

Point-variety incidences, unit distances and Zarankiewicz's problem for algebraic graphs

Aleksa Milojević*, Benny Sudakov*, István Tomon†

Abstract

In this paper we study the number of incidences between m points and n varieties in \mathbb{F}^d , where \mathbb{F} is an arbitrary field, assuming the incidence graph contains no copy of $K_{s,s}$. We also consider the analogous problem for algebraically defined graphs and unit distance graphs.

First, we prove that if \mathcal{P} is a set of m points and \mathcal{V} is a set of n varieties in \mathbb{F}^d , each of dimension d and degree at most Δ , and in addition the incidence graph is $K_{s,s}$ -free, then the number of incidences satisfies

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

This bound is tight when s, Δ are sufficiently large with respect to d , with an appropriate choice of $\mathbb{F} = \mathbb{F}(m, n)$. We give two proofs of this upper bound, one based on the framework of the induced Turán problems and the other based on VC-dimension theory. In the second proof, 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.

We also resolve the problem of finding the maximum number of unit distances which can be spanned by a set of n points \mathcal{P} in \mathbb{F}^d whose unit-distance graph is $K_{s,s}$ -free, showing that it is

$$\Theta_{d,s}(n^{2 - \frac{1}{\lceil d/2 \rceil + 1}}).$$

Finally, we obtain tight bounds on the maximum number of edges of a $K_{s,s}$ -free algebraic graph defined over a finite field, thus resolving the Zarankiewicz problem for this class of graphs.

1 Introduction

In a recent work [19], we established sharp bounds on the maximum number of incidences between m points and n hyperplanes in a d -dimensional space \mathbb{F}^d for an arbitrary field \mathbb{F} , under the standard non-degeneracy condition that the incidence graph contains no copy of the complete bipartite graph $K_{s,s}$. In this paper, we continue this line of research by establishing sharp bounds on the maximum number of edges in incidence graphs between points and varieties, algebraic graphs, and unit distance graphs, under the same non-degeneracy condition.

1.1 Point-Variety incidences

Given a set of points \mathcal{P} and a family of sets \mathcal{V} , the *incidence graph* of $(\mathcal{P}, \mathcal{V})$ is the bipartite graph $G(\mathcal{P}, \mathcal{V})$ with vertex classes \mathcal{P} and \mathcal{V} in which an edge is drawn between $x \in \mathcal{P}$ and $V \in \mathcal{V}$ if $x \in V$. The number of incidences in $(\mathcal{P}, \mathcal{V})$ is the number of edges of $G(\mathcal{P}, \mathcal{V})$, and it is denoted by $I(\mathcal{P}, \mathcal{V})$. First, we consider the case when the elements of \mathcal{V} are algebraic varieties, which are defined as common zero sets of a collection of polynomials, see Section 2 for precise definitions.

The number of incidences between points and varieties in real space have been extensively studied. Pach and Sharir [21, 22] 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

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

*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. Research supported in part by the Swedish Research Council grant VR 2023-03375

Zahl [28] and Basu and Sombra [2] extended this bound to dimension 3 and 4, respectively. Subsequently, Fox, Pach, Sheffer, Suk, and Zahl [13] 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

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

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 [24].

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). We obtain the following theorem.

Theorem 1.1. *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 (1). As $s \rightarrow \infty$, the bound in (1) approaches the shape $m^{\frac{D-1}{D}}n + m + n$, which matches the one in Theorem 1.1 in case $d = D - 1$. For smaller d , Theorem 1.1 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 1.2. *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*

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

The proof of Theorem 1.1 uses the framework of induced Turán problems, as introduced in [16, 19]. Namely, 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 [16] to prove bounds on the number of edges of induced $H_{d,\Delta}$ -free and $K_{s,s}$ -free bipartite graphs.

Furthermore, we provide an alternative proof of Theorem 1.1 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 [23] 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 we think that 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.

Finally, in order to prove Theorem 1.2 as well as other constructions in this paper, we use the results from [26], which are based on the random algebraic method originally introduced by [4].

1.2 Algebraic graphs

Fox, Pach, Sheffer, Suk, and Zahl [13] considered the classical Zarankiewicz's problem for semi-algebraic graphs. This type of problems comes from an old question of Zarankiewicz, who asked

what is the largest number of edges that an n -vertex graph can have without containing a copy of $K_{s,s}$. In this setting, 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 . The authors of [13] showed 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

$$|E(G)| = O_{D, t, s, \varepsilon} \left((mn)^{\frac{D}{D+1} + \varepsilon} + m + n \right). \quad (2)$$

This bound is only known to be tight for $D = 2$. Furthermore, it is worth highlighting that this upper bound matches (up to the ε error term) the best known upper bound on the number of incidences between m points and n hyperplanes in \mathbb{R}^D , originally proved by Apfelbaum and Sharir [1].

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 1.3. *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.

1.3 Unit distances in finite fields

The celebrated Erdős unit distance problem [8, 9] 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 still a large gap between the best known lower and upper bounds. 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 [9], and the upper bound is due to Spencer, Szemerédi and Trotter [25]. Furthermore, $\Omega(n^{4/3} \log \log n) < f_3(n) < n^{1.498}$ by [9] and [29], respectively. On the other hand, for $d \geq 4$, it is not difficult to construct a set of points achieving $f_d(n) = \Theta_d(n^2)$ [17]. 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 [13], 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 $\varepsilon > 0$. The key observation is that such unit distance graphs are semi-algebraic of complexity at most 2 containing no copy of $K_{s,s}$, in which case the previous bound follows simply from (2). See also Frankl and Kupavskii [11] for further strengthening.

Here, we completely resolve the analogous 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 .

Using Theorem 1.3 and noting that unit distance graphs are algebraic of complexity 2 in $(\mathbb{F}^d, \mathbb{F}^d)$, we immediately get that the maximum number of unit distances spanned by n points in \mathbb{F}^d such that the unit distance graph contains no copy of $K_{s,s}$ is $O_{d,s}(n^{2-1/d})$. However, as we show in the next theorem, this can be significantly improved.

Theorem 1.4. *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 1.5. *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.*

Organization. For the reader's convenience, we begin by introducing the definitions and basics of algebraic geometry in Section 2. Then, in Section 3, we present two proofs of Theorem 1.1, one based on extremal graph theory and another slightly weaker variant using shatter functions and VC dimension. We postpone the constructions showing tightness of these bounds to Section 4, where we also prove Theorem 1.3. Finally, in Section 5, we address the unit-distance problem and prove Theorems 1.4 and 1.5.

2 Preliminaries from algebraic geometry

Much of the following material is based on the Appendix of [6] and the excellent book of Cox, Little and O'Shea [5].

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 [5]).

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 $\overline{\mathbb{F}}[x_1, \dots, x_D]/I(Y)$. The same definition works for projective algebraic sets $Y \subseteq \mathbb{P}^D(\overline{\mathbb{F}})$, where $I(Y)$ stands for 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 [14]). 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. 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 $\overline{\mathbb{F}}$. Furthermore, 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 integer function $t \mapsto \dim(\overline{\mathbb{F}}_t[x_0, x_1, \dots, x_D]/I_t)$ 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 [5]). 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 [10]).

Theorem 2.1. *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 [20]).

Theorem 2.2. *Let V be a projective variety and let $I = I(V)$ be the associated ideal. Then, for every positive integer t , $\dim(\overline{\mathbb{F}}_t[x_0, x_1, \dots, x_D]/I_t) \leq \deg(V)t^{\dim V} + \dim V$.*

3 Point-variety incidences

3.1 Upper bounds via extremal graph theory

In this section, we prove Theorem 1.1 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 $k = 2^{\Delta} + 1$. 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 3.1. *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 2.1, 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 2.1, 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 2.1 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 [16], 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 [16].

Lemma 3.2 (Lemma 2.3 in [16]). *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 3.3 (Proposition 2.4 in [16]). *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 3.4. *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 [12].

Lemma 3.5. *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 3.4, 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 3.3 to find an (H, k, s) -rich set $S \subset X$ of size $|B|$. But then, by Lemma 3.2, there is an embedding of H in which B is embedded to S , finishing the proof. \square

Proof of Theorem 1.1. By Proposition 3.1, 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 3.5 with $k = d+1$ and $H = H_{d,\Delta}$ gives the desired bound. \square

3.2 Varieties containment patterns and dual shatter functions

In this section, we show a variant of Theorem 1.1 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 [27]).

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 3.6. *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 [23] 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 3.7. *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 3.7 combines the ideas from [23] with Hilbert's theorem. Hence, before presenting the proof, let us recall the result of Rónyai, Babai and Ganapathy [23]. 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

$$|\mathcal{Z}(f_1, \dots, f_k)| \leq \binom{k\Delta}{D}. \quad (3)$$

Note that (3) can be used as a black box to derive a weaker version of Lemma 3.7, 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 3.7.

Proof of Lemma 3.7. 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 the Preliminaries).

Let $x_1, \dots, x_N \in V_1$ be points witnessing the N distinct zero-patterns, and for $j \in [N]$, let $S_j \subset \{f_{2,1}, \dots, f_{k,t}\}$ be the set of polynomial 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 $S_j \not\subseteq S_\ell$, 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 $S_j \not\subseteq S_\ell$, 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 2.2, and so $M \leq N \leq O_{t,D}(k^d)$. \square

In the proof of Theorem 3.6, we also use the following result of [13], whose proof follows from the Haussler packing lemma [15].

Theorem 3.8 (Theorem 2.1 from [13]). *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 3.6. 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 3.7. 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 3.8 with $r = d + 1$, we get

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

\square

4 Algebraic Zarankiewicz's problem

In this section, we present the proof of our bound on Zarankiewicz's problem in algebraic graphs of bounded complexity. We will prove the upper and lower bounds separately, and we will then use the lower bound to show Theorem 1.2. In the proof of the upper bound, we combine Theorem 3.8 with the estimate (3).

For the lower bounds, we use the random polynomial method, which was pioneered by Bukh [4]. 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 [26]).

Theorem 4.1. *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 3.8. 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 3.8 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 (3) 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 4.2. *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

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

Claim 3.4 of [26] 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 4.3. *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 4.3 suffices to show the second part of Theorem 1.3.

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 [4]. 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 4.2 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 1.2, which demonstrates that our upper bounds on point-variety incidences are tight.

Proof of Theorem 1.2. 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 4.3.

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 [5], 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:

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

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

5 Unit distances

5.1 Geometry of spheres

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 Upper bounds

In this section, we prove Theorem 1.4. 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 $O_{d,s}(n^{2 - \frac{1}{\lceil d/2 \rceil + 1}})$.*

Following the ideas of [19], we prove that there exists a simple family \mathcal{F}_d of graphs with the property that an incidence graph of points and spheres in \mathbb{F}^{d-1} does not contain any member of \mathcal{F}_d as an induced subgraph. The family \mathcal{F}_d is defined to contain 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 $i = 1, \dots, d - 2$.

A simpler way to visualise this family is through the notion of a *pattern*. A *pattern* is an edge labeling of a complete bipartite graph, where every edge is labeled 0, 1 or *. A pattern Π corresponds to a family of graphs \mathcal{F}_Π on the same vertex set, which contains the graphs F such that 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. Finally, 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 .

One can express the condition that the incidence graph does not contain an induced member of the family \mathcal{F}_d by forbidding the following pattern. Namely, let Π_d be $2d$ 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.

We begin by presenting the basic geometric lemma which shows that incidence graphs of points and unit spheres in \mathbb{F}^d avoid the pattern Π_{d+1} .




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

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 $O_{d,s}\left(n^{2 - \frac{1}{\lceil d/2 \rceil + 1}}\right)$ edges.*

Proof. In [19], 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 with many unit distances

In this section, our goal is to prove Theorem 1.5. 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 $\Omega\left(n^{2 - \frac{1}{\lceil d/2 \rceil + 1}}\right)$ 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 [7].

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 [18]). 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 2.1 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 1.5. 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 [3], 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

Acknowledgements. We would like to thank Larry Guth for suggesting to study point-variety incidences and stimulating discussions.

References

- [1] 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.
- [2] S. Basu, M. Sombra. *Polynomial partitioning on varieties and point-hypersurface incidences in four dimensions*. Discrete Comput. Geom. 55 (2016): 158–184.
- [3] R. Breusch. *Zur Verallgemeinerung des Bertrandschen Postulates*. dass zwischen x und $2x$ stets Primzahlen liegen, Math. Z. 34 (1932): 505–526.
- [4] B. Bukh. *Random algebraic construction of extremal graphs*. Bulletin of London Math. Soc. 47(6) (2015): 939–945.
- [5] 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.
- [6] Z. Dvir, A. Gabizon, and A. Wigderson. *Extractors And Rank Extractors For Polynomial Sources*. Comput. Complex. 18 (2009): 1–58.
- [7] Z. Dvir, J. Kollár, and S. Lovett. *Variety Evasive Sets*. Comput. Complex. 23 (2014): 509–529.
- [8] P. Erdős. *On sets of distances of n points*. Amer. Math. Monthly 53 (1946): 248–250.
- [9] P. Erdős. *On sets of distances of n points in Euclidean space*. Magyar Tud. Akad. Mat. Kutató Int. Közl. 5 (1960): 165–169.
- [10] W. Fulton, *Intersection Theory*, 2nd ed. New York, NY: Springer, 1998.
- [11] N. Frankl, and A. Kupavskii. *On the Erdős-Purdy problem and the Zarankiewicz problem for semialgebraic graphs*. preprint, arXiv:2112.10245.
- [12] J. Fox, and B. Sudakov. *Dependent random choice*. Random Structures & Algorithms 38 (2011): 1–32.
- [13] 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.
- [14] R. Hartshorne. *Algebraic Geometry*. 1st ed. New York, NY: Springer, 1997.
- [15] D. Haussler. *Sphere packing numbers for subsets of the boolean n -cube with bounded Vapnik-Chervonenkis dimension*. 69 (2) (1995): 217–232.
- [16] Z. Hunter, A. Milojević, B. Sudakov, and I. Tomon. *Kővári-Sós-Turán theorem for hereditary families*. preprint, arXiv:2401.10853 (2024).
- [17] H. Lenz. *Zur Zerlegung von Punktmengen in solche kleineren Durchmessers*. Arch. Math. 5 (1955): 413–416.
- [18] R. Lidl, and H. Niederreiter. *Finite Fields*. Second ed. Cambridge: Cambridge University Press, 1997.
- [19] A. Milojević, B. Sudakov, and I. Tomon. *Point-hyperplane incidences via extremal graph theory*. preprint, arXiv:2401.06670 (2024).

- [20] J.V. Nesterenko, and F. Amoroso. *Introduction to Algebraic Independence Theory*. Berlin: Springer, 2001.
- [21] J. Pach, M. Sharir. *Repeated angles in the plane and related problems*. J. Combin. Theory Ser. A 59 (1992): 12–22.
- [22] J. Pach, M. Sharir. *On the number of incidences between points and curves*. Combin. Probab. Comput. 7 (1998): 121–127.
- [23] 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
- [24] A. Sheffer. *Lower bounds for incidences with hypersurfaces*. Discrete Anal. 1 (16) (2016): 14pp.
- [25] 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.
- [26] 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
- [27] T. Tao. *Higher Order Fourier Analysis*. Providence, RI: American Mathematical Society, 2012.
- [28] J. Zahl. *An improved bound on the number of point-surface incidences in three dimensions*. Contrib. Discrete Math. 8 (2013): 100–121.
- [29] J. Zahl. *Breaking the $3/2$ barrier for unit distances in three dimensions*. Int. Math. Res. Not. 2019.20 (2019): 6235–6284.