of the other.

Proposition 4.3. *If \mathcal{P} is a set of points and \mathcal{V} is a set of d -dimensional varieties of degree at most Δ , the incidence graph $G(\mathcal{P}, \mathcal{V})$ does not contain $H_{d,\Delta}$ as an induced subgraph, where the vertices of A correspond to points and the vertices of B correspond to varieties.*

Proof. Assume that $G(\mathcal{P}, \mathcal{V})$ contains an induced copy of $H_{d,\Delta}$ and let $V_\ell^{(i_3, \dots, i_\ell)}$ be the variety corresponding to the vertex $v_\ell^{(i_3, \dots, i_\ell)}$. We show by induction on $2 \leq \ell \leq d + 1$ that one can choose a sequence i_3, \dots, i_ℓ with the following property. Let Z_1, \dots, Z_t be the irreducible components of the intersection $V_1 \cap V_2 \cap \dots \cap V_\ell^{(i_3, \dots, i_\ell)}$ that satisfy $\dim(Z_i) \leq d - \ell + 1$. Then the union $\bigcup_{i=1}^t Z_i$ contains all points corresponding to the common neighbours of $v_1, \dots, v_\ell^{(i_3, \dots, i_\ell)}$ in A .

For $\ell = 2$, this statement is obvious, since all components of $V_1 \cap V_2$ have dimension at most $d - 1$ by Theorem 4.1, noting that V_1 and V_2 are distinct varieties of dimension d . To show the inductive step, suppose we found a sequence i_3, \dots, i_ℓ satisfying the above property. Let Z_1, \dots, Z_t be the irreducible components of the intersection of these varieties satisfying $\dim(Z_i) \leq d - \ell + 1$. By iterated application of Theorem 4.1, one can deduce that $t \leq \Delta^\ell$.

For each variety $V_{\ell+1}^{(i_3, \dots, i_\ell, j)}$, $j \in \{1, \dots, k\}$, let us denote by T_j the set of components among Z_1, \dots, Z_t which are fully contained in $V_{\ell+1}^{(i_3, \dots, i_\ell, j)}$. As $k = 2^{\Delta^d} > 2^t$, there exist two sets T_j and $T_{j'}$ which are identical. In other words, the varieties $V_{\ell+1}^{(i_3, \dots, i_\ell, j)}$ and $V_{\ell+1}^{(i_3, \dots, i_\ell, j')}$ contain the exact same components Z_i . We claim that the sequence i_3, \dots, i_ℓ, j then satisfies the required conditions and suffices to perform the inductive step.

To see this, note that $V_{\ell+1}^{(i_3, \dots, i_\ell, j)}$ intersects all components $Z_i \notin T_j$ in subvarieties of dimension at most $d - \ell$ (or in the empty set). Thus, to complete the inductive step, it suffices to show that all common neighbours of $v_1, \dots, v_{\ell+1}^{(i_3, \dots, i_\ell, j)}$ are contained in the components $Z_i \notin T_j$. If there is a common neighbour P of these vertices in a component $Z_i \in T_j$, it belongs to the variety $V_{\ell+1}^{(i_3, \dots, i_\ell, j')}$, and thus it is connected to the vertex $v_{\ell+1}^{(i_3, \dots, i_\ell, j')}$. But this is impossible, since $v_{\ell+1}^{(i_3, \dots, i_\ell, j)}$ and $v_{\ell+1}^{(i_3, \dots, i_\ell, j')}$ have no common neighbours in $H_{d,\Delta}$.

Taking $\ell = d + 1$, we get a sequence i_3, \dots, i_{d+1} such that $V_1, \dots, V_{d+1}^{(i_3, \dots, i_{d+1})}$ has the following property. If Z_1, \dots, Z_t are the 0-dimensional components of $V_1 \cap \dots \cap V_{d+1}^{(i_3, \dots, i_{d+1})}$, then $\bigcup_{i=1}^t Z_i$ contains k points, corresponding to the common neighbours of $v_1, \dots, v_{d+1}^{(i_3, \dots, i_{d+1})}$. However, note that an iterated application of Theorem 4.1 implies $t \leq \Delta^{d+1}$ and so $t < k$. But this is a contradiction, since each Z_i is a single point. \square

Next, given a bipartite graph H , we prove a general upper bound on the number of edges in a $K_{s,s}$ -free and induced H -free bipartite graph. A similar statement was proved recently by Hunter and the authors of this paper [30], but their result is not directly applicable in case the host graph is bipartite with parts of unequal sizes. Fortunately, we can reuse most of the key auxiliary lemmas to adapt to this scenario.

Let $H = (A, B, E)$ be a bipartite graph such that every vertex in A has degree at most k . Given a graph G , we say that a set of vertices $S \subset V(G)$ is (H, k, s) -rich if for every $T \subset S$, $|T| \leq k$, we

have

$$|\{v \in V(G) \setminus S : N(v) \cap S = T\}| \geq (4|A|s)^{|A|}.$$

We need the following two results from [30].

Lemma 4.4 (Lemma 2.3 in [30]). *Let G be a graph not containing $K_{s,s}$. If G contains an (H, k, s) -rich independent set S of size $|B|$, then G contains H as an induced subgraph in which B is embedded into S .*

Proposition 4.5 (Proposition 2.4 in [30]). *Let G be a $K_{s,s}$ -free graph and $X \subseteq V(G)$ a set of at least $(4|A||B|s)^{4|B|+10}$ vertices in which every k -tuple of vertices has at least $(4|A||B|s)^{2|V(H)|}$ common neighbours. Then X contains an (H, k, s) -rich independent set of size $|B|$.*

Furthermore, we use the following simple statement about independent sets in hypergraphs.

Proposition 4.6. *Let \mathcal{H} be a k -uniform hypergraph with N vertices and M edges. Then \mathcal{H} contains an independent set of size at least $N^{\frac{k}{k-1}} / (4(M+N)^{1/(k-1)})$.*

Proof. Let $p = (N/2(M+N))^{1/(k-1)}$. Let $X \subseteq V(\mathcal{H})$ be a random sample in which each vertex is included independently with probability p . Then $\mathcal{H}[X]$ contains an independent set of size at least $|X| - e(\mathcal{H}[X])$, as we can remove a single vertex from every edge to get an independent set. But $\mathbb{E}(|X| - e(\mathcal{H}[X])) = pN - p^k M \geq \frac{pN}{2}$, so there is a choice for X such that $|X| - e(\mathcal{H}[X]) \geq pN/2 \geq N^{\frac{k}{k-1}} / (4(M+N)^{1/(k-1)})$. \square

Now we are ready to prove our bound, which is a simple combination of the previous two results and the dependant random choice method [24].

Lemma 4.7. *Let $G = (U, V; E)$ be a bipartite graph, $|U| = m$ and $|V| = n$, and let $H = (A, B; E)$ be a bipartite graph such that every vertex in A has degree at most $k \geq 2$. If G contains no induced copy of H in which A is embedded to U , and G is $K_{s,s}$ -free, then*

$$E(G) = O_{H,s} \left(m^{\frac{k-1}{k}} n + m \right).$$

Proof. Assume $e(G) \geq K(m^{\frac{k-1}{k}} n + m)$, where $K = K(H, s)$ is specified later. Let $t = (4|A||B|s)^{2|V(H)|}$ and $z = (4|A||B|s)^{4|B|+10}$ be constants depending only on H and s . We show that if G contains no $K_{s,s}$, then G contains an induced copy of H in which A is embedded to U .

Let u be a vertex in U , and let $X_0 = N_G(u)$. Define the k -uniform hypergraph \mathcal{H} on vertex set X_0 such that a k -element set $D \subset X_0$ is an edge if D has at most t common neighbours in G . First, we show that there is a choice for u such that \mathcal{H} contains an independent set of size at least z . By Proposition 4.6, there is independent set of size at least $|X_0|^{\frac{k}{k-1}} / (4(|X_0| + e(\mathcal{H}))^{\frac{1}{k-1}})$. Hence, it is enough to show that there exists a choice of $u \in U$ for which $|X_0|^k > (4z)^{k-1} (|X_0| + e(\mathcal{H}))$.

Choose u uniformly at random from U . Then $\mathbb{E}[|X_0|] = \frac{e(G)}{m} \geq K \frac{n}{m^{1/k}} + K$. On the other hand, for any fixed k element set $D \subset V$ with less than t common neighbours, we have $\mathbb{P}(D \subset X_0) \leq \frac{t}{m}$. Since there are at most n^k such k element sets in V , we conclude that $\mathbb{E}[e(\mathcal{H})] \leq \frac{tn^k}{m}$. By Jensen's inequality, we have $\mathbb{E}[|X_0|^k] \geq \mathbb{E}[|X_0|]^k \geq K^k \frac{n^k}{m} + K^k$, and therefore

$$\mathbb{E}[|X_0|^k - (4z)^{k-1} (e(\mathcal{H}) + |X_0|)] > 0$$

when K is sufficiently large. Hence, there exists a choice of u for which \mathcal{H} contains an independent set of size z . Fix such a choice, and let $X \subset X_0$ be such an independent set.

As every k element subset of X has at least t common neighbors, we can apply Proposition 4.5 to find an (H, k, s) -rich set $S \subset X$ of size $|B|$. But then, by Lemma 4.4, there is an embedding of H in which B is embedded to S , finishing the proof. \square

Proof of Theorem 1.4. By Proposition 4.3, the incidence graph $G(\mathcal{P}, \mathcal{V})$ does not contain an induced copy of the bipartite graph $H_{d,\Delta}$, where part A is corresponds to points. Also, every vertex of A has degree at most $d+1$, so applying Lemma 4.7 with $k = d+1$ and $H = H_{d,\Delta}$ gives the desired bound. \square

4.3 Proofs of upper bounds via VC-dimension

In this section, we show a variant of Theorem 1.4 using a different approach. Note that we use a slightly stronger assumption that the varieties have bounded description complexity, instead of just bounded degree. However, since the description complexity can be bounded by a function of the degree and the ambient dimension, this assumption does not make a big difference (see e.g. Theorem 2.1.16 in [53]).

We say that the variety V has *description complexity* at most t if it can be defined using at most t polynomials $f_1, \dots, f_t \in \mathbb{F}[x_1, \dots, x_D]$ of degree at most t .

Theorem 4.8. *Let \mathcal{P} be a set of m points in \mathbb{F}^D and let \mathcal{V} be a set of n varieties in \mathbb{F}^D , each of dimension d and description complexity at most t . If the incidence graph $G(\mathcal{P}, \mathcal{V})$ is $K_{s,s}$ -free, then*

$$I(\mathcal{P}, \mathcal{V}) \leq O_{D,t,s}(m^{\frac{d}{d+1}}n + m).$$

Let us prepare the proof of this theorem. First, we recall some basic notions from the theory of VC-dimension. Let X be a ground set and let \mathcal{F} be a family of subsets of X . The *shatter function* of the system \mathcal{F} , denoted by $\pi_{\mathcal{F}}(k)$, is defined as the maximum number of distinct intersections of a k -element set with members of \mathcal{F} . Formally,

$$\pi_{\mathcal{F}}(k) = \max_{A \subset X, |A|=k} |\{A \cap B : B \in \mathcal{F}\}|.$$

For a sequence of varieties V_1, \dots, V_k and a point x in \mathbb{F}^D , we define the *containment pattern* of V_1, \dots, V_k at x as the set of indices $I \subseteq [k]$ of those varieties V_i which contain x , that is $I = \{i \in [k] : x \in V_i\}$. Our first key lemma is an extension of a celebrated result of Rónyai, Babai and Ganapathy [42] on the number of zero-patterns of polynomials. This statement, which might be of independent interest, implies that the family of possible containment-patterns has a polynomial shatter function.

Lemma 4.9. *Let V_1, \dots, V_k be a sequence of varieties in \mathbb{F}^D , each of dimension d and description complexity at most t . The number of distinct containment patterns of V_1, \dots, V_k is at most $O_{D,t}(k^{d+1})$.*

The proof of Lemma 4.9 combines the ideas from [42] with Hilbert's theorem. Hence, before presenting the proof, let us recall the result of Rónyai, Babai and Ganapathy [42]. Using our terminology, given sequence of polynomials $f_1, \dots, f_k \in \mathbb{F}[x_1, \dots, x_D]$ and point $x \in \mathbb{F}^D$, the *zero-pattern* of f_1, \dots, f_k at x is the set of indices $I \subset [k]$ defined as $I = \{i \in [k] : f_i(x) = 0\}$. If $\mathcal{Z}(f_1, \dots, f_k)$ is the set of all zero-patterns of f_1, \dots, f_k and f_1, \dots, f_k are polynomials of degree at most Δ , then the number of distinct zero-patterns is at most

$$|\mathcal{Z}(f_1, \dots, f_k)| \leq \binom{k\Delta}{D}. \quad (4)$$

Note that (4) can be used as a black box to derive a weaker version of Lemma 4.9, since every containment patterns of V_1, \dots, V_k corresponds to a zero-patterns of their defining polynomials. However, this argument would lead to the weaker upper bound $O_t(k^D)$ in Lemma 4.9.

Proof of Lemma 4.9. We work in the algebraic closure $\overline{\mathbb{F}}$ of the field \mathbb{F} , which can only increase the number of containment patterns. We show that for every $a \in [k]$, the number of containment patterns $I \subset [k]$ with $a \in I$, i.e. the patterns coming from points $x \in V_a$, is at most $O_{t,D}(k^d)$. Then we are clearly done. Without loss of generality, let $a = 1$, and let M be the number of containment patterns $I \subseteq [k]$ with $1 \in I$.

For $i = 2, \dots, k$, let the defining polynomials of the variety V_i be $f_{i,1}, \dots, f_{i,t} \in \overline{\mathbb{F}}[x_1, \dots, x_D]$, where we may repeat some of the polynomials to ensure that we have exactly t polynomials. Since every containment pattern corresponds to a unique zero-pattern of polynomials $f_{2,1}, \dots, f_{k,t}$, we have $M \leq N$, where N is the number of zero-patterns of $f_{2,1}, \dots, f_{k,t}$ on the variety V_1 . Let $\tilde{f}_{i,j}$ be

the polynomial function induced on V_1 by the polynomial $f_{i,j}$, i.e. $\tilde{f}_{i,j} = f_{i,j} \pmod{I(V_1)}$, where we recall that $I(V_1)$ is the ideal of the variety V_1 (see Section 4.1).

Let $x_1, \dots, x_N \in V_1$ be points witnessing the N distinct zero-patterns, and for $j \in [N]$, let $S_j \subset \{\tilde{f}_{2,1}, \dots, \tilde{f}_{k,t}\}$ be the set of polynomial \tilde{f} for which $\tilde{f}(x_j) \neq 0$. Define the polynomial

$$g_j(x) = \prod_{\tilde{f} \in S_j} \tilde{f}(x).$$

Clearly, we have $g_j(x_j) \neq 0$, and

$$\deg g_j \leq \sum_{i,\ell} \deg \tilde{f}_{i,\ell} \leq kt^2.$$

Furthermore, note that $g_j(x_\ell) = 0$ if $S_j \not\subset S_\ell$, since for every $\tilde{f} \in S_j \setminus S_\ell$, we have $\tilde{f}(x_\ell) = 0$.

We now argue that the polynomials g_1, \dots, g_N are linearly independent. Otherwise, there exist scalars $\lambda_1, \dots, \lambda_N \in \overline{\mathbb{F}}$, not all zero, for which

$$\sum_{j=1}^n \lambda_j g_j = 0.$$

Choose ℓ such that $\lambda_\ell \neq 0$ and $|S_\ell|$ is minimal. If we plug $x = x_\ell$ into the above equation, all terms except $g_\ell(x_\ell)$ vanish. Indeed, for all $j \neq \ell$ we either have $\lambda_j = 0$ or $S_j \not\subset S_\ell$, meaning that $g_j(x_\ell) = 0$. But $g_\ell(x_\ell) \neq 0$, which is a contradiction. Thus, the polynomials g_1, \dots, g_N are linearly independent. But the dimension of the space of polynomials of degree at most kt^2 on the variety V_1 is bounded by $O_{t,D}(k^d)$ by Theorem 4.2, and so $M \leq N \leq O_{t,D}(k^d)$. \square

In the proof of Theorem 4.8, we also use the following result of [22], whose proof follows from the Haussler packing lemma [29].

Theorem 4.10 (Theorem 2.1 from [22]). *Let $c > 0$ and $r, s \in \mathbb{N}$, then there exists $c_1 = c_1(c, D, s) > 0$ such that the following holds. Let $G = (A, B, E)$ be a bipartite graph with $|A| = m$ and $|B| = n$ such that the set system $\mathcal{F}_1 = \{N(q) : q \in A\}$ satisfies $\pi_{\mathcal{F}_1}(k) \leq ck^r$ for every positive integer k . Then, if G is $K_{s,s}$ -free, we have*

$$|E(G)| \leq c_1(nm^{1-1/r} + m).$$

Proof of Theorem 4.8. Let $\mathcal{V} = \{V_1, \dots, V_n\}$, and let \mathcal{F} be the set of containment patterns of V_1, \dots, V_n . Then $\pi_{\mathcal{F}}(k) = O_{D,t}(k^{d+1})$ by Lemma 4.9. For every $x \in \mathcal{P}$, the neighborhood $N(x)$ corresponds to the containment pattern $I_x = \{i \in [n] : x \in V_i\}$. Therefore, if $\mathcal{F}_1 = \{N(x) : x \in \mathcal{P}\}$, then $\pi_{\mathcal{F}_1}(k) \leq \pi_{\mathcal{F}}(k) = O_{D,t}(k^{d+1})$. Applying Theorem 4.10 with $r = d + 1$, we get

$$I(\mathcal{P}, \mathcal{V}) = O_{D,t,s}(nm^{\frac{d}{d+1}} + m).$$

\square

5 Unit distances

5.1 Geometry of spheres over arbitrary fields

In this section, we discuss basic properties of spheres in finite fields, highlighting some differences with the real space. Throughout this section, we consider a non-degenerate bilinear form $\langle \cdot, \cdot \rangle : \mathbb{F}^d \rightarrow \mathbb{F}$, with respect to which we define the notions of orthogonality and distance. The form is non-degenerate if for every $v \neq 0$ there exists some w such that $\langle v, w \rangle \neq 0$. It is standard to associate a norm to a bilinear form by defining $\|v\|^2 = \langle v, v \rangle$. In what follows, a *unit sphere* with center w is defined to be the set of points $x \in \mathbb{F}^d$ for which $\|x - w\|^2 = 1$.

Furthermore, we say that vectors $u, v \in \mathbb{F}^d$ are *orthogonal* if $\langle u, v \rangle = 0$. One can generalize the notion to affine flats and say that $U, V \subseteq \mathbb{F}^d$ are *orthogonal* if for all points $u, u' \in U$ and $v, v' \in V$ one has $\langle u - u', v - v' \rangle = 0$. A key property of orthogonal subspaces we use is that they satisfy $\dim U + \dim V \leq d$. Furthermore, for vectors $v_1, \dots, v_k \in \mathbb{F}^d$, we denote the *affine span* of v_1, \dots, v_k by $\text{Aff}\{v_1, \dots, v_k\}$. More formally, we have $\text{Aff}\{v_1, \dots, v_k\} = \{\sum_{i=1}^d \lambda_i v_i \mid \lambda_i \in \mathbb{F}_q, \lambda_1 + \dots + \lambda_d = 1\}$.

Finally, we define a specific bilinear form $\langle u, v \rangle_d$ and the associated norm $\|v\|_d$ on the space \mathbb{F}_q^d as follows. When $d \not\equiv 1 \pmod{4}$, we let $\langle u, v \rangle_d$ be the standard inner product defined by $\langle u, v \rangle_d = u_1 v_1 + \dots + u_d v_d$, while for $d \equiv 1 \pmod{4}$ we define $\langle u, v \rangle_d = u_1 v_1 + \dots + u_{d-1} v_{d-1} - u_d v_d$. In both cases, we define the norm $\|v\|_d^2 = \langle v, v \rangle_d$.

Lemma 5.1. *Let $\langle \cdot, \cdot \rangle$ be a nondegenerate bilinear form on \mathbb{F}^d and S_1, \dots, S_k be unit spheres in \mathbb{F}^d , defined with respect to $\langle \cdot, \cdot \rangle$. Then there exists an affine subspace $U \subseteq \mathbb{F}^d$ such that $S_1 \cap \dots \cap S_k = S_1 \cap U$. Moreover, U is orthogonal to the affine flat spanned by the centers of S_1, \dots, S_k .*

Proof. We prove the statement by induction on k . Let us denote the centers of the spheres S_1, \dots, S_k by w_1, \dots, w_k .

When $k = 1$, we may take $U = \mathbb{F}^d$. For $k > 1$, we assume by our induction hypothesis that $S_1 \cap \dots \cap S_{k-1} = S_1 \cap U$ for some affine subspace $U \subseteq \mathbb{F}^d$, which is orthogonal to $\text{Aff}\{w_1, \dots, w_{k-1}\}$. Hence, we have $S_1 \cap \dots \cap S_{k-1} \cap S_k = S_1 \cap S_k \cap U$. Consider the defining equations for S_1 and S_k , which are

$$\langle x - w_1, x - w_1 \rangle = 1 \text{ and } \langle x - w_k, x - w_k \rangle = 1.$$

By subtracting these two equations, we have $2\langle x, w_k - w_1 \rangle + \langle w_1, w_1 \rangle - \langle w_k, w_k \rangle = 0$ for any $x \in S_1 \cap S_k$. Note that this equation is linear in x and hence it defines a hyperplane H , whose normal vector is $w_1 - w_k$. We conclude that $S_1 \cap S_k = S_1 \cap H$ and therefore $S_1 \cap \dots \cap S_k = S_1 \cap S_k \cap U = S_1 \cap (U \cap H)$, where $U \cap H$ is an affine space.

Let us now argue that $U \cap H$ is orthogonal to $\text{Aff}\{w_1, \dots, w_k\}$. Let $u, u' \in U \cap H$ and $v, v' \in \text{Aff}\{w_1, \dots, w_k\}$ be arbitrary vectors. One may write $v - v' = \sum_{i=2}^k \lambda_i (w_i - w_1)$ with suitable $\lambda_2, \dots, \lambda_k \in \mathbb{F}$. Therefore, we have

$$\langle u - u', v - v' \rangle = \left\langle u - u', \sum_{i=2}^{k-1} \lambda_i (w_i - w_1) \right\rangle + \langle u - u', \lambda_k (w_k - w_1) \rangle.$$

Note that the first term is zero since $u, u' \in U$, while the second term is zero because $u, u' \in H$. Hence, $U \cap H$ and $\text{Aff}\{w_1, \dots, w_k\}$ are orthogonal. \square

One counter-intuitive feature of spheres over finite fields is that they can contain flats. In the next lemma, we study some properties of such flats.

Lemma 5.2. *Let \mathbb{F} be a field of characteristic different from 2, let V be a flat, and let S be a unit sphere in \mathbb{F}^d centered at w . If $V \subseteq S$, then for any $x, y \in V$ we have $\langle x - y, x - y \rangle = 0$ and $\langle x - w, x - y \rangle = 0$.*

Proof. If $x, y \in V$, then $x + \lambda(y - x) \in V$ for every $\lambda \in \mathbb{F}$. Since $S = \{v \in \mathbb{F}^d : \langle v - w, v - w \rangle = 1\}$ and $V \subseteq S$, we must have

$$1 = \langle x + \lambda(y - x) - w, x + \lambda(y - x) - w \rangle = \langle x - w, x - w \rangle + 2\lambda \langle x - w, y - x \rangle + \lambda^2 \langle y - x, y - x \rangle.$$

Since $\langle x - w, x - w \rangle = 1$, we obtain $\lambda^2 \langle y - x, y - x \rangle + 2\lambda \langle x - w, y - x \rangle = 0$ for all λ , implying that $\langle y - x, y - x \rangle = 0$ and $\langle x - w, y - x \rangle = 0$, just as claimed. \square

A flat V for which $\langle x - y, x - y \rangle = 0$ for all $x, y \in V$ is called *totally isotropic*.

Lemma 5.3. *Let \mathbb{F} be a field containing no roots of the equation $x^2 = -1$ and let $d = 2k + 1$. Furthermore, consider the bilinear form $\langle \cdot, \cdot \rangle_d$ on \mathbb{F}^d . Then \mathbb{F}^d does not contain a pair (V, w) , where V is a totally isotropic flat of dimension k , and w is a vector of norm 1 orthogonal to V .*

Proof. Assume that \mathbb{F}^d contains such a pair (V, w) . After translation, we may assume that V contains the origin. Let v_1, \dots, v_k be a basis of the subspace V which satisfies $\langle v_i, e_j \rangle_d = \delta_{ij}$ for all $i, j \in [k]$, where e_1, \dots, e_d are the standard basis vectors. Here, δ_{ij} is defined to be 1 when $i = j$ and 0 otherwise. After possibly rearranging the coordinates, Gaussian elimination shows that it is always possible to find such a basis. Furthermore, since V is totally isotropic, any vector $w' = w + v$ with $v \in V$ is also of unit norm and orthogonal to V . Therefore, we may subtract an appropriate linear combination of v_1, \dots, v_k from w and assume that $\langle w, w \rangle_d = 1$ and $\langle w, e_j \rangle_d = 0$ for all $j \in [k]$.

To derive a contradiction, we compute the Gram matrix G of the vectors $v'_1 = v_1 - e_1, \dots, v'_k = v_k - e_k, w$. Note that the first k coordinates of each of these $k + 1$ vectors are equal to 0 and therefore one might consider v'_1, \dots, v'_k, w as vectors in \mathbb{F}^{k+1} without changing their Gram matrix. Hence, we have $G(i, j) = \langle v'_i, v'_j \rangle_d = \langle v_i, v_j \rangle_d - \langle e_i, v_j \rangle_d - \langle v_i, e_j \rangle_d + \langle e_i, e_j \rangle_d = -\delta_{ij}$ for all $i, j \in [k]$. Similarly, one can compute $G(i, k + 1) = G(k + 1, i) = \langle v'_i, w \rangle_d = \langle v_i, w \rangle_d - \langle e_i, w \rangle_d = 0$ and $G(k + 1, k + 1) = \langle w, w \rangle_d = 1$. In conclusion, $G = \text{diag}[-1, \dots, -1, 1]$.

On the other hand, let A be a $(k + 1) \times (k + 1)$ the matrix whose columns are the vectors v'_1, \dots, v'_k, w (recall that we may consider these vectors as elements of \mathbb{F}^{k+1}). Recalling the definition of the bilinear form $\langle v'_i, v'_j \rangle_d$, we note that we can rewrite it as $\langle v'_i, v'_j \rangle_d = (v'_i)^T M v'_j$, where $M = I_{k+1}$ if $d \equiv 3 \pmod{4}$ and $M = \text{diag}[1, \dots, 1, -1]$ if $d \equiv 1 \pmod{4}$. Hence, one can express the Gram matrix as $G = A^T M A$ and so $A^T M A = \text{diag}[-1, \dots, -1, 1]$. By taking determinants, we see that $\det(A)^2 \det(M) = \det(\text{diag}[-1, \dots, -1, 1]) = (-1)^k$.

The final observation is that $\det(M) = -1$ when $d \equiv 1 \pmod{4}$ and $\det(M) = 1$ otherwise, meaning that $\det(M) = (-1)^{k-1}$. Hence, no matter the parity of k , we have $\det(A)^2 = -1$. But this is impossible, since the equation $x^2 = -1$ has no solution in \mathbb{F} . \square

5.2 Proofs of the upper bounds

In this section, we prove Theorem 1.6. In particular, we prove the following incidence bound between points and unit spheres defined with respect to any nondegenerate bilinear form, from which the theorem immediately follows.

Proposition 5.4. *Let \mathcal{P} be a set of n points and let \mathcal{S} be a collection of n unit spheres in \mathbb{F}^d . If the incidence graph $G(\mathcal{P}, \mathcal{S})$ is $K_{s,s}$ -free, then the number of incidences between \mathcal{P} and \mathcal{S} is at most $O_{d,s}(n^{2 - \frac{1}{\lceil d/2 \rceil + 1}})$.*

The proof follows the ideas from Section 2, where we showed that incidence graphs of points and hyperplanes in \mathbb{F}^d avoid the pattern Π_d . Here, we will show that the point-sphere incidence graphs in \mathbb{F}^d avoid the pattern Π_{d+1} and use this to deduce Proposition 5.4.

Lemma 5.5. *Let \mathcal{P} be a set of n points and let \mathcal{S} be a set of n spheres in \mathbb{F}^d . The incidence graph $G(\mathcal{P}, \mathcal{S})$ does not contain the pattern Π_{d+1} defined above.*

Proof. For contradiction, let us assume that $G(\mathcal{P}, \mathcal{S})$ does contain Π_{d+1} , whose vertices a_1, \dots, a_{d+1} correspond to points P_1, \dots, P_{d+1} and b_1, \dots, b_{d+1} correspond to unit spheres S_1, \dots, S_{d+1} .

We show by induction that $S_1 \cap \dots \cap S_k = S_1 \cap U$ for a affine flat U of dimension at most $d - k + 1$, for all $k \leq d + 1$. For $k = 1$ this statement is trivial. For $k \geq 2$, we may assume that $S_1 \cap \dots \cap S_{k-1} = S_1 \cap U$ for some flat U of dimension at most $d - k + 2$. By Lemma 5.1, there exists a flat H for which $S_1 \cap S_k = S_1 \cap H$. Let $V = U \cap H$. Then $S_1 \cap \dots \cap S_k = S_1 \cap U \cap S_k = S_1 \cap (U \cap H) = S_1 \cap V$. The point P_{k-2} is contained in the spheres S_1, \dots, S_{k-1} , but not in S_k . Therefore, P_{k-2} is a point of U not contained in V , implying that $\dim(V) \leq \dim(U) - 1 \leq d - k + 1$.

Applying this claim with $k = d + 1$, we conclude that there exists a 0-dimensional affine space U for which $S_1 \cap \dots \cap S_{d+1} = S_1 \cap U$. In other words, $S_1 \cap \dots \cap S_{d+1}$ is a single point. Then, it is not possible for both P_d and P_{d+1} to be contained in $S_1 \cap \dots \cap S_{d+1}$, presenting a contradiction. \square

Proposition 5.6. *Let G be a bipartite graph on $2n$ vertices which does not contain $K_{s,s}$ or the pattern Π_{d+1} . Then, G has at most $O_{d,s}(n^{2 - \frac{1}{\lceil d/2 \rceil + 1}})$ edges.*

Proof. Lemma 3.2 it is proved that a balanced bipartite graph on $2n$ vertices containing no $K_{s,s}$ or Π_{d+1} has at most $O_{d,s}\left(n^{2-\frac{1}{\lceil(d+2)/2\rceil}}\right)$. \square

5.3 Constructions

In this section, our goal is to prove Theorem 1.7. More precisely, for every n and d , we construct a set \mathcal{P} of n points in \mathbb{F}_q^d , where $q = p^r$ is some prime power, with the properties that \mathcal{P} spans $\Omega(n^{2-\frac{1}{\lceil d/2\rceil+1}})$ unit distances and that the unit distance graph does not contain $K_{s,s}$.

Our construction is based on so called variety-evasive sets. For fixed parameters $1 \leq k \leq d$ and $\Delta, s > 0$, we say that a set U is (k, s) -variety-evasive if every variety $V \subseteq \mathbb{F}_q^d$ of dimension k and degree at most Δ intersects U in less than s points. Here we only consider varieties of degree at most 2 and therefore we fix $\Delta = 2$. Dvir, Kollár and Lovett constructed large variety-evasive sets and in particular the following theorem is a special case of Corollary 6.1 from [15].

Theorem 5.7. *For every positive integer d and prime power q , there exist a constant $s = s(d)$ and set U of q^{d-k} points in \mathbb{F}_q^d such that any k -dimensional variety of degree at most 2 intersects U in less than s points.*

Let us now explain the idea behind our constructions. We set $k = \lfloor d/2 \rfloor$ and choose p to be a prime with $p \equiv 3 \pmod{4}$ and $p \approx n^{\frac{1}{\lceil d/2\rceil+1}}$. Then, we use Theorem 5.7 to find a $(k-1, s)$ -variety evasive set $U \subseteq \mathbb{F}_p^d$ of size $p^{d-k+1} = p^{\lceil d/2\rceil+1}$. Since U is $(k-1, s)$ -variety evasive, we are able to show that the unit distance graph does not contain $K_{s,s}$ as a subgraph. However, in order to ensure many unit distances, we need to take the union of U with a random shift, i.e. we set $\mathcal{P}_x = U \cup (U+x)$ where $x \in \mathbb{F}_p^d$ is chosen uniformly at random.

Proposition 5.8. *There exists $x \in \mathbb{F}_p^d$ for which \mathcal{P}_x spans $\Omega\left(|\mathcal{P}_x|^{2-\frac{1}{\lceil d/2\rceil+1}}\right)$ unit distances with respect to the norm $\|\cdot\|_d$.*

Proof. Choose x from the uniform distribution on \mathbb{F}_p^d . The number of unit distances spanned by \mathcal{P}_x is at least the number of pairs $u, v \in U$ such that $\|u - (v+x)\|_d = 1$. For a fixed pair $u, v \in U$, the equation $\|u - v - x\|_d = 1$ is quadratic in x and therefore it has $(1+o(1))p^{d-1}$ solutions (see e.g. Theorems 6.26 and 6.27 in [34]). Hence, for fixed $u, v \in U$ we have $\mathbb{P}[\|u - (v+x)\|_d = 1] = (1+o(1))p^{-1}$ and so the expected number of pairs $u, v \in U$ with $\|u - (v+x)\|_d = 1$ is $(1+o(1))|U|^2/p$. Therefore, x can be chosen such that the number of unit distances spanned by \mathcal{P}_x is $\Omega(|U|^2/p) = \Omega\left(|\mathcal{P}_x|^{2-\frac{1}{\lceil d/2\rceil+1}}\right)$. \square

Let us fix the value x from Proposition 5.8 from now on and write $\mathcal{P} = \mathcal{P}_x$. Note that the set \mathcal{P} is $(k-1, 2s)$ -variety evasive, since no variety of dimension $k-1$ and degree at most 2 can contain more than s points of either U or $U+x$.

Proposition 5.9. *The unit distance graph of \mathcal{P} with respect to the norm $\|\cdot\|_d$ does not contain $K_{4s,4s}$ as a subgraph.*

Proof. We argue by contradiction, assuming that there exist points $v_1, \dots, v_{4s}, w_1, \dots, w_{4s} \in \mathcal{P}$ such that $\|v_i - w_j\|_d = 1$ for all $i, j \in [4s]$. Let $V = \text{Aff}\{v_1, \dots, v_{4s}\}$ and $W = \text{Aff}\{w_1, \dots, w_{4s}\}$.

We claim that the flats V and W are orthogonal, with respect to the bilinear form $\langle \cdot, \cdot \rangle_d$. If we denote by S_1, \dots, S_{4s} the unit spheres centered at w_1, \dots, w_{4s} , the assumption implies that $v_1, \dots, v_{4s} \in \bigcap_{i=1}^{4s} S_i$. By Lemma 5.1, there exists an affine flat V' , orthogonal to W , such that $\bigcap_{j=1}^{4s} S_j = S_1 \cap V'$. Since $V \subseteq V'$, we conclude V and W are orthogonal, which implies that $\dim V + \dim W \leq d$.

Since \mathcal{P} is $(k-1, 2s)$ -variety evasive, no $(k-1)$ -dimensional flat contains $2s$ points of \mathcal{P} (note that an affine flat is a variety of degree 1). In particular, this means that $\dim V \geq k$ and $\dim W \geq k$. Combining this with $\dim V + \dim W \leq d = 2k+1$, we conclude that $\dim V = k$ or $\dim W = k$. By symmetry, we may assume that $\dim V = k$.

We claim that $V \subseteq S_j$, for every $j \in [4s]$. Suppose this was not the case, i.e. that there was some j for which $V \not\subseteq S_j$. Since V is a variety of dimension k and degree 1 and S_j is a variety of dimension $d-1$ and degree 2, Theorem 4.1 implies that if $V \cap S_j$ is a union of irreducible components Z_1, \dots, Z_t , then $\sum_{i=1}^t \deg Z_i \leq 2$. Hence, $V \cap S_j$ consists either of one $(k-1)$ -dimensional component of degree 2 or of at most two $(k-1)$ -dimensional components of degree 1. However, recall that $v_1, \dots, v_{4s} \in S_j \cap V$ and that every $(k-1)$ -dimensional variety of degree at most 2 contains less than $2s$ points of \mathcal{P} . Hence, in both cases we have a contradiction and so we must have $V \subseteq S_j$. Therefore, by Lemma 5.2, V must be a totally isotropic flat and we have $\langle w_j - y, z - y \rangle_d = 0$ for all $y, z \in V$.

Since $\langle w_j - v_i, w_j - v_i \rangle_d = 1$, the point w_j does not lie in the flat V . Also, $\text{Aff}(\{w_j\} \cup V)$ is orthogonal to V , since $\langle w_j - x, y - x \rangle_d = 0$ for all $x, y \in V$. Hence, $\dim \text{Aff}(\{w_j\} \cup V) = \dim V + 1 = k + 1$. If d is even, this is a contradiction since $\dim V + \dim \text{Aff}(\{w_j\} \cup V) = 2k + 1 > d$.

If d is odd, we found a totally isotropic k -flat with a vector $w_j - v_i$ which has norm 1 and is orthogonal to this k -flat. But this is impossible by Lemma 5.3. \square

Let us now put all the ingredients together and show the general construction.

Proof of Theorem 1.7. Let p be the smallest prime such that $p \equiv 3 \pmod{4}$ and $p > n^{\frac{1}{\lceil d/2 \rceil + 1}}$. By a theorem of Breusch [7], we have $p \leq 2n^{\frac{1}{\lceil d/2 \rceil + 1}}$, as long as $n^{\frac{1}{\lceil d/2 \rceil + 1}} \geq 7$. Let \mathcal{P} be the set of at least $p^{\lceil d/2 \rceil + 1} \geq n$ points constructed in this section. For $d \not\equiv 1 \pmod{4}$, taking a random n element subset of \mathcal{P} gives the desired set. However, for $d \equiv 1 \pmod{4}$ the standard norm and the norm $\|\cdot\|_d$ are different as we need one final step to complete the construction.

If $d \equiv 1 \pmod{4}$, we let $q = p^2$, we choose $\alpha \in \mathbb{F}_q$ to be a solution of $\alpha^2 = -1$, and consider the map $\varphi : \mathbb{F}_p^d \rightarrow \mathbb{F}_q^d$ given by $\varphi((x_1, \dots, x_d)) = (x_1, \dots, x_{d-1}, \alpha x_d)$. Note that $\varphi(u)$ and $\varphi(v)$ are at unit distance measured with respect to the standard inner product if and only if $\|u - v\|_d = 1$. Hence, we conclude that the set $\{\varphi(u) | u \in \mathcal{P}\} \subseteq \mathbb{F}_q^d$ spans $\Omega\left(|\mathcal{P}|^{2 - \frac{1}{\lceil d/2 \rceil + 1}}\right)$ unit distances and has no $K_{s,s}$ in the unit distance graph, where unit distances are measured with respect to the standard inner product. Now taking a random n element subset of \mathcal{P} completes the proof. \square

6 The algebraic Zarankiewicz problem

In this section, we present the proof of our bound on Zarankiewicz's problem in algebraic graphs of bounded complexity. We prove the upper and lower bounds separately, and then we use the lower bound to show Theorem 1.5. In the proof of the upper bound, we combine Theorem 4.10 with the estimate (4).

For the lower bounds, we use the random polynomial method, which was pioneered by Bukh [8]. The main idea is to choose the polynomial defining the algebraic graph to be a random low-degree polynomial, chosen uniformly among all polynomials of degree at most Δ in $\mathbb{F}[x_1, \dots, x_D]$, where \mathbb{F} is a finite field of appropriate size. The key observation is that for any fixed $x_1, \dots, x_s \in \mathbb{F}^D$ and random polynomial f of degree at most Δ , the random variables $\{f(x_j) | j \in [s]\}$ are independent and uniform in \mathbb{F} as long as $s \leq \min\{\Delta, |\mathbb{F}|^{1/2}\}$ (see Claims 3.3 and 3.4 in [51]).

Theorem 6.1. *Let G be an algebraic graph of description complexity at most t on the vertex set $\mathcal{P} \cup \mathcal{Q}$, where $\mathcal{P} \subseteq \mathbb{F}^{d_1}$, $\mathcal{Q} \subseteq \mathbb{F}^{d_2}$ and $|\mathcal{P}| = m$, $|\mathcal{Q}| = n$. If G is $K_{s,s}$ -free, the number of edges in G is at most*

$$O_{d_1, d_2, t, s}(\min\{m^{1-1/d_1}n, mn^{1-1/d_2}\}).$$

Proof. Let the graph G be defined by the polynomials $f_1, \dots, f_t : \mathbb{F}^{d_1} \times \mathbb{F}^{d_2} \rightarrow \mathbb{F}$ and a boolean formula Φ , where $\deg f_i \leq t$ for $i = 1, \dots, t$.

Since the roles of m and n are symmetric, it suffices to show that any $K_{s,s}$ -free algebraic graph G in $(\mathbb{F}^{d_1}, \mathbb{F}^{d_2})$ has at most $O_{t,s}(mn^{1-1/d_2} + n)$ edges. To do this, we use Theorem 4.10. Define the set system $\mathcal{F} = \{N(v) : v \in \mathcal{Q}\}$ on the ground set \mathcal{P} . Our main goal is to show $\pi_{\mathcal{F}}(k) \leq O_{t,d_2}(k^{d_2})$ for all k , since in this case Theorem 4.10 shows that $e(G) \leq O_{t,s}(mn^{1-1/d_2} + n)$. To this end, we

fix k points $x_1, \dots, x_k \in \mathcal{P}$ and bound the number of possible intersections $\{x_1, \dots, x_k\} \cap N(v)$ for $v \in \mathcal{Q}$. For every $i \in [k]$ and $j \in [t]$, consider the polynomial $y \mapsto f_j(x_i, y)$ over \mathbb{F}^{d_2} . Let F be the set of these kt polynomials.

The number of distinct intersections $\{x_1, \dots, x_k\} \cap N(v)$ is bounded by the number of zero-patterns of polynomials in F , since the zero-pattern of polynomials in F at a point $v \in \mathbb{F}$ can be used to determine which x_i is adjacent to v . Since F is a set of kt polynomials of degree at most t , by (4) we have that the number of zero-patterns of the polynomials in F is bounded by $\binom{kt^2}{d_2}$. Therefore, $\pi_{\mathcal{F}}(k) \leq \binom{kt^2}{d_2} = O_{t,d_2}(k^{d_2})$, finishing the proof. \square

To show the lower bounds, we will need the following simple lemma.

Lemma 6.2. *Let $\Delta, D \geq 3$ be fixed integers and let $p \geq 4$ be a prime. Let $f \in \mathbb{F}_p[x_1, \dots, x_D]$ be chosen uniformly at random among polynomials of degree at most Δ , then f has at least $p^{D-1}/2$ zeros in \mathbb{F}_p^D with probability at least $\frac{3}{4}$.*

Proof. Since $f(x)$ is a uniformly distributed in \mathbb{F}_p for any fixed $x \in \mathbb{F}_p^D$, the expected number of zeros of f is $\mathbb{E}[|V(f)|] = p^{D-1}$. On the other hand, using the second moment method, one can also show that the random variable $|V(f)|$ is concentrated around its mean. More precisely, if we denote by $\mathbf{1}_x$ the indicator random variable of $f(x) = 0$, we have

$$\text{Var}(|V(f)|) = \sum_{x,y \in \mathbb{F}^D} \text{Cov}(\mathbf{1}_x, \mathbf{1}_y).$$

Claim 3.4 of [51] states that $f(x)$ and $f(y)$ are independent random variables when $x \neq y$, if $p \geq 4$ and $\Delta \geq 2$. Moreover, for $x = y$, we have $\text{Cov}(\mathbf{1}_x, \mathbf{1}_x) \leq \mathbb{P}[f(x) = 0] = p^{-1}$. Hence, $\text{Var}(|V(f)|) \leq p^{D-1}$. By Chebyshev's inequality, we have $\mathbb{P}[|V(f)| < p^{D-1} - \lambda \sqrt{p^{D-1}}] \leq \frac{1}{\lambda^2}$, which for $\lambda = \frac{1}{2} \sqrt{p^{D-1}}$ gives

$$\mathbb{P}\left[|V(f)| < \frac{p^{D-1}}{2}\right] \leq \frac{4}{p^{D-1}} \leq \frac{1}{4}.$$

\square

Proposition 6.3. *For any d_1, d_2 and sufficiently large integers m, n , there exists a field \mathbb{F} , sets of points $\mathcal{P} \subseteq \mathbb{F}^{d_1}, \mathcal{Q} \subseteq \mathbb{F}^{d_2}$, $|\mathcal{P}| = m, |\mathcal{Q}| = n$ and a polynomial $f : \mathbb{F}^{d_1} \times \mathbb{F}^{d_2} \rightarrow \mathbb{F}$ of degree at most $\Delta = (d_1 + d_2)^2$ with the following property. The algebraic graph G defined on $\mathcal{P} \cup \mathcal{Q}$ by the equation $f(x, y) = 0$ is $K_{s,s}$ -free, where $s = d_1 + d_2$, and has $\Omega(\min\{m^{1-1/d_1}n, mn^{1-1/d_2}\})$ edges.*

Observe that the graph G has description complexity at most $(d_1 + d_2)^2$ and therefore Proposition 6.3 suffices to show the second part of Theorem 1.8.

Proof. By symmetry, we may assume that $mn^{1-1/d_2} < m^{1-1/d_1}n$, which is equivalent to $m \leq n^{d_1/d_2}$. Let p be the smallest prime larger than n^{1/d_2} , which satisfies $p < 2n^{1/d_2}$ by Bertrand's postulate.

We define an algebraic graph G_0 in $(\mathbb{F}_p^{d_1}, \mathbb{F}_p^{d_2})$ as follows. Let $\mathcal{P}_0 = \mathbb{F}_p^{d_1}$ and $\mathcal{Q}_0 = \mathbb{F}_p^{d_2}$. Furthermore, let $f : \mathbb{F}_p^{d_1} \times \mathbb{F}_p^{d_2} \rightarrow \mathbb{F}_p$ be a polynomial chosen randomly from the uniform distribution on all polynomials of degree at most $\Delta = (d_1 + d_2)^2$. Finally, set $x \in \mathcal{P}_0$ and $y \in \mathcal{Q}_0$ to be adjacent in G_0 if and only if $f(x, y) = 0$.

We begin by showing that the probability G_0 contains $K_{s,s}$ is at most $\frac{1}{4}$. Given $2s$ points $x^{(1)}, \dots, x^{(s)} \in \mathcal{P}_0, y^{(1)}, \dots, y^{(s)} \in \mathcal{Q}_0$, the s^2 points $(x^{(i)}, y^{(j)})$ are independent uniform random variables as long as $s^2 \leq \min\{\Delta, p^{1/2}\}$, see [8]. Therefore, the probability that $f(x^{(i)}, y^{(j)}) = 0$ for all $i, j \in [s]$ is exactly p^{-s^2} . Hence, we can apply the union bound over all subsets of \mathcal{P}_0 and \mathcal{Q}_0 of size s to deduce

$$\mathbb{P}[K_{s,s} \subseteq G] \leq \binom{p^{d_1}}{s} \binom{p^{d_2}}{s} p^{-s^2} \leq \frac{p^{d_1 s}}{s!} \frac{p^{d_2 s}}{s!} p^{-s^2} \leq \frac{1}{(s!)^2} \leq \frac{1}{4}.$$

On the other hand, Lemma 6.2 shows that with probability at least $3/4$, the polynomial f has at least $p^{d_1+d_2-1}/2$ zeros, which means that G_0 has at least $p^{d_1+d_2-1}/2$ edges with probability at

least $3/4$. Therefore, there exists a polynomial f such that the graph G_0 is $K_{s,s}$ -free and has at least $\frac{1}{2p}|\mathcal{P}_0||\mathcal{Q}_0|$ edges. Let us fix this graph G_0 .

Then we get our final graph G by sampling an m element subset \mathcal{P} of \mathcal{P}_0 , and an n element subset \mathcal{Q} of \mathcal{Q}_0 . Then, the expected number of edges of G is $\frac{|\mathcal{P}|}{|\mathcal{P}_0|} \frac{|\mathcal{Q}|}{|\mathcal{Q}_0|} e(G_0) \geq \frac{mn}{2p}$. Hence, there exists a choice of \mathcal{P}, \mathcal{Q} for which $e(G) \geq \frac{mn}{2p}$, let us fix this choice. This graph G is clearly algebraic and $K_{s,s}$ -free, since it is an induced subgraph of a $K_{s,s}$ -free algebraic graph. Furthermore, $e(G) \geq \frac{mn}{2p} \geq \frac{mn}{2}(2n)^{-1/d_2} = \Omega(mn^{1-1/d_2})$. This completes the proof. \square

We conclude the section by proving Theorem 1.5, which demonstrates that our upper bounds on point-variety incidences are tight.

Proof of Theorem 1.5. Consider the case $d = D - 1$. The general case follows by embedding \mathbb{F}_p^{d+1} into a higher dimensional space. Let p be a prime between $m^{1/D}$ and $2m^{1/D}$, which exists by Bertrand's postulate. Further, let $D' = \lceil \alpha D \rceil$, and let G be an algebraic graph in $(\mathbb{F}_p^D, \mathbb{F}_p^{D'})$ with parts \mathcal{P}, \mathcal{Q} of size m, n obtained through Proposition 6.3.

We construct a set of hypersurfaces \mathcal{V} from the set of points \mathcal{Q} . For each point $q \in \mathcal{Q}$, we consider the polynomial $f_q : \mathbb{F}_p^D \rightarrow \mathbb{F}_p$ given by $f_q(x) = f(x, q)$. Since $f_q(x)$ may not be irreducible, let $g_q(x)$ be an irreducible factor of $f_q(x)$ with the largest number of roots in \mathcal{P} . Finally, define the algebraic set $V_q = \{x \in \mathbb{F}_p^D | g_q(x) = 0\}$. We show that the point set \mathcal{P} and the collection $\mathcal{V} = \{V_q | q \in \mathcal{Q}\}$ satisfy all conditions of the theorem.

Since $g_q(x)$ is an irreducible polynomial, the ideal $\langle g_q(x) \rangle \subseteq \mathbb{F}_p[x_1, \dots, x_D]$ is a prime ideal, and therefore by Proposition 3, page 207 of [11], the algebraic set defined by $g_q(x)$ is irreducible and therefore a variety. Furthermore, the dimension of the variety $V_q = \{x \in \mathbb{F}_p^D | g_q(x) = 0\}$ is $D - 1 = d$. Hence, \mathcal{V} is indeed a collection of varieties of required dimension.

Let us now argue that there are no distinct varieties $V_{q_1}, \dots, V_{q_s} \in \mathcal{V}$ with at least s points of \mathcal{P} in common, where $s = (D + D')^2$. If such s varieties exist, the common neighbourhood of points q_1, \dots, q_s in the graph G contains at least s points, which is not possible since G is $K_{s,s}$ -free. Hence, the incidence graph $G(\mathcal{P}, \mathcal{V})$ is $K_{s,s}$ -free.

Finally, we argue that there are many incidences between \mathcal{P} and \mathcal{V} , i.e. that $I(\mathcal{P}, \mathcal{V}) \geq \Omega_{D,\alpha}(m^{\frac{D-1}{D}}n)$. The main observation is that the polynomials $f_q(x)$ have at most $(D + D')^2$ irreducible factors, since $\deg f_q(x) \leq (D + D')^2$. Hence, the number of roots of $g_q(x)$ in \mathcal{P} is at least $\frac{1}{(D+D')^2}$ times the number of roots of $f_q(x)$ in \mathcal{P} . But the number of roots of $f_q(x)$ in \mathcal{P} summed over all $q \in \mathcal{Q}$ is the number of edges of G and therefore we have:

$$I(\mathcal{P}, \mathcal{V}) = \sum_{q \in \mathcal{Q}} |V_q \cap \mathcal{P}| \geq \sum_{q \in \mathcal{Q}} \frac{|\{f_q(x) = 0\} \cap \mathcal{P}|}{(D + D')^2} \geq \frac{e(G)}{(D + \lceil \alpha D \rceil)^2} \geq \Omega_{D,\alpha} \left(\min\{m^{\frac{D-1}{D}}n, mn^{\frac{D'-1}{D'}}\} \right).$$

It is not hard to see that $m^{-1/D} \leq n^{-1/D'}$ since $m^{D'} = m^{\lceil \alpha D \rceil} \geq m^{\alpha D} = n^D$. Therefore, $I(\mathcal{P}, \mathcal{V}) \geq \Omega_{D,\alpha} \left(m^{\frac{D-1}{D}}n \right)$, which suffices to complete the proof. \square

7 Concluding remarks

In this paper, we proposed a novel combinatorial approach to study incidence problems over arbitrary fields. We considered several natural problems, such as determining the maximum number of point-hyperplane incidences, point-variety incidences, and unit distances, under standard non-degeneracy assumptions.

Our bounds are also optimal in the sense that they cannot be improved unless assumptions on the underlying field are made. Hence, it would be interesting to decide whether the techniques in this paper can be combined with some field-dependent information to improve some of our bounds. Such improved bounds are known over the real numbers [2, 22], and also in some cases over finite fields [5, 43]. For example, Rudnev's theorem [43] shows that when $n \leq p^2$, there are at most $O(m\sqrt{n} + ms)$ incidences between m points and n planes in \mathbb{F}_p^3 , assuming the incidence graph is $K_{2,s}$ -free.

Finally, it would be also interesting to find other geometric problems for which the methods of this paper can be applied. For example, it may be possible to use our approach to address other non-degeneracy conditions beyond forbidding $K_{s,s}$ in the incidence graph.

Acknowledgements. We would like to thank Larry Guth for suggesting to study point-variety incidences and stimulating discussions. We would also like to thank Joshua Zahl and Misha Rudnev for pointing out additional references.

References

- [1] E. Aksoy Yazici, B. Murphy, M. Rudnev and I. Shkredov. *Growth estimates in positive characteristic via collisions*, Int. Math. Res. Not. 23 (2017): 7148–7189.
- [2] R. Apfelbaum, and M. Sharir. *Large Complete Bipartite Subgraphs In Incidence Graphs Of Points And Hyperplanes*. SIAM Journal on Discrete Mathematics 21.3 (2007): 707–725.
- [3] M. Balko, J. Cibulka, and P. Valtr, *Covering Lattice Points by Subspaces and Counting Point-Hyperplane Incidences*, Discrete Computational Geometry 61 (2019): 325–354.
- [4] S. Basu, M. Sombra. *Polynomial partitioning on varieties and point-hypersurface incidences in four dimensions*. Discrete Comput. Geom. 55 (2016): 158–184.
- [5] J. Bourgain, N. Katz, and T. Tao. *A sum-product estimate in finite fields, and applications*. Geom. funct. anal. 14 (2004): 27–57.
- [6] P. Brass, and C. Knauer. *On counting point-hyperplane incidences*. Comput. Geom. Theory Appl. 25 (2003): 13–20.
- [7] R. Breusch. *Zur Verallgemeinerung des Bertrandischen Postulates dass zwischen x und $2x$ stets Primzahlen liegen*, Math. Z. 34 (1932): 505–526.
- [8] B. Bukh. *Random algebraic construction of extremal graphs*. Bulletin of London Math. Soc. 47(6) (2015): 939–945.
- [9] B. Chazelle. *Cutting hyperplanes for divide-and-conquer*. Discrete Comput. Geom. 9 (1993): 145–158.
- [10] K. L. Clarkson. *New applications of random sampling in computational geometry*. Discrete Comput. Geom. 2 (1987): 195–222.
- [11] D. Cox, J. Little, and D. O’Shea. *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Communitative Algebra*. Fourth ed. New York: Springer, 2015.
- [12] T. Do. *The Largest Complete Bipartite Subgraph in Point-Hyperplane Incidence Graphs*. The Electronic Journal of Combinatorics (2020): 27.1.
- [13] Z. Dvir. *Incidence Theorems and Their Applications*. Foundations and Trends in Theoretical Computer Science 6 (2012): 257–393.
- [14] Z. Dvir, A. Gabizon, and A. Wigderson. *Extractors And Rank Extractors For Polynomial Sources*. Comput. Complex. 18 (2009): 1–58.
- [15] Z. Dvir, J. Kollár, and S. Lovett. *Variety Evasive Sets*. Comput. Complex. 23 (2014): 509–529.
- [16] Z. Dvir, and S. Lovett. *Subspace evasive sets*. in Proceedings of the forty-fourth annual ACM symposium on Theory of computing (2012): 351–358.

- [17] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. volume 10 of EATCS Monographs on Theoretical Computer Science. Springer-Verlag, 1987.
- [18] G. Elekes. *On the number of sums and products*. Acta Arithmetica 81.4 (1997): 365–367.
- [19] P. Erdős. *On sets of distances of n points*. Amer. Math. Monthly 53 (1946): 248–250.
- [20] P. Erdős. *On sets of distances of n points in Euclidean space*. Magyar Tud. Akad. Mat. Kutató Int. Közl. 5 (1960): 165–169.
- [21] J. Erickson. *New lower bounds for Hopcroft’s problem*. Discret. Comput. Geom., 16(4) (1996): 389–418.
- [22] J. Fox, J. Pach, A. Sheffer, A. Suk, and J. Zahl. *A semi-algebraic version of Zarankiewicz’s problem*. J. Eur. Math. Soc. 19 (2017): 1785–1810.
- [23] J. Fox, J. Pach, and A. Suk. *A Polynomial Regularity Lemma for Semialgebraic Hypergraphs and Its Applications in Geometry and Property Testing*. SIAM Journal on Computing 45.6 (2016): 2199–2223.
- [24] J. Fox, and B. Sudakov. *Dependent random choice*. Random Structures and Algorithms 38 (2011): 1–32.
- [25] N. Frankl, and A. Kupavskii. *On the Erdős-Purdy problem and the Zarankiewicz problem for semialgebraic graphs*. preprint, arXiv:2112.10245.
- [26] W. Fulton, Intersection Theory, 2nd ed. New York, NY: Springer, 1998.
- [27] L. Guth, and N. H. Katz. *On the Erdős distinct distances problem in the plane*. Annals of Math. 181 N1 (2015): 155–190.
- [28] R. Hartshorne. *Algebraic Geometry*. 1st ed. New York, NY: Springer, 1997.
- [29] D. Haussler. *Sphere packing numbers for subsets of the boolean n -cube with bounded Vapnik-Chervonenkis dimension*. 69 (2) (1995): 217–232.
- [30] Z. Hunter, A. Milojević, B. Sudakov, and I. Tomon. *Kővári-Sós-Turán theorem for hereditary families*. preprint, arXiv:2401.10853 (2024).
- [31] J. Kollár. *Szemerédi–Trotter-type theorems in dimension 3*. Advances in Mathematics 271, (2015): 30–61.
- [32] T. Kővári, V. Sós, and P. Turán. *On a problem of K. Zarankiewicz*. Colloq. Math. 3 (1954): 50–57.
- [33] H. Lenz. *Zur Zerlegung von Punktmengen in solche kleineren Durchmessers*. Arch. Math. 5 (1955): 413–416.
- [34] R. Lidl, and H. Niederreiter. Finite Fields. Second ed. Cambridge: Cambridge University Press, 1997.
- [35] L. Lovász, and M. Saks. *Lattices, Möbius functions and communication complexity*. Annual Symposium on Foundations of Computer Science (1988): 81–90.
- [36] S. Lovett. *Communication Is Bounded by Root of Rank*. In: Journal of the ACM (2016) 63.1. Conference version in STOC (2014): 1–9.
- [37] A. Mohammadi, and S. Stevens. *Attaining the Exponent $5/4$ for the Sum-Product Problem in Finite Fields*, Int. Math. Res. Not. (2023): 3516–3532.

- [38] J.V. Nesterenko, and F. Amoroso. *Introduction to Algebraic Independence Theory*. Berlin: Springer, 2001.
- [39] J. Pach, M. Sharir. *Repeated angles in the plane and related problems*. J. Combin. Theory Ser. A 59 (1992): 12–22.
- [40] J. Pach, M. Sharir. *On the number of incidences between points and curves*. Combin. Probab. Comput. 7 (1998): 121–127.
- [41] O. Roche-Newton, M. Rudnev, and I. D. Shkredov. *New sum-product type estimates over finite fields*. Advances in Mathematics 293 (2016): 589–605.
- [42] L. Rónyai, L. Babai, and M. K. Ganapathy. *On the number of zero-patterns of a sequence of polynomials*. J. Amer. Math. Soc. 14 (2001): 717–735
- [43] M. Rudnev. *On the number of incidences between points and planes in three dimensions*. Combinatorica 38 (2018): 219–254.
- [44] M. Rudnev, *On incidences of lines in regular complexes*. Eur. J. Comb. 98 (2021).
- [45] A. Sheffer. *Lower bounds for incidences with hypersurfaces*. Discrete Anal. 1 (16) (2016): 14pp.
- [46] N. Singer, and M. Sudan. *Point-Hyperplane Incidence Geometry and the Log-Rank Conjecture*. preprint, arXiv:2101.09592.
- [47] J. Solymosi, and T. Tao. *An Incidence Theorem in Higher Dimensions*, Discrete Comput. Geom. 48 (2012): 255–280.
- [48] J. Spencer, E. Szemerédi, and W. Trotter. *Unit distances in the Euclidean plane*. In Graph Theory and Combinatorics (ed., B. Bollobás), Academic Press, 1984, 293–308.
- [49] S. Stevens, and F. de Zeeuw. *An improved point-line incidence bound over arbitrary fields*, Bull. London Math. Soc., 49 (2017): 842–858.
- [50] B. Sudakov, and I. Tomon. *Ramsey properties of algebraic graphs and hypergraphs*. Forum of Mathematics, Sigma 10 (2022): e95.
- [51] B. Sudakov, and I. Tomon. *Evasive sets, covering by subspaces, and point-hyperplane incidences*. Discrete Comput. Geom. (2023). doi:10.1007/s00454-023-00601-1
- [52] E. Szemerédi, and W. T. Trotter. *Extremal problems in discrete geometry*. Combinatorica 3(3-4) (1983): 381–392.
- [53] T. Tao. *Higher Order Fourier Analysis*. Providence, RI: American Mathematical Society, 2012.
- [54] J. Zahl. *An improved bound on the number of point-surface incidences in three dimensions*. Contrib. Discrete Math. 8 (2013): 100–121.
- [55] J. Zahl. *Breaking the 3/2 barrier for unit distances in three dimensions*. Int. Math. Res. Not. 2019.20 (2019): 6235–6284.
- [56] J. Zahl. *Sphere tangencies, line incidences and Lie’s line-sphere correspondence*. Math. Proc. Camb. Philos. Soc. 172(2) (2022):401-421.
- [57] F. de Zeeuw. *A short proof of Rudnev’s point-plane incidence bound*. preprint, arXiv:1612.02719.