On Incidences Between Points and Hyperplanes*

Roel Apfelbaum† Micha Sharir‡

September 26, 2005

Abstract

We show that if the number I of incidences between m points and n planes in $\mathbb{R}^3$ is sufficiently large, then the incidence graph (that connects points to their incident planes) contains a large complete bipartite subgraph involving r points and s planes, so that $rs \geq \frac{I^2}{mn} - a(m+n)$, for some constant $a > 0$. This is shown to be almost tight in the worst case because there are examples of arbitrarily large sets of points and planes where the largest complete bipartite incidence subgraph records only $\frac{I^2}{mn} - \frac{m+n}{16}$ incidences. We also make some steps towards generalizing this result to higher dimensions.

On the way, we slightly improve upon a result of Brass and Knauer [BK] about the representation complexity of incidences between m points and n hyperplanes in $\mathbb{R}^d$, and get rid of the logarithmic factor in their upper bound.

1 Introduction

Let P be a set of m points, and let Π be a set of n hyperplanes in $\mathbb{R}^d$. We denote by $G(P, \Pi)$ their *incidence graph*, that is, the set of all point-hyperplane pairs $(p, \pi) \in P \times \Pi$, such that $p \in \pi$. We let $I(P, \Pi)$ denote the total number of these incidences. There have been several works on point-hyperplane incidences in the past 15 years [AA, BK, EGS, ET], which we shall review later on. The reader can also consult the recent survey by Pach and Sharir [PS], which reviews some of these results.

As we show, an interesting property of point-hyperplane incidence graphs is that, if the number of incidences is large (close to mn), then the incidence graph contains large complete bipartite subgraphs. Such a subgraph is in fact a configuration consisting of many hyperplanes of Π intersecting at a common lower dimensional affine subspace H, together with many points of P, all incident to H. This property arises, in one way or another, in almost all previous works; see below for details. In this paper we continue to study this property and ask: Given a point-hyperplane configuration

*Work on this paper has been supported by a grant from the Israeli Academy of Sciences for a Center of Excellence in Geometric Computing at Tel Aviv University, Work by Micha Sharir was also supported by by a grant from the U.S.–Israel Binational Science Foundation, by NSF Grant CCR-00-98246, and CCF-05-14079, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. This paper is part of the first author's M.Sc. and Ph.D. dissertations, prepared under the supervision of the second author.

†School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; `roel6@hotmail.com`.

‡School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; `michas@post.tau.ac.il`.

with many incidences, what is the size of the largest complete bipartite incidence subgraph? To state the question more precisely, we define

$$\mathrm{rs}(P,\Pi) = \max\{ rs \mid K_{r,s} \subseteq G(P,\Pi) \},$$

where $K_{r,s}$ denotes the complete bipartite subgraph with r vertices on one side and s vertices on the other, and the notation $K_{r,s} \subseteq G(P,\Pi)$ means that $K_{r,s}$ is a subgraph of $G(P,\Pi)$, such that there are some r points of P and s hyperplanes of Π all incident to one another. We let

$$\mathrm{rs}_d(m,n,I) = \min_{\substack{|P|=m \\ |\Pi|=n \\ I(P,\Pi)\geq I}} \mathrm{rs}(P,\Pi)$$

denote the minimum of $\mathrm{rs}(P,\Pi)$ over all choices of a set $P \subset \mathbb{R}^d$ of m points and a set Π of n hyperplanes in $\mathbb{R}^d$, such that $I(P,\Pi) \geq I$. Note that $\mathrm{rs}_d(m,n,I) \geq \max\left\{\frac{I}{m},\frac{I}{n}\right\}$, since there always exists a point incident to at least I/m hyperplanes, and a hyperplane incident to at least I/n points, which give rise to both subgraphs $K_{I/n,1}$ and $K_{1,I/m}$. We thus have $\mathrm{rs}_d(m,n,I) \geq \frac{I}{\min\{m,n\}} = \Omega\left(\frac{I}{m}+\frac{I}{n}\right)$, and any non-trivial estimate must exceed this lower bound.

1.1 Our results

For the case $d=3$, we can estimate $\mathrm{rs}_d(m,n,I)$ almost exactly:

Theorem 1.1 *(i) If $I = \Omega(m\sqrt{n}+n\sqrt{m})$, with a sufficiently large multiplicative constant, then*

$$\mathrm{rs}_3(m,n,I) = \frac{I^2}{mn} - \Theta(m+n).$$

(ii) If $m \leq n$, $I = O(n\sqrt{m})$, and $I = \Omega((mn)^{3/4})$, for appropriate multiplicative constants, then

$$\mathrm{rs}_3(m,n,I) = \Theta\left(\frac{I^4}{m^2n^3}+\frac{I}{m}\right).$$

(iii) Symmetrically, if $m \geq n$, $I = O(m\sqrt{n})$, and $I = \Omega((mn)^{3/4})$, then

$$\mathrm{rs}_3(m,n,I) = \Theta\left(\frac{I^4}{m^3n^2}+\frac{I}{n}\right).$$

(iv) If $I = O(m^{3/4}n^{3/4}+m+n)$, then

$$\mathrm{rs}_3(m,n,I) = \Theta\left(\frac{I}{m}+\frac{I}{n}\right).$$

The interesting case is (i), where the number of incidences is largest. The upper bound construction for this case consists of almost disjoint complete bipartite subconfigurations.

As the dimension d increases beyond 3, there are progressively more different ranges of I (as a function of m,n) where the lower or upper bounds for rs_d might change qualitatively. A complete analysis of all these cases seems hard at this point, already in four dimensions, and even the simpler question of focusing on the range where I is largest seems quite hard. There are actually two subproblems here: (a) Determine the range itself (i.e., how large must I be to ensure the best available bounds on rs_d), and (b) determine these best bounds for rs_d in this range.

We have obtained some non-trivial results for the higher dimensional case, but not tight ones:

Theorem 1.2 (Lower bound) *If* $I = \Omega(mn^{1-\frac{1}{d-1}} + m^{1-\frac{1}{d-1}}n)$*, then*

$$\mathrm{rs}_d(m,n,I) = \Omega\left(\left(\frac{I}{mn}\right)^{d-1} mn\right),$$

where the constant of proportionality depends on d*.*

Theorem 1.3 (Upper bound) *If* $I = \Omega((mn)^{1-\frac{1}{d-1}})$*, then*

$$\mathrm{rs}_d(m,n,I) = O\left(\left(\frac{I}{mn}\right)^{\frac{d+1}{2}} mn\right),$$

where the constant of proportionality depends on d*.*

Note that for $d = 3$, both theorems yield the same bound on rs_3, which is also identical (up to multiplicative constants) to that in Theorem 1.1(i). Moreover, all three theorems apply within the same (asymptotic) range $I = \Omega(m\sqrt{n} + n\sqrt{m})$ (Theorem 1.3 applies within a wider range).

It is interesting to compare these bounds to the equivalent bounds for general graphs. There are (m,n)-bipartite graphs with $\frac{1}{2}mn$ edges, such that the largest complete bipartite subgraph has fewer than $2(m+n)$ edges. For example, take a random subgraph H of the complete (m,n)-bipartite graph, in which each of the possible mn edges is chosen independently at random with probability $\frac{1}{2}$. The probability of a particular pair of subsets of r vertices from one side, and of s vertices from the other side, with $rs \geq 2(m+n)$, to induce a complete bipartite subgraph of H is at most $2^{-2(m+n)}$. The number of possible pairs of subsets as above is less than 2^{m+n}. Thus, the expected number of subsets, which induce a complete subgraph with at least $2(m+n)$ edges, is less than $2^{-(m+n)}$. Namely, for most choices of H there are no such subgraphs (see [Bo] for similar constructions). In contrast, Theorems 1.1 and 1.2 assert that a point-hyperplane incidence graph with these many edges has a complete bipartite subgraph with $\Omega(mn)$ edges.

On a different but closely related subject, we have also slightly improved upon a result of Brass and Knauer [BK], who have shown that any d-dimensional point-hyperplane incidence graph is the disjoint union of complete bipartite subgraphs, so that the overall size of their vertex sets is $O(((mn)^{1-\frac{1}{d+1}} + m + n)\log(m+n))$. We remove the logarithmic factor and show that the total size is actually

$$O((mn)^{1-\frac{1}{d+1}} + m + n).$$

This bound is not known to be tight, even in $d = 3$ dimensions. Brass and Knauer derive weaker upper bounds, and we refine their lowe bound for $d = 3$.

1.2 Previous work

The problem of bounding the number of incidences between points and curves or surfaces is one of the classical problems in combinatorial geometry, and has been studied extensively during the past 20 years; see the recent survey [PS] for a comprehensive review of the state of the art in this area. Most of the study has focused on incidences in the plane, but a considerable amount of work has also been devoted to higher-dimensional problems. The specific problem of analyzing and bounding the number $I(P,\Pi)$ of incidences between a set P of m points and a set Π of n hyperplanes in d dimensions has already been studied in [AA, BK, EGS, ET].

The main technical issue that arises in the study of point-hyperplane incidences in $d \geq 3$ dimensions is the possible presence of many points of P incident to many hyperplanes of Π. This happens when the intersection of many of the hyperplanes is a nonzero-dimensional affine subspace, and many of the points lie in that subspace. In this case the incidence graph $G(P, \Pi)$ can be a complete bipartite graph, or contain large such subgraphs, and then $I(P, \Pi)$ can be as high as (the trivial upper bound) mn.

Several attempts can be (and have been) made to study this problem in more restricted settings. One is to assume that, in $\mathbb{R}^3$, not too many points and/or not too many planes are collinear (or, for hyperplanes in higher dimensions, affinely dependent); see [EGS]. Another is to restrict the problem only to points that are vertices of the arrangement of the hyperplanes [AA, EGS]. Under these assumptions, better (nontrivial) upper bounds on $I(P, \Pi)$ can be obtained. For example:

- The maximum number of incidences between n hyperplanes in $\mathbb{R}^d$ and $m \geq n^{d-2}$ vertices of their arrangement is $\Theta(m^{2/3}n^{d/3} + n^{d-1})$ [AA].
- For m points and n planes in $\mathbb{R}^3$, if no three points are collinear, the number of incidences is $O(m^{3/5}n^{4/5} + m + n)$ (see [EGS][1]). The symmetric bound $O(m^{4/5}n^{3/5} + m + n)$ holds when no three planes are collinear. Brass and Knauer [BK] give a construction from which it follows that the latter bound is tight in the worst case, when no three planes are collinear[2].
- As a consequence of our analysis, for m points and n hyperplanes in $\mathbb{R}^d$, and for any fixed $r, s > 0$, if the incidence graph does not contain $K_{r,s}$ as a subgraph, the number of incidences is $O((mn)^{1-\frac{1}{d+1}} + m + n)$; see a remark at the end of Section 4.2.
- Elekes and Tóth [ET] have studied incidences between points and "non-degenerate" hyperplanes, where a hyperplane is considered degenerate if it contains a lower dimensional affine subspace that contains at least a constant fraction, say β, of its incident points. Elekes and Tóth show that the number of such hyperplanes that are incident to at least k points in a given set of m points in $\mathbb{R}^d$ is $O(m^d/k^{d+1} + (m/k)^{d-1})$, where the constant of proportionality depends on d and β. This, in turn, implies that the number of incidences between m points and n non-degenerate hyperplanes in $\mathbb{R}^d$ is $O((mn)^{1-\frac{1}{d+1}} + mn^{1-\frac{1}{d-1}})$.

Brass and Knauer [BK] considered the general case, where the incidence graph $G(P, \Pi)$ can contain large complete bipartite subgraphs. Rather than bounding $I(P, \Pi)$ itself, they have obtained an upper bound for the overall minimum possible complexity of a representation of $G(P, \Pi)$ as the disjoint union of complete bipartite graphs, that is, $G(P, \Pi) = \bigcup_{i=1}^{s} A_i \times B_i$, where $A_i \subseteq P$ and $B_i \subseteq \Pi$, for all $i = 1, \ldots, s$, and each incidence is recorded exactly once in this union. The *complexity* of such a representation of $G(P, \Pi)$ is defined to be $\sum_{i=1}^{s}(|A_i| + |B_i|)$, and the smallest complexity of such a representation, or the *representation complexity* of $G(P, \Pi)$, is denoted by $J(P, \Pi)$. We let $J_d(m, n)$ denote the maximum of $J(P, \Pi)$ over all sets P of m points, and Π of n hyperplanes in $\mathbb{R}^d$. Brass and Knauer [BK] have shown that

$$J_d(m, n) = O(((mn)^{1-\frac{1}{d+1}} + m + n)\log(m + n)). \quad (1)$$

[1] In the original paper, this bound is multiplied by a subpolynomial factor of the form $m^{\delta}n^{\delta}$, for any $\delta > 0$. This factor, however, can be eliminated using a more refined analysis.

[2] Brass and Knauer do not derive this specific bound, although it is implicit in their construction; see later in this section and in Appendix A. We remark that the symmetric case, where no three points are collinear, is *not* known to be tight in the worst case, because of some subtle aspects of point-hyperplane duality; see Appendix A.

As noted, we improve this result and show that

$$J_d(m,n) = O((mn)^{1-\frac{1}{d+1}} + m + n). \tag{2}$$

One way to interpret (1) and (2) is that if the number of incidences $I(P,\Pi)$ is much larger than $J_d(m,n)$, then $G(P,\Pi)$ should contain large complete bipartite subgraphs (or else the succint representation would not be possible). This has been one of our main motivations to study how large must these complete bipartite subgraphs be.

On the flip side of the same coin, one would like to obtain constructions of sets P of m points, and Π of n hyperplanes, so that $G(P,\Pi)$ contains no large complete bipartite subgraphs and $I(P,\Pi)$ is as large as possible. Here too one would hope to obtain such constructions with $I(P,\Pi)$ close to $J_d(m,n)$, or, in three dimensions, to $\Theta(m^{3/4}n^{3/4}+m+n)$. The best three-dimensional construction to date is due to Brass and Knauer [BK], where $G(P,\Pi)$ does not contain any $K_{2,3}$, and $I(P,\Pi) = \Omega(m^{7/10}n^{7/10})$ in the balanced case $m \approx n$. We note that their construction actually yields the bound $I(P,\Pi) = \Omega(m^{4/5}n^{3/5})$, and has the property that no three planes are collinear. Thus, the known upper bound $I(P,\Pi) = O(m^{4/5}n^{3/5}+m+n)$ for this restricted case (which, as noted above, follows from the analysis in [EGS]) is worst-case tight. For the sake of completeness, we present the construction in Appendix A.

2 Many Incidences Yield Large Complete Bipartite Incidence Subgraphs (in $\mathbb{R}^3$)

In this section we prove of Theorem 1.1. The main result here is the following lower bound.

Theorem 2.1 (Cf. Theorem 1.1(i)—lower bound) *Let P be a set of m points and Π a set of n planes in $\mathbb{R}^3$, with I incidences between them. Then there exists a line ℓ containing r points of P and contained in s planes of Π, such that*

$$\sqrt{rs} \geq \frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I},$$

where $a > 0$ is some sufficiently large constant.

This inequality, when squared, implies that $rs \geq \frac{I^2}{mn} - 2a(m+n)$. This establishes the lower bound of Theorem 1.1(i). Note that here there is no lower bound requirement on I, as opposed to Theorem 1.1(i), where it is required that $I = \Omega(m\sqrt{n}+n\sqrt{m})$. However, if $I < \sqrt{amn(m+n)}$, then the right hand side is negative. Thus the theorem is interesting only for point-plane configurations with $I > \sqrt{amn(m+n)} = \Omega(m\sqrt{n}+n\sqrt{m})$.

On the upper bound side, the situation is much simpler, so we first dispose of this case.

Lemma 2.2 (Theorem 1.1(i)—upper bound) *There exist arbitrarily large configurations of m points and n planes in $\mathbb{R}^3$ with $I = \Omega(m\sqrt{n}+n\sqrt{m})$ incidences, such that the largest $K_{r,s}$ incidence subgraph satisfies*

$$rs \leq \frac{I^2}{mn} - \frac{1}{16}(m+n).$$

Proof: Without loss of generality, we present the construction for $m \geq n$. Fix two arbitrarily large numbers, $r \geq s \geq 2$, and a third number $2 \leq k \leq s$. Take a set L of k parallel lines such that no three lines are coplanar. Then each pair of lines in L determine a distinct plane. We include all these $\binom{k}{2}$ planes in the set Π of n planes. We include in Π additional planes, each of which contains just one of the lines of L, so that each line is incident to exactly s planes. The set P of points consists of $m = rk$ elements, so that each line of L contains r points. The set Π consists of $sk - \binom{k}{2}$ planes, and $I(P, \Pi) = krs$. Put $n_0 = sk = n + \binom{k}{2}$. Note that $\mathrm{rs}(P, \Pi) = rs$, because the corresponding subgraph $K_{r,s}$ cannot contain points from three lines—no plane passes through three lines of L, and if it contains points from two lines then there is only one plane that passes through both lines. This gives

$$\mathrm{rs}(P, \Pi) = rs = \frac{I^2}{mn_0} = \frac{I^2}{m} \cdot \frac{1}{n + \binom{k}{2}} \leq \frac{I^2}{m} \cdot \frac{1}{n + k^2/4}.$$

We now use the inequality

$$\frac{1}{x+h} \leq \frac{1}{x} - \frac{h}{2x^2},$$

which holds for all $x \geq h > 0$, with $x = n$ and $h = k^2/4 \leq n$, to get

$$rs \leq \frac{I^2}{m} \cdot \left(\frac{1}{n} - \frac{k^2/4}{2n^2} \right) = \frac{I^2}{mn} - \frac{(kI)^2}{8mn^2},$$

and since $k = \frac{mn_0}{I}$, we get

$$rs \leq \frac{I^2}{mn} - \frac{(mn_0)^2}{8mn^2} \leq \frac{I^2}{mn} - \frac{m}{8} \leq \frac{I^2}{mn} - \frac{1}{16}(m+n),$$

as claimed. Note that $I = \frac{mn_0}{k} \geq m\left(\frac{n}{k} + \frac{k}{4}\right) \geq m\sqrt{n} = \Omega(m\sqrt{n} + n\sqrt{m})$ is in the required range.

The case $m \leq n$ is handled in a symmetric manner, using duality between points and planes. This completes the proof. □

Remark: A simpler construction, consisting of *disjoint* copies of $K_{r,s}$, yields the lower bound $I^2/(mn)$. Our construction shows that a lower order term proportional to $m + n$ is unavoidable.

Lemma 2.3 (Theorem 1.1(iv)—upper bound) *For any $m, n, I > 0$, such that $\Omega(m+n) = I = O(m^{3/4}n^{3/4})$, there exist configurations of m points and n planes in $\mathbb{R}^3$ with at least I incidences, such that the largest $K_{r,s}$ incidence subgraph satisfies*

$$rs \leq \frac{6I}{\min\{m,n\}}.$$

Remark: The construction of [BK] (see Appendix A), in which there are no $K_{3,2}$ or $K_{2,3}$ incidence subgraphs, provides us with an example where $rs = O(I/\min\{m,n\})$, showing the bound is asymptotically tight. This construction, however, is good only for the range $I = O(m^{3/5}n^{4/5} + m^{4/5}n^{3/5})$, and cannot be used for larger values of I within the assumed larger range $O((mn)^{3/4})$. In contrast, our construction is good for the entire range specified in Lemma 2.3, but may have complete bipartite incidence subgraphs with an arbitrarily large number of elements on both sides.

Proof: This construction resembles similar constructions of Elekes [El]. Put $k = \left\lfloor \sqrt{2I/n} \right\rfloor$, $l = \left\lfloor \sqrt{6I/m} \right\rfloor$, and $t = \lfloor (mn)^{3/2}/(12\sqrt{3}I^2) \rfloor$. With an appropriate choice of the constants of proportionality, we may assume that $k, l, t \geq 3$. Define

$$P = \{ (x_1, x_2, x_3) \mid x_1, x_2 \in \{1, \ldots, k\}, \text{ and } x_3 \in \{1, \ldots, 3klt\} \},$$

and

$$\Pi = \{ x_3 = a_1 t x_1 + a_2 t x_2 + b \mid a_1, a_2 \in \{1, \ldots, l\}, \text{ and } b \in \{1, \ldots, klt\} \}.$$

The set P consists of $3k^3lt \leq m$ points, and the set Π consists of $kl^3t \leq n$ planes. Each plane is incident to k^2 points, and each point is incident to at most l^2 planes, so the number of incidences is $k^3l^3t \geq I$.

Now there are three types of complete bipartite incidence subgraphs $K_{r,s}$:

1. Between a point and all its incident planes. Then $r = 1$, $s \leq l^2$, and $rs \leq l^2 \leq 6I/m$.
2. Between a plane and all its incident points. Then $r = k^2$, $s = 1$, and $rs = k^2 \leq 2I/n$.
3. Between some r collinear points and s collinear planes, all incident to the same line. Then $r \leq k$, $s \leq l$, and $rs \leq kl = 2\sqrt{3}I/\sqrt{mn} \leq 2\sqrt{3}I/\min\{m, n\}$.

In either case, we have $rs \leq \frac{6I}{\min\{m,n\}}$. This completes the proof. □

To prove Theorem 2.1, we use the result of Elekes and Tóth [ET]. We define a hyperplane π to be *β-degenerate* with respect to a point set P, if some lower dimensional flat $F \subset \pi$ contains at least a β-fraction of the points of P incident to π, i.e., if

$$|P \cap F| \geq \beta\, |P \cap \pi|,$$

for some lower dimensional flat $F \subset \pi$. If no such flat exists, then the hyperplane is said to be *β-non-degenerate*.[3] A hyperplane π is called *k-rich* (with respect to P) if it contains at least k points of P. Elekes and Tóth show:

Theorem 2.4 ([ET]) *For any integer $d \geq 3$, there are constants $\beta_d > 0$ and $C_1(d) > 0$, such that for any set of m points in $\mathbb{R}^d$, the number of k-rich β_d-non-degenerate hyperplanes is at most*

$$C_1(d) \left(\frac{m^d}{k^{d+1}} + \frac{m^{d-1}}{k^{d-1}} \right),$$

and this bound is best possible.

This theorem can be rephrased in terms of an upper bound on the number of incidences between m points and n β_d-non-degenerate hyperplanes. For the sake of completeness, we include a proof of this fact.

[3] We caution the reader that this notation is the *opposite* to that used in [ET].

Corollary 2.5 *For any integer $d \geq 3$, there is a constant $C_2(d)$, such that the number of incidences between any set of m points and a set of n β_d-non-degenerate hyperplanes (with respect to the given point set) in $\mathbb{R}^d$ is at most*

$$C_2(d)\left((mn)^{1-\frac{1}{d+1}} + mn^{1-\frac{1}{d-1}}\right).$$

Proof: Let P be a set of m points in $\mathbb{R}^d$, let Π be a set of n β_d-non-degenerate hyperplanes with respect to P, and let $I = I(P, \Pi)$ denote the number of their incidences.

Delete from Π all hyperplanes that are incident to fewer than $I/(2n)$ points of P. This removes at most $I/2$ of the incidences, se we are left with a set Π' of $n' \leq n$ hyperplanes, such that $I(P, \Pi') \geq I/2$, and such that each $\pi \in \Pi'$ is $(I/(2n))$-rich (and remains β_d-non-degenerate).

If $n' \leq n/4$ we use induction (the induction basis $n = O(1)$ is trivial to establish), and have

$$\frac{I}{2} \leq I(P, \Pi') \leq C_2(d)\left(\left(\frac{mn}{4}\right)^{1-\frac{1}{d+1}} + m\left(\frac{n}{4}\right)^{1-\frac{1}{d-1}}\right),$$

or

$$\begin{aligned} I &\leq 2C_2(d)\left(\frac{(mn)^{1-\frac{1}{d+1}}}{4^{1-\frac{1}{d+1}}} + \frac{mn^{1-\frac{1}{d-1}}}{4^{1-\frac{1}{d-1}}}\right) \\ &\leq C_2(d)\left((mn)^{1-\frac{1}{d+1}} + mn^{1-\frac{1}{d-1}}\right). \end{aligned}$$

Otherwise, Theorem 2.4 implies that

$$\frac{n}{4} \leq n' \leq C_1(d)\left(\frac{m^d}{(I/(2n))^{d+1}} + \frac{m^{d-1}}{(I/(2n))^{d-1}}\right).$$

If the first term in the right hand side dominates, then $n\left(\frac{I}{2n}\right)^{d+1} \leq 8C_1(d)m^d$, or

$$I \leq 2(8C_1(d))^{\frac{1}{d+1}}(mn)^{1-\frac{1}{d+1}}.$$

If the second term dominates, then $n\left(\frac{I}{2n}\right)^{d-1} \leq 8C_1(d)m^{d-1}$, or

$$I \leq 2(8C_1(d))^{\frac{1}{d-1}}mn^{1-\frac{1}{d-1}}.$$

Hence, choosing $C_2(d) > 2(8C_1(d))^{\frac{1}{d-1}}$ yields the asserted bound. □

Using this bound, we can prove the following result, which is slightly weaker than Theorem 2.1; see below for a more detailed comparison.

Lemma 2.6 *Let P be a set of m points and Π a set of n planes in $\mathbb{R}^3$, such that $I = I(P, \Pi) = \Omega((mn)^{3/4} + m\sqrt{n})$, with a sufficiently large multiplicative constant. Then there exists a line ℓ containing r points of P and contained in s planes of Π, such that*

$$rs = \Omega\left(\min\left\{\frac{I^4}{m^2n^3}, \frac{I^2}{mn}\right\}\right).$$

This bound is asymptotically tight in the worst case.

Proof: Applying Corollary 2.5 with $d = 3$, we see that, when the constant of proportionality is chosen sufficiently large, most incidences are with planes of Π that are β_3-degenerate, i.e., for each such plane, at least a β_3-fraction of its incident points are contained in a single line.

Put $\beta = \beta_3$, and let $\Pi' \subseteq \Pi$ be the subset of those planes in Π that contain at least $I/(2n)$ points each, and are β-degenerate. By the preceding argument, if the constant of proportionality in the assumed lower bound on I is sufficiently large, then $I(P, \Pi') \geq I/3$, say. We replace each plane of Π' with a line that lies on it and contains a β-fraction of its incident points. Thus, each such line contains at least $\beta I/(2n)$ points of P. By projecting these lines and the points of P onto some generic plane, and applying the Szemerédi-Trotter bound for planar incidences [ST], the number of incidences between the points of P and these lines is

$$I' = O\left(\frac{m^2}{(\beta I/(2n))^2} + m\right) = O\left(\frac{m^2n^2}{I^2} + m\right).$$

Note that $I(P, \Pi')$ differs from I', because we count in $I(P, \Pi')$ each line with its *multiplicity*, equal to the number of planes of Π' that contain it. The average multiplicity of a line is thus

$$s = \frac{I(P, \Pi')}{I'} = \Omega\left(\frac{I}{I'}\right) = \Omega\left(\min\left\{\frac{I^3}{m^2n^2}, \frac{I}{m}\right\}\right).$$

By the pigeonhole principle, some line ℓ does have at least this multiplicity, i.e., it is contained in at least s planes. By construction, it also contains $r = \Omega(I/n)$ points. Altogether, we get

$$rs = \Omega\left(\min\left\{\frac{I^4}{m^2n^3}, \frac{I^2}{mn}\right\}\right).$$

We have thus found a line ℓ with the asserted properties.

This bound was proved using the bounds of Szemerédi and Trotter [ST] and of Elekes and Tóth [ET], which are both tight and have matching lower bound constructions. We can combine these constructions to obtain a point-plane configuration that has a matching upper bound on $\mathrm{rs}(P, \Pi)$ of the same order of magnitude as the lower bound we have just proved, i.e., $\mathrm{rs}(P, \Pi) = O(\min\{I^4/(m^2n^3), I^2/(mn)\})$ (that is, unless the trivial bound $\mathrm{rs}(P, \Pi) \geq \max\{I/m, I/n\}$ dominates). This is done as follows. We take m points spanning the maximal number of lines incident to $r = \Theta(I/n)$ of the points, which, by the lower bound of [ST], is

$$\Theta\left(\frac{m^2}{r^3} + \frac{m}{r}\right) = \Theta\left(\frac{m^2n^3}{I^3} + \frac{mn}{I}\right).$$

We then let each such line occur on $s = \Theta(\min\{I^3/(m^2n^2), I/m\})$ planes. The constants of proportionality are chosen so that the total number of planes is n, and the number of incidences is I. We have thus shown that the bound asserted in the lemma is asymptotically tight in the worst case. □

Proof of Theorem 1.1(ii,iii): If $I = O(n\sqrt{m})$, for an appropriate multiplicative constant, then, in the bound of Lemma 2.6, the first term $I^4/(m^2n^3)$ is smaller than the second term $I^2/(mn)$. Moreover, the lower bound on I that the lemma requires, holds under the assumptions $I = \Omega((mn)^{3/4})$ and $m \leq n$. Hence, under the assumptions of part (ii) of the theorem, $\mathrm{rs}_3(m, n) = \Omega(I^4/(m^2n^3))$, which clearly implies the lower bound of Theorem 1.1(ii). The upper bound also follows easily from Lemma 2.6. Finally, Theorem 1.1(iii) follows by point-plane duality.

□

On the other hand, if $I = \Omega(n\sqrt{m})$, then the second term in Lemma 2.6, namely $I^2/(mn)$, dominates. We thus get:

Corollary 2.7 *Let P be a set of m points and Π a set of n planes in $\mathbb{R}^3$, such that $I = I(P,\Pi) = \Omega(m\sqrt{n} + n\sqrt{m})$, with a sufficiently large multiplicative constant. Then there exists a line ℓ containing r points of P and contained in s planes of Π, such that*

$$rs = \Omega\left(\frac{I^2}{mn}\right).$$

(The lemma is applicable since $(mn)^{3/4}$ is always dominated by $m\sqrt{n} + n\sqrt{m}$.)

This is already very close to the bound we are trying to prove, except for the multiplicative constant. We shall now get rid of this constant and finish the proof of Theorem 2.1. Recall that the theorem states that

$$\sqrt{rs} \geq \frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I},$$

for some r points and s planes all incident to one another, and for some constant $a > 0$.

Proof of Theorem 2.1: Let P be a set of m points and Π a set of n planes in $\mathbb{R}^3$, with $I = I(P,\Pi)$ incidences. By Corollary 2.7, there exist positive absolute constants A, k, and β, such that for all $m > k$ and $n > k$, if $I > A(m\sqrt{n} + n\sqrt{m})$, then

$$\sqrt{\mathrm{rs}(P,\Pi)} \geq \frac{\beta I}{\sqrt{mn}}.$$

We choose the constant a so that it satisfies $a \geq \max\left\{4, 2A^2, k, \frac{2}{\beta}\right\}$.

The proof proceeds by induction on m and n. It is easy to see that the theorem holds for sufficiently small values of m, n, or I. More precisely, if $I < \sqrt{amn(m+n)}$, then $\frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I} < 0$, and the theorem holds trivially. Moreover, if $m \leq a$ or $n \leq a$, then $I \leq mn < \sqrt{amn(m+n)}$. Hence, the theorem holds for all m and n such that $m \leq a$ or $n \leq a$.

Suppose then that $m > a$ and $n > a$ are arbitrary, and that the claim holds for all (m', n') satisfying $m' < m$ and $n' < n$. Let P be a set of m points and Π a set of n planes in $\mathbb{R}^3$ with $I > \sqrt{amn(m+n)}$ incidences between them. Let ℓ be a line that maximizes rs, where $r = |\ell \cap P|$, and $s = |\{\pi \in \Pi \mid \pi \supset \ell\}|$.

Remove from the setting all the points and planes incident to ℓ. We are left with $m - r$ points and $n - s$ planes, and denote by I_1 the number of incidences among them. We note that

$$I_1 \geq I - rs - (m+n) + (r+s). \tag{3}$$

Indeed, by removing the elements incident to ℓ, we lose the rs incidences between these elements. We may also lose incidences between the removed points and the surviving planes, and between the removed planes and the surviving points. However, each surviving plane can be incident to at most one removed point, and each surviving point can be incident to at most one removed plane. This implies the asserted inequality (3).

We next choose a line ℓ_1 incident to r_1 of the remaining points, and to s_1 of the remaining planes, such that r_1s_1 is maximized. If $\sqrt{r_1s_1} \geq \frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I}$, then we are done, since, by construction, $rs \geq r_1s_1$. Otherwise, we may write

$$\frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I} > \sqrt{r_1s_1} \geq \frac{I_1}{\sqrt{(m-r)(n-s)}} - \frac{a((m+n)-(r+s))\sqrt{(m-r)(n-s)}}{I_1},$$

where the right inequality follows from the induction hypothesis, or

$$\frac{I}{\sqrt{mn}} > \frac{I_1}{\sqrt{(m-r)(n-s)}} + a\left(\frac{(m+n)\sqrt{mn}}{I} - \frac{((m+n)-(r+s))\sqrt{(m-r)(n-s)}}{I_1}\right).$$

Put

$$h = \frac{(m+n)\sqrt{mn}}{I} - \frac{((m+n)-(r+s))\sqrt{(m-r)(n-s)}}{I_1},$$

so we have

$$\frac{I}{\sqrt{mn}} > \frac{I_1}{\sqrt{(m-r)(n-s)}} + ah.$$

We now distinguish between the two cases $h \geq 0$ and $h < 0$. If $h \geq 0$, then we have

$$\frac{I}{\sqrt{mn}} > \frac{I_1}{\sqrt{(m-r)(n-s)}}, \tag{4}$$

or, using the inequality $(m-r)(n-s) \leq (\sqrt{mn} - \sqrt{rs})^2$, and applying (3),

$$\frac{I}{\sqrt{mn}} > \frac{I - rs - (m+n)}{\sqrt{mn} - \sqrt{rs}},$$

or

$$I\sqrt{mn} - I\sqrt{rs} > I\sqrt{mn} - \sqrt{mn}rs - (m+n)\sqrt{mn},$$

or

$$rs - \frac{I}{\sqrt{mn}}\sqrt{rs} + (m+n) > 0.$$

This quadratic inequality in the variable $\sqrt{rs}$ solves to

$$\begin{aligned}\sqrt{rs} &> \frac{\frac{I}{\sqrt{mn}} + \sqrt{\frac{I^2}{mn} - 4(m+n)}}{2}, \quad \text{or} \\ \sqrt{rs} &< \frac{\frac{I}{\sqrt{mn}} - \sqrt{\frac{I^2}{mn} - 4(m+n)}}{2}.\end{aligned}$$

Note that, since $a \geq 4$, it follows that $\frac{I^2}{mn} - 4(m+n) \geq 0$ for the assumed range of I. We can then use the inequality $\sqrt{x - \Delta x} \geq \sqrt{x} - \frac{\Delta x}{\sqrt{x}}$, which holds for $0 \leq \Delta x \leq x$, to obtain

$$\begin{aligned}\sqrt{rs} &> \frac{I}{\sqrt{mn}} - \frac{2(m+n)\sqrt{mn}}{I}, \quad \text{or} \\ \sqrt{rs} &< \frac{2(m+n)\sqrt{mn}}{I}.\end{aligned}$$

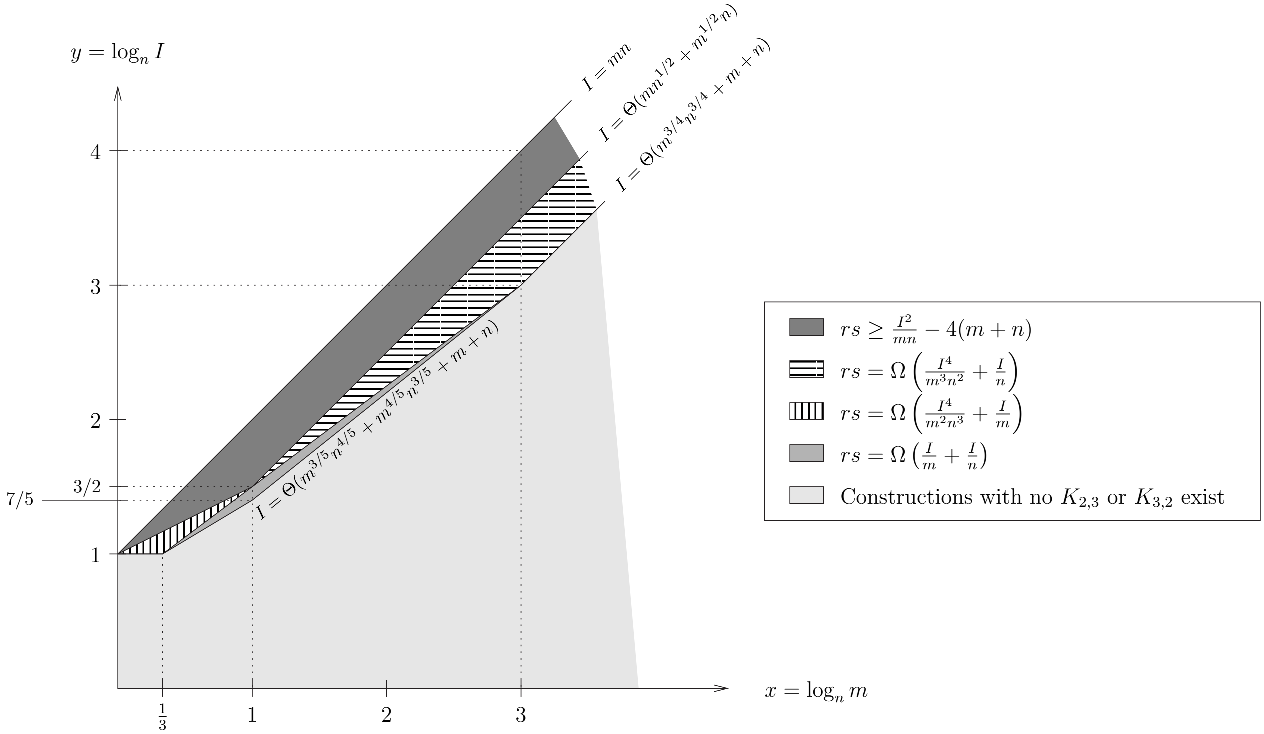

Figure 1: The known lower bounds for the maximum number of edges in a complete bipartite incidence subgraph in $\mathbb{R}^3$.

Since $a \geq 2A^2$, it is easily checked that Corollary 2.7 is applicable for the assumed range of I, and implies that $\sqrt{rs} \geq \frac{\beta I}{\sqrt{mn}}$. Hence, if the second case were possible, we would have $\frac{\beta I}{\sqrt{mn}} < \frac{2(m+n)\sqrt{mn}}{I}$, or $I < \sqrt{\frac{2}{\beta}mn(m+n)}$, which, having chosen $a \geq \frac{2}{\beta}$, would contradict our assumption on I. Hence, only the first inequality is possible, and the theorem holds in this case (since $a > 2$).

Consider now the case $h < 0$. We have

$$\begin{aligned} \frac{(m+n)\sqrt{mn}}{I} &< \frac{((m+n)-(r+s))\sqrt{(m-r)(n-s)}}{I_1} \\ &< \frac{(m+n)\sqrt{(m-r)(n-s)}}{I_1}, \end{aligned}$$

or

$$\frac{I}{\sqrt{mn}} > \frac{I_1}{\sqrt{(m-r)(n-s)}}.$$

But this is exactly inequality (4), which, as we have already seen, implies $\sqrt{rs} > \frac{I}{\sqrt{mn}} - \frac{a(m+n)\sqrt{mn}}{I}$, so the theorem holds in this case too.

This completes the induction step, and thus the proof of the theorem. □

Figure 1 summarizes our findings. Each differently-shaded region represents certain values of m, n and I, and has a different lower bound for rs.

3 Large Complete Bipartite Incidence Subgraphs in Higher Dimensions

In Lemma 2.6 we require that $I = \Omega((mn)^{3/4} + m\sqrt{n})$, because we want to ensure that most planes are 'degenerate' in the sense that they can be replaced by lines, and the number of incidences will stay roughly the same. However, the lemma holds in a considerably more general setting, involving any family of 'degenerate' subsets of points in any dimension. Specifically, we call a finite set of points $S \subset \mathbb{R}^d$ (β, j)-*degenerate*, if some j-flat contains at least a β-fraction of the points of S. In other words, if

$$|F \cap S| \geq \beta |S|$$

for some j-flat F of $\mathbb{R}^d$. If no such j-flat exists, we call S (β, j)-*non-degenerate*.[4] With this notion of degeneracy, Lemma 2.6 becomes a special case of the following lemma (with each plane $\pi \in \Pi$ being mapped to the set $\pi \cap P$, and the entire set of planes Π being mapped to a multiset of subsets of P).

Lemma 3.1 *Let $P \subset \mathbb{R}^d$ be a set of m points, let $\mathcal{T} \subseteq 2^P$ be a multiset of n subsets of P, and let $0 < \beta < 1$ be some constant, such that all the members of $\mathcal{T}$ are $(\beta, 1)$-degenerate. Then there exist a subset $R \subseteq P$ of $|R| = r$ points and a subfamily $\mathcal{S} \subset \mathcal{T}$ of $|\mathcal{S}| = s$ subsets (counted with multiplicity), such that $R \subseteq S$ for each $S \in \mathcal{S}$, and*

$$rs = \Omega\left(\min\left\{\frac{I^4}{m^2 n^3}, \frac{I^2}{mn}\right\}\right),$$

where $I = \sum_{T \in \mathcal{T}} |T|$.

In particular, the multiset $\mathcal{T}$ need not be induced by planes, as in Lemma 2.6, but can be induced by hyperplanes of any dimension. The proof is essentially identical to that of Lemma 2.6: We replace each subset $S \in \mathcal{S}$ by a line that contains a fraction of its points, and estimate the average multiplicity of the lines using the Szemerédi-Trotter bound within a generic 2-plane onto which we project the points and lines.

We next obtain the following generalization of Lemma 3.1.

Lemma 3.2 *Let $P \subset \mathbb{R}^d$ be a set of m points, let $\mathcal{T} \subseteq 2^P$ be a multiset of n subsets of P, and let $\beta > 0$ and $j \geq 1$ be some constants, such that all the members of $\mathcal{T}$ are (β, j)-degenerate. Then there exist a subset $R \subseteq P$ of $|R| = r$ points and a subfamily $\mathcal{S} \subset \mathcal{T}$ of $|\mathcal{S}| = s$ subsets (again, counted with multiplicity), such that $R \subseteq S$ for each $S \in \mathcal{S}$, and*

$$rs = \Omega\left(\min\left\{\frac{I^{j+3}}{m^{j+1} n^{j+2}}, \frac{I^{j+1}}{m^j n^j}\right\}\right),$$

where $I = \sum_{T \in \mathcal{T}} |T|$, and the constant of proportionality depends on β and j.

Proof: The proof proceeds by double induction on j and n. The base case $j = 1$ is given by Lemma 3.1 (for any n). Suppose now that the lemma holds for $j - 1 \geq 1$, and also for j and for $n' < n$, and we shall see that it also holds for j and for n. (The base case for n, at any fixed j, is trivial, with an appropriate choice of the constants of proportionality.)

[4] Again, this notation is opposite to that of [ET].

Delete from $\mathcal{T}$ all the members containing fewer than $I/(2n)$ points, and let $\mathcal{T}'$ denote the multiset of the remaining sets. We have $I' = \sum_{T \in \mathcal{T}'} |T| \geq I/2$. If $|\mathcal{T}'| < n/4$, then, by induction on n, we have subsets $R \subseteq P$ and $\mathcal{S} \subseteq \mathcal{T}'$, such that $R \subseteq S$ for each $S \in \mathcal{S}$, and

$$\begin{aligned} |R| \cdot |\mathcal{S}| &= \Omega\left(\min\left\{\frac{(I/2)^{j+3}}{m^{j+1}(n/4)^{j+2}}, \frac{(I/2)^{j+1}}{m^j(n/4)^j}\right\}\right) \\ &= \Omega\left(\min\left\{2^{j+1}\frac{I^{j+3}}{m^{j+1}n^{j+2}}, 2^{j-1}\frac{I^{j+1}}{m^j n^j}\right\}\right). \end{aligned}$$

Since $j \geq 2$, we obtain R and $\mathcal{S}$ that satisfy the asserted lower bound. We can therefore assume that there are at least $n/4$ remaining sets in $\mathcal{T}'$.

For each set $T \in \mathcal{T}'$, let π_T be a j-flat (which exists by assumption) containing at least $\beta|T| \geq \frac{\beta I}{2n}$ points of P. Project these j-flats and the points of P onto some generic $(j+1)$-space Q, and partition $\mathcal{T}'$ into two subfamilies:

$$\mathcal{T}_1 = \left\{ T \in \mathcal{T}' \;\middle|\; \pi_T \text{ is } \beta_{j+1}\text{-non-degenerate in } Q \right\}, \quad \text{and} \quad \mathcal{T}_2 = \mathcal{T}' \setminus \mathcal{T}_1.$$

Note that all the members of $\mathcal{T}_2$ are $(\beta\beta_{j+1}, j-1)$-degenerate, that is, informally, they are 'more degenerate' than the other members of $\mathcal{T}'$. One of these two families contains at least half of the members of $\mathcal{T}'$. If $|\mathcal{T}_2| \geq |\mathcal{T}'|/2 \geq n/8$, we have, by induction on j,

$$\begin{aligned} rs &= \Omega\left(\min\left\{\frac{I^{j+2}}{m^j n^{j+1}}, \frac{I^j}{m^{j-1}n^{j-1}}\right\}\right) \\ &= \Omega\left(\min\left\{\frac{I^{j+3}}{m^{j+1}n^{j+2}}, \frac{I^{j+1}}{m^j n^j}\right\}\right), \end{aligned}$$

with an appropriate careful choice of the constants of proportionality, and the lemma holds in this case.

Suppose then that $|\mathcal{T}_1| \geq |\mathcal{T}'|/2 \geq n/8$. Put $\Pi = \{\pi_T \mid T \in \mathcal{T}_1\}$. Since the j-flats $\pi \in \Pi$ are $\frac{\beta I}{2n}$-rich and β_{j+1}-non-degenerate with respect to P (in the space Q of projection), Theorem 2.4 implies that the number of these j-flats is upper-bounded by

$$\begin{aligned} |\Pi| &= O\left(\frac{m^{j+1}}{(\beta I/(2n))^{j+2}} + \frac{m^j}{(\beta I/(2n))^j}\right) \\ &= O\left(\frac{m^{j+1}n^{j+2}}{I^{j+2}} + \frac{m^j n^j}{I^j}\right). \end{aligned}$$

Taking into account that $|\mathcal{T}_1| \geq n/8$, the average multiplicity of an element of Π is

$$\frac{|\mathcal{T}_1|}{|\Pi|} = \Omega\left(\min\left\{\frac{I^{j+2}}{m^{j+1}n^{j+1}}, \frac{I^j}{m^j n^{j-1}}\right\}\right).$$

Let $\pi \in \Pi$ be a j-flat with at least this multiplicity. Define $R = \pi \cap P$, and $\mathcal{S} = \{T \in \mathcal{T}_1 \mid \pi_T = \pi\}$. We have (i) $r = |R| \geq \frac{\beta I}{2n} = \Omega(I/n)$, (ii) $s = |\mathcal{S}| \geq |\mathcal{T}_1|/|\Pi|$, (iii) R is contained in every member of $\mathcal{S}$, and

$$rs = \Omega\left(\frac{I}{n} \cdot \frac{|\mathcal{T}_1|}{|\Pi|}\right) = \Omega\left(\min\left\{\frac{I^{j+3}}{m^{j+1}n^{j+2}}, \frac{I^{j+1}}{m^j n^j}\right\}\right),$$

as asserted by the lemma. □

As a corollary, we obtain:

Theorem 3.3 *If $I = \Omega((mn)^{1-\frac{1}{d+1}} + mn^{1-\frac{1}{d-1}})$, with a sufficiently large multiplicative constant, then*

$$\mathrm{rs}_d(m,n,I) = \Omega\left(\min\left\{\frac{I^{d+1}}{m^{d-1}n^d}, \frac{I^{d-1}}{m^{d-2}n^{d-2}}\right\}\right).$$

Proof: Let P be a set of m points in $\mathbb{R}^d$ and Π a set of n hyperplanes in $\mathbb{R}^d$, with $I = I(P,\Pi)$ in the assumed range. By Corollary 2.5, an appropriate choice of constants implies that most incidences are with hyperplanes of Π that are β_d-degenerate with respect to P. We map each hyperplane $\pi \in \Pi$, which is β_d-degenerate, to the set $T_\pi = P \cap \pi$, and let $\mathcal{T}$ be the multiset of all those T_π's. This multiset has $n' < n$ elements, all of which are $(\beta_d, d-2)$-degenerate, and $I' \geq I/2$ incidences. By Lemma 3.2, there are subsets $R \subseteq P$ and $\mathcal{S} \subseteq \mathcal{T}$, such that $R \subseteq S$ for each $S \in \mathcal{S}$ and

$$\begin{aligned} rs &= \Omega\left(\min\left\{\frac{(I')^{d+1}}{m^{d-1}(n')^d}, \frac{(I')^{d-1}}{m^{d-2}(n')^{d-2}}\right\}\right) \\ &= \Omega\left(\min\left\{\frac{I^{d+1}}{m^{d-1}n^d}, \frac{I^{d-1}}{m^{d-2}n^{d-2}}\right\}\right), \end{aligned}$$

where $r = |R|$ and $s = |\mathcal{S}|$.

We map each member $S \in \mathcal{S}$ back to the hyperplane $\pi \in \Pi$ that satisfies $S = T_\pi$ (by the multiset structure of $\mathcal{S}$, this inverse mapping can be assumed to be well defined). We denote the resulting set of hyperplanes by Σ. Then $G(R,\Sigma) = K_{r,s}$ and rs has the asserted lower bound. This completes the proof. □

We can now prove Theorem 1.2, which states that in the range $I = \Omega(mn^{1-\frac{1}{d-1}} + nm^{1-\frac{1}{d-1}})$, we have the lower bound

$$\mathrm{rs}_d(m,n) = \Omega\left(\left(\frac{I}{mn}\right)^{d-1} mn\right);$$

That is, in this range the minimum in the expression provided by Theorem 3.3 is attained by the second term.

Proof of Theorem 1.2: Let P be a set of m points and Π a set of n hyperplanes in $\mathbb{R}^d$, with $I = I(P,\Pi) = \Omega(mn^{1-\frac{1}{d-1}} + nm^{1-\frac{1}{d-1}})$ incidences. This lower bound is larger than the one required in Theorem 3.3. Indeed, we have $(mn)^{1-\frac{1}{d+1}} \leq mn^{1-\frac{1}{d-1}}$ when $n \leq m^{(d-1)/2}$, and, symmetrically, $(mn)^{1-\frac{1}{d+1}} \leq nm^{1-\frac{1}{d-1}}$ when $m \leq n^{(d-1)/2}$; since $(d-1)/2 \geq 1$, one of the latter inequalities must hold. Therefore, we have in this range

$$\mathrm{rs}_d(m,n,I) = \Omega\left(\min\left\{\frac{I^{d+1}}{m^{d-1}n^d}, \frac{I^{d-1}}{m^{d-2}n^{d-2}}\right\}\right).$$

However, the minimum is attained by the second term when $I = \Omega(mn^{1/2})$, which certainly holds for I in the assumed range, which therefore yields the asserted lower bound. □

Next, we give an upper bound construction showing that

$$\mathrm{rs}_d(m,n) = O\left(\left(\frac{I}{mn}\right)^{\frac{d+1}{2}} mn\right),$$

as asserted in Theorem 1.3.

Proof of Theorem 1.3: We start with the following d-dimensional structure, which is closely related to similar constructions of Elekes [El]. For arbitrary integers $k, l > 0$, let $P_{d,k,l}$ and $\Pi_{d,k,l}$ denote the following respective sets of points and hyperplanes in $\mathbb{R}^d$:

$$P_{d,k,l} = \{ (x_1, \ldots, x_d) \mid x_1, \ldots, x_{d-1} \in \{1, \ldots, k\}, \text{ and } x_d \in \{1, \ldots, dkl\} \},$$

$$\Pi_{d,k,l} = \left\{ x_d = \sum_{i=1}^{d-1} a_i x_i + b \;\middle|\; a_1, \ldots, a_{d-1} \in \{1, \ldots, l\}, \text{ and } b \in \{1, \ldots, kl\} \right\}.$$

Note that $|P_{d,k.l}| = dk^d l$ and $|\Pi_{d,k,l}| = kl^d$. For any hyperplane $\pi \in \Pi_{d,k,l}$, and for each choice of $x_1, \ldots, x_{d-1} \in \{1, \ldots k\}$, there is a point $(x_1, \ldots, x_{d-1}, x_d) \in P_{d,k,l} \cap \pi$. The set $P_{d,k,l} \cap \pi$ is thus a $(d-1)$-lattice isomorphic to the hypercube $\{1, \ldots, k\}^{d-1}$, and contains k^{d-1} points. Hence the number of incidences between $P_{d,k,l}$ and $\Pi_{d,k,l}$ is $I = k^d l^d$.

Each j-flat $F \subset \mathbb{R}^d$, which is the intersection of some $d-j$ or more hyperplanes of $\Pi_{d,k,l}$, is the image of some j-flat of the hypercube, as embedded into any of the hyperplanes $\pi \in \Pi_{d,k,l}$ that contain F. Since any j-flat of the hypercube contains at most k^j points, we have $|F \cap P_{d,k,l}| \leq k^j$. Furthermore, we have:

Observation 3.4 *A j-flat F as above is contained in at most l^{d-j-1} hyperplanes of $\Pi_{d,k,l}$.*

Proof: Any j-flat ($j < d$) in $\mathbb{R}^d$ is the image of some affine mapping $T : \mathbb{R}^j \to \mathbb{R}^d$, that is, $T(y) = My + v$, for some matrix $M \in \mathbb{R}^{d \times j}$, with rank $\rho(M) = j$, and vector $v \in \mathbb{R}^d$. Let M and v be such a pair of a matrix and a vector, so that the image of the affine mapping $y \mapsto My + v$ is F.

Let $\pi \in \Pi_{d,k,l}$ be a hyperplane containing F, given by the linear equation $x_d = \sum_{i=1}^{d-1} a_i x_i + b$, for some $a_1, \ldots, a_{d-1} \in \{1, \ldots, l\}$ and $b \in \{1, \ldots, kl\}$. Put $a_d = -1$, and $a = (a_1, \ldots, a_d) \in \mathbb{R}^d$. Thus we can write $\pi = \left\{ x \in \mathbb{R}^d \;\middle|\; a^{\mathrm{T}} x + b = 0 \right\}$.

Since $\pi \supset F$, we have $a^{\mathrm{T}}(My + v) + b = 0$ for all $y \in \mathbb{R}^j$. In particular, for $y = 0$, we have

$$a^{\mathrm{T}} v + b = 0.$$

This gives $a^{\mathrm{T}} M y = 0$, for all $y \in \mathbb{R}^j$, which is equivalent to

$$M^{\mathrm{T}} a = 0.$$

Thus, a is in the kernel of $M^{\mathrm{T}} \in \mathbb{R}^{j \times d}$. We have

$$\dim \mathrm{Ker}(M^{\mathrm{T}}) = d - \dim \mathrm{Im}(M^{\mathrm{T}}) = d - \rho(M) = d - j.$$

Hence a lies in the $(d-j)$-flat $K = \mathrm{Ker}(M^{\mathrm{T}})$. In addition, the requirement $a_d = -1$ constrains a to a hyperplane H. Note that $H \not\supset K$, since $0 \in K$, but $0 \notin H$. Hence a lies in the $(d-j-1)$-flat $K \cap H$. This flat can contain at most l^{d-j-1} points of the $l \times \cdots \times l \times 1$ lattice section. Hence there are at most l^{d-j-1} possible values of a. Once a has been determined, $b = -a^{\mathrm{T}} v$ is also uniquely determined. Thus, there are at most l^{d-j-1} possible hyperplanes $\pi \in \Pi_{d,k,l}$ containing F, and the observation is established. $\square$

By adding another dimension to the construction, an x_{d+1}-axis, we turn every point of $P_{d,k,l}$ into a line parallel to the x_{d+1}-axis, and every $(d-1)$-hyperplane of $\Pi_{d,k,l}$ into a d-hyperplane parallel to the x_{d+1}-axis. We denote the resulting set of lines by $P'_{d,k,l}$, and the set of d-hyperplanes by $\Pi'_{d,k,l}$. These sets have the same incidence relations as the original sets of points and $(d-1)$-hyperplanes. In particular, every j-flat in $\mathbb{R}^{d+1}$, which is the intersection of some $d-j+1$ or more d-hyperplanes of $\Pi'_{d,k,l}$ contains at most k^{j-1} lines of $P'_{d,k,l}$ (all parallel to the x_{d+1}-axis), and is contained in at most l^{d-j} d-hyperplanes of $\Pi'_{d,k,l}$.

To construct an example that attains the asserted bound $rs = O\left(\left(\frac{I}{mn}\right)^{\frac{d+1}{2}} mn\right)$, we proceed as follows. Let $P' = P'_{d-2,k,k}$ and $\Pi' = \Pi'_{d-2,k,k}$ be sets of $(d-2)k^{d-1}$ lines and k^{d-1} $(d-2)$-flats in $\mathbb{R}^d$ (these lines and flats are constructed in $\mathbb{R}^{d-1}$, but we embed them in a natural way in $\mathbb{R}^d$). For every line $\ell \in P'$, choose μ arbitrary points on ℓ, and let P denote the overall resulting set of points, and put $m = |P| = (d-2)\mu k^{d-1}$. For every $(d-2)$-flat $\pi' \in \Pi'$, choose ν distinct arbitrary hyperplanes, i.e., $(d-1)$-flats, containing π', and let Π denote the overall resulting set of hyperplanes. The hyperplanes are chosen so that no two hyperplanes containing two different flats from Π' coincide. Put $n = |\Pi| = \nu k^{d-1}$.

Now every hyperplane $\pi \in \Pi$ contains one flat $\pi' \in \Pi'$, which contains k^{d-3} lines of P', yielding a total of μk^{d-3} points of P incident to π. The number of incidences between P and Π is thus $I = \mu\nu k^{2d-4} = \Theta(k^{-2}mn)$, or $\frac{I}{mn} = \Theta(k^{-2})$. Note that the freedom of choice of the parameters k, μ and ν allows I to have almost any asymptotic value from $\Theta((mn)^{1-\frac{1}{d-1}})$ (choose $\mu = \nu = 1$) up to $\Theta(mn)$ (choose $k = 1$). In particular, we may assume $I = \Omega(mn^{1-\frac{2}{d+1}} + m^{1-\frac{2}{d+1}}n)$. Suppose now that $G(P, \Pi)$ contains a $K_{r,s}$ subgraph, that is, there exists some j-flat F (for some $j = 1, \ldots, d-2$) containing r points of P, and contained in s hyperplanes of Π. Without loss of generality, we may take F to be the intersection of these s hyperplanes. Thus, F is parallel to the x_{d+1}-axis, so any line of P' that meets F is fully contained in F. F contains at most k^{j-1} lines of P', hence, $r \le \mu k^{j-1}$. Also, F is contained in at most k^{d-j-2} flats of Π', hence, $s \le \nu k^{d-j-2}$. Altogether,

$$rs \le \mu\nu k^{d-3} = \underbrace{\mu k^{d-1}}_{\approx m} \cdot \underbrace{\nu k^{d-1}}_{=n} \cdot \underbrace{k^{-d-1}}_{\approx\left(\frac{I}{mn}\right)^{\frac{d+1}{2}}} = O\left(\left(\frac{I}{mn}\right)^{\frac{d+1}{2}} mn\right),$$

as claimed. □

We leave it as an open problem to close the gap between the bounds in Theorem 1.2 and Theorem 1.3, for $d \ge 4$.

4 Improved Compact Representation of Point-Hyperplane Incidence Graphs

In this section we obtain a compact bipartite representation of point-hyperplane incidence graphs, whose size is slightly smaller than that obtained by Brass and Knauer [BK].

4.1 Preliminaries

Let π be a hyperplane in $\mathbb{R}^d$, and let Δ be a k-dimensional convex polyhedron, for some $1 \le k \le d$. We say that π *crosses* Δ, if π intersects, but does not contain, Δ.

Let Π be a set of n hyperplanes in $\mathbb{R}^d$. A $\frac{1}{r}$-*cutting* for Π is a decomposition of $\mathbb{R}^d$ into a set Ξ of pairwise disjoint relatively open simplices of dimensions $0, \ldots, d$, so that each simplex is crossed by at most $\frac{n}{r}$ hyperplanes of Π.

The following lemma is a variant of well-known results (see [AE, Ch, CF, Ma]).

Lemma 4.1 *Let P be a set of m distinct points and let Π be a set of n distinct hyperplanes in $\mathbb{R}^d$. Then, for any $r < n$, there exists a $\frac{1}{r}$-cutting, Ξ, consisting of $O(r^d)$ relatively open pairwise disjoint simplices of dimensions $0, \ldots, d$, so that each cell of Ξ contains at most m/r^d points of P.*

For the sake of completeness, we sketch a proof of the lemma, even though it is similar to the proof in [CF]. We do so because the proof of Theorem 4.2 depends on several features of the construction, and requires a slight modification of it.

To construct a $\frac{1}{r}$-cutting for Π that satisfies the properties of the lemma, we proceed as follows. We choose a random sample $R_0 \subset \Pi$ of $c_0 r$ hyperplanes (c_0 being a sufficiently large constant), construct the arrangement $\mathcal{A}(R_0)$ and triangulate it into a set Ξ_0 of relatively open simplices. For each simplex $\Delta \in \Xi_0$, let $\Pi_\Delta \subset \Pi$ denote the set of hyperplanes that cross Δ. The *weight* of Δ is defined as $|\Pi_\Delta| \cdot \frac{r}{n}$. Note that a cell of weight ≤ 1 is crossed by at most $\frac{n}{r}$ hyperplanes of Π. Thus, we only need to handle cells of weight > 1. This we do by a refinement of the cutting within each of these cells.

For each $\Delta \in \Xi_0$ of weight $t > 1$, we choose an additional random sample $R_\Delta \subset \Pi_\Delta$ of $ct \log t$ hyperplanes among the ones crossing Δ, where c is a sufficiently large constant. Let $\mathcal{A}_\Delta(R_\Delta)$ denote the arrangement of R_Δ within Δ, i.e., the subdivision of Δ into cells by the hyperplanes of R_Δ. We triangulate the cells of $\mathcal{A}_\Delta(R_\Delta)$ into a set Ξ_Δ of simplices. For simplices $\Delta \in \Xi_0$ of weight ≤ 1, we simply define $\Xi_\Delta = \{\Delta\}$. Finally, we define $\Xi = \bigcup_{\Delta \in \Xi_0} \Xi_\Delta$ as the final triangulation. We denote by R the union of the samples R' in all the stages of the cutting construction.

By the theory of random sampling, it follows that, with high probability, Ξ is a $\frac{1}{r}$-cutting. Also, as is shown in [CF], the expected number of cells with weight $\ge t$ in Ξ_0 is $O(r^d 2^{-t})$. This, in turn, implies that the expected total number of cells in Ξ is $O(r^d)$.

In the original constructions [AE, Ch, CF, Ma] concerning hyperplanes in higher dimensions, the triangulation method employed is bottom-vertex triangulation. Here, however, we modify the construction as follows. Let R' be one of the random samples (within the entire space, or within some $\Delta \in \Xi_0$). For every cell τ of $\mathcal{A}(R')$ which is not a simplex, let v be a point in the relative interior of τ and let ξ be a triangulation of the relative boundary of τ. We extend every $(k-1)$-dimensional simplex $\Delta' \in \xi$ to the k-dimensional simplex Δ which is the convex hull of Δ' and v. ξ itself is constructed recursively using the same method in lower dimensions. The number of new vertices introduced in the triangulation process is at most the number of faces of all dimensions in $\mathcal{A}(R')$. Summing over all the random samples R' of the different stages of triangulation, gives a total of $O(r^d)$ new vertices, and the total complexity of Ξ remains $O(r^d)$. The theory of random sampling implies that the resulting triangulation preserves (with high probability) the properties of the cutting.

Consider now the triangulation Ξ together with a given set of points P. A k-face (i.e., a k-

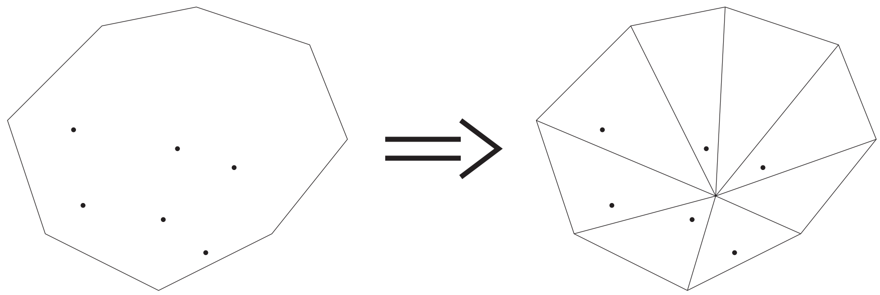

Figure 2: Triangulation by an interior point.
The interior vertex is chosen so that the new faces of the triangulation avoid the points of P that lie in the interior of the cell.

dimensional face) of a cell of Ξ may be either a simplex portion of a k-face of $\mathcal{A}(R')$, for some of the random samples R' used to construct Ξ, or a new face created during the triangulation process by taking the convex hull of an interior point v of a j-dimensional cell τ, for $j > k$, and a $(k-1)$-dimensional simplex Δ' on the relative boundary of τ.

Now a crucial point needed for our analysis is that the points of P must not lie on these new faces. However, this can easily be achieved by an appropriate choice of the interior points v. See Figure 2. Finally, each cell τ of Ξ that contains more than m/r^d points of P is further partitioned, by cutting it with a generic collection of hyperplanes, into subcells, each containing at most m/r^d points. Clearly the new cells still constitute a $\frac{1}{r}$-cutting of Π, and their overall number is still $O(r^d)$. We thus obtain a $\frac{1}{r}$-cutting consisting of $O(r^d)$ simplices with the above properties, including the property that the new faces added in the triangulation do not contain points of P.

Finally, we note that the construction of Ξ is randomized. However, since our analysis is non-algorithmic, we simply assume that each of the samples taken at each step of the construction meets or beats the expected values of the various parameters that it controls.

To summarize, in addition to the properties asserted in the lemma itself, the cutting has the property that a k-dimensional simplex of Ξ contains points of P only if it is contained in the intersection of (at least) $d-k$ sample hyperplanes of Π.

4.2 Constructing a compact representation of the incidence graph

Theorem 4.2

$$J_d(m,n) = O((mn)^{1-\frac{1}{d+1}} + m + n).$$

Proof: Let $P \subset \mathbb{R}^d$ be a set of m points, and let Π be a set of n hyperplanes. We shall construct a representation of $G(P,\Pi)$ with the claimed complexity.

Fix a parameter $r \leq \min\{m^{1/d}, n\}$, whose value will be determined later, and construct, using Lemma 4.1, a $\frac{1}{r}$-cutting Ξ of $\mathbb{R}^d$, that consists of $t = O(r^d)$ relatively open simplices $\Delta_1, \ldots, \Delta_t$, of dimensions $0, \ldots, d$, so that each Δ_i contains $m_i \leq m/r^d$ points of P, and is crossed by $n_i \leq n/r$ hyperplanes of Π.

For each $i = 1, \ldots, t$, put $P_i = P \cap \Delta_i$, and let Π_i (resp., Π_i^*) denote the set of all hyperplanes of Π that cross Δ_i (resp., fully contain Δ_i). We clearly have

$$G(P, \Pi) = \left(\bigcup_{i=1}^{t} P_i \times \Pi_i^* \right) \cup \left(\bigcup_{i=1}^{t} G(P_i, \Pi_i) \right).$$

However, we can compactify the first union still further. For each $k = 0, \ldots, d-1$, we consider all the k-faces of Ξ that contain points of P, and all the k-flats obtained as the intersection of $d-k$ affinely independent hyperplanes of R. We lump together all the k-faces that lie on the same k-flat, f, into a single common group, and note that any hyperplane that contains one of these faces contains all of them. (Note that lower-dimensional faces on f are not part of this group, but are collected in the appropriate lower-dimensional sub-flats of f. Note also that any k-face of Ξ that contains points of P is of this form, by construction.) We repeat this process for all dimensions k, and can thus write

$$G(P, \Pi) = \left(\bigcup_{f} P_f \times \Pi_f^* \right) \cup \left(\bigcup_{i=1}^{t} G(P_i, \Pi_i) \right),$$

where the first union is over all k-flats f, of dimensions $k = 0, \ldots, d-1$, which are the intersection of at least $d-k$ sample hyperplanes of R and which contain k-faces of Ξ with points of P. Here P_f is the set of all points of P that lie in f and are not contained in lower dimensional faces of Ξ, and Π_f^* is the set of all hyperplanes of Π that contain f.

Hence, we obtain the recurrence

$$J(P, \Pi) \le \sum_{f} (|P_f| + |\Pi_f^*|) + \sum_{i=1}^{t} J(P_i, \Pi_i),$$

where the first summation is only over those f for which the corresponding set P_f is nonempty. Since $|P_i| \le m/r^d$, and $|\Pi_i| \le n/r$, for each $i = 1, \ldots, t$, the last sum in the right-hand side can be upper bounded by

$$O(r^d) \cdot J_d \left(\frac{m}{r^d}, \frac{n}{r} \right).$$

Consider next the first sum $\sum_f (|P_f| + |\Pi_f^*|)$. Since the (open) cells of Ξ are pairwise disjoint, we have

$$\sum_{f} |P_f| \le m.$$

To estimate $\sum_f |\Pi_f^*|$, we argue as follows. Let R_0 denote the set of $O(r)$ hyperplanes in the initial random sample. Each hyperplane $\pi \in \Pi$ can be incident to at most $O(r^{d-k-1})$ of the k-flats of $\mathcal{A}(R_0)$. This is because the hyperplanes of R_0 intersect π in at most $|R_0|$ $(d-2)$-flats (or $|R_0| - 1$ if $\pi \in R_0$) and induce on π a $(d-1)$-dimensional arrangement of these flats, which has at most $\binom{|R_0|}{d-k-1} = O(r^{d-k-1})$ k-flats. Since every k-flat of $\mathcal{A}(R_0)$, which is contained in π, appears as a k-flat of the induced arrangement, the claim follows. Any k-face φ of Ξ, that is the intersection of $d-k$ sample hyperplanes and has not been counted so far, must fully lie inside some cell $\Delta \in \Xi_0$ of the first stage of the cutting, and all the sample hyperplanes that form φ are those sampled in Δ (including the $O(1)$ hyperplanes that define Δ). Suppose that Δ has weight t. Then there are $ct \log t$ sample hyperplanes in Δ, and, arguing as above, they form $O((t \log t)^{d-k-1})$ k-flats of the kind that we consider on each hyperplane $\pi \in \Pi$ that crosses Δ. Since there are only tn/r such hyperplanes, the total number of such additional k-faces within Δ is $O(\frac{tn}{r}(t \log t)^{d-k-1})$. Summing over all cells $\Delta \in \Xi_0$, and rounding up each weight to the nearest integer, we obtain a total of

$$O\left(\sum_{j\geq 2} \frac{jn}{r}(j\log j)^{d-k-1} r^d 2^{-j}\right) = O\left(nr^{d-1}\sum_{j\geq 2} j(j\log j)^{d-k-1}2^{-j}\right) = O\left(nr^{d-1}\right),$$

where we use the fact that the expected number of cells of weight $j \geq t \geq j-1 \geq 1$ is $O(r^d 2^{-j})$ [CF], and assume that the samples in our construction meet (or beat) these expectations. Hence we have

$$\sum_f |\Pi_f^*| = \sum_{k=0}^{d-1} O(nr^{d-1}) = O(nr^{d-1}).$$

Putting everything together, we obtain the following recurrence:

$$J_d(m,n) = O(r^d)\cdot J_d\left(\frac{m}{r^d}, \frac{n}{r}\right) + O(nr^{d-1} + m). \tag{5}$$

We next note that, using point-hyperplane duality, $J_d(m,n)$ is symmetric, i.e., $J_d(m,n) = J_d(n,m)$. In particular, restricting r further to satisfy $r \leq \min\left\{m^{1/(d+1)}, n^{1/(d+1)}\right\}$, we have

$$J_d\left(\frac{m}{r^d}, \frac{n}{r}\right) = J_d\left(\frac{n}{r}, \frac{m}{r^d}\right) = O\left(r^d J_d\left(\frac{n}{r^{d+1}}, \frac{m}{r^{d+1}}\right) + \frac{m}{r} + \frac{n}{r}\right).$$

Substituting that in (5) we get

$$\begin{aligned} J_d(m,n) &= O\left(r^d\left(r^d J_d\left(\frac{m}{r^{d+1}}, \frac{n}{r^{d+1}}\right) + \frac{m}{r} + \frac{n}{r}\right) + nr^{d-1} + m\right) \\ &= O\left(r^{2d} J_d\left(\frac{m}{r^{d+1}}, \frac{n}{r^{d+1}}\right) + (m+n)r^{d-1}\right). \end{aligned} \tag{6}$$

Lemma 4.3 *The solution of* (5) *and* (6) *is*

$$J_d(m,n) = O((mn)^{1-\frac{1}{d+1}} + m + n).$$

Proof: Without loss of generality, we may assume that $m \geq n$. Put $s = m/n \geq 1$, and rewrite (5) as

$$J_d(m,n) = O\left(r^d J_d\left(\frac{sn}{r^d}, \frac{n}{r}\right) + nr^{d-1} + sn\right). \tag{7}$$

We then choose $r = s^{\frac{1}{d-1}}$ and get

$$J_d(m,n) = O\left(s^{\frac{d}{d-1}} J_d\left(\frac{n}{s^{1/(d-1)}}, \frac{n}{s^{1/(d-1)}}\right) + sn\right). \tag{8}$$

Note that for this choice of parameters to make sense we must have $s^{1/(d-1)} = (m/n)^{1/(d-1)} \leq \min\{m^{1/d}, n\}$, or $m \leq n^d$. If $m > n^d$ we argue as follows. The complexity of a compact representation of the incidence graph can be bounded exactly as in the analysis that has led to the derivation of (5), but using Π as the entire "sample", i.e., with $r = n$. This yields the bound $J_d(m,n) = O(m + nr^{d-1}) = O(m + n^d) = O(m)$. Hence, $J_d(m,n) = O(m)$ when $m > n^d$.

We may thus assume that $m \leq n^d$, and estimate $J_d\left(\frac{n}{s^{1/(d-1)}}, \frac{n}{s^{1/(d-1)}}\right)$ as follows. Apply (6) with $m = n = q$, to obtain

$$J_d(q,q) = O\left(r^{2d} J_d\left(\frac{q}{r^{d+1}}, \frac{q}{r^{d+1}}\right) + qr^{d-1}\right).$$

Choosing $r = q^{1/(d+1)}$, we obtain

$$J_d(q,q) = O\left(q^{2d/(d+1)} J_d(1,1) + q^{2d/(d+1)}\right) = O\left(q^{2d/(d+1)}\right).$$

Substituting this solution in (8), with $q = n/s^{1/(d-1)}$and taking into account also the case $m > n^d$, we get

$$J_d(m,n) = O\left(s^{\frac{d}{d-1}} \left(\frac{n}{s^{1/(d-1)}}\right)^{\frac{2d}{d+1}} + sn + m\right) = O((mn)^{1-\frac{1}{d+1}} + m). \tag{9}$$

Symmetrically, for $n \geq m$ we get

$$J_d(m,n) = O((mn)^{1-\frac{1}{d+1}} + n). \tag{10}$$

Combining (9) and (10), we get the desired bound. □

This completes the proof of Theorem 4.2. □

Remarks:

1. The best known lower bound constructions, presented in [BK], are point-hyperplane configurations with:
   - $J_d(m,n) = \Omega\left((mn)^{1-\frac{2}{d+3}-\varepsilon}\right)$ for d odd, and for any $\varepsilon > 0$,
   - $J_d(m,n) = \Omega\left((mn)^{1-\frac{2(d+1)}{(d+2)^2}-\varepsilon}\right)$ for d even, and for any $\varepsilon > 0$, and
   - $J_d(m,n) = \Omega(m^{7/10}n^{7/10})$ for $d = 3$.

   These bounds apply for the balanced case, where $m \approx n$. For the non-balanced case, there are slightly better lower bounds; see [BK]. Closing the gaps between these bounds and our upper bound remains an open problem. See also Appendix A for more details concerning the 3-dimensional construction.

2. One corollary of Theorem 4.2, already noted in the introduction, is that if the incidence graph does not contain a $K_{r,s}$ subgraph, i.e., no r points lie inside any intersection of s hyperplanes, then the number of incidences is $O((mn)^{1-\frac{1}{d+1}} + m + n)$, where the constant of proportionality depends on r and s. This slightly improves (by a logarithmic factor) a similar observation in [BK].

5 Conclusion

We have studied the structure of point-hyperplane incidence graphs, and have shown that whenever the number of incidences is large, the incidence graph contains large complete bipartite subgraphs. Specifically,

1. We have obtained an improved upper bound on the overall size of a compact representation of the incidence graph as the union of complete bipartite subgraphs (Theorem 4.2).

2. We have derived lower bounds on the number of edges in the largest complete bipartite incidence subgraph in three dimensions (Theorems 1.1) and in higher dimensions (Theorem 1.2).

3. We have obtained matching upper bound constructions for these lower bounds. The three-dimensional constructions (Lemma 2.2, Lemma 2.3, and Lemma 2.6) are worst-case tight, whereas the higher-dimensional one (Theorem 1.3) is not known to be tight.

4. For each of these bounds, we have provided an estimate of how many incidences must there be in order to ensure the existence of large complete bipartite incidence subgraphs that attain the asserted lower bounds. The three-dimensional estimates are tight, whereas the higher-dimensional ones are not known to be tight.

We leave as open problems:

(i) To close the gap between the higher-dimensional bounds on the number of edges in the largest complete bipartite point-hyperplane incidence subgraph.

(ii) To establish a tight bound on the representation complexity of point-hyperplane incidence graphs. Even the case $d = 3$ is still open.

(iii) To find a maximal point-hyperplane incidence graph that does not contain some fixed complete bipartite subgraph, or to prove that the known constructions of [BK] for this problem are maximal. It is not clear whether or not this problem is equivalent to problem (ii).

References

[AA] P. K. Agarwal and B. Aronov, Counting facets and incidences, *Discrete Comput. Geom.* 7 (1992), 359–369.

[AE] P. K. Agarwal and J. Erickson, Geometric range searching and its relatives, in: *Advances in Discrete and Computational Geometry* (B. Chazelle, J. E. Goodman and R. Pollack, eds.), AMS Press, Providence, RI, 1998, 1–56.

[BHPT] I. Bárány, G. Harcos, J. Pach and G. Tardos, Covering lattice points by subspaces, *Periodica Mathematica Hungarica* 43 (2001), 93–103.

[Bo] B. Bollobas, *Extremal Graph Theory*, London Academic Press, 1978.

[BK] P. Brass and C. Knauer, On counting point-hyperplane incidences, *Comput. Geom. Theory Appls.* 25 (2003), 13–20.

[Ch] B. Chazelle, Cutting hyperplanes for divide-and-conquer, *Discrete Comput. Geom.* 9 (1993), 145–158.

[CF] B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica* 10 (1990), 229–249.

[EGS] H. Edelsbrunner, L. Guibas and M. Sharir, The complexity of many cells in arrangements of planes and related problems, *Discrete Comput. Geom.* 5 (1990), 197–216.

[El] G. Elekes, Sums versus products in number theory, algebra and Erdős geometry — a survey, *Paul Erdős and his Mathematics II*, Bolyai Math. Soc. Stud. 11, Budapest, 2002, 241–290.

[ET] G. Elekes and C. D. Tóth, Incidences of not too degenerate hyperplanes, *Proc. 21st Annu. ACM Sympos. Comput. Geom.*, 2005, 16–21.

[Ma] J. Matoušek, Cutting hyperplane arrangements, *Discrete Comput. Geom.* 6 (1991), 385–406.

[PS] J. Pach and M. Sharir, Geometric incidences, in *Towards a Theory of Geometric Graphs* (J. Pach, ed.), *Contemporary Mathematics*, Vol. 342, Amer. Math. Soc., Providence, RI, 2004, 185–223.

[ST] E. Szemerédi and W. Trotter, Extremal problems in discrete geometry, *Combinatorica* 3 (1983), 381–392.

Appendix A: Incidences Between Points and Planes in $\mathbb{R}^3$ with no Three Collinear Planes

An upper bound on the number of incidences between m points and n planes in $\mathbb{R}^3$ with no three collinear planes (and a symmetric bound for the dual problem, where no three points are collinear) has been known for a while. As discussed in the introduction, we attribute the result to [EGS]; the bound there is slightly weaker, but can be cleaned-up using a more careful analysis.

Theorem 5.1 (Edelsbrunner *et al.* [EGS]) *Let $P \subset \mathbb{R}^3$ be a set of m points and let Π be a set of n planes in $\mathbb{R}^3$, such that no three planes of Π are collinear. Then the number of incidences is bounded by*

$$I(P, \Pi) = O(m^{4/5}n^{3/5} + m + n).$$

The symmetric bound $I(P, \Pi) = O(m^{3/5}n^{4/5} + m + n)$ holds in the dual case, where no three points of P are collinear.

The proof of the first bound uses the fact that if no three planes are collinear, then the incidence graph does not contain a $K_{2,3}$, i.e., two distinct points lying in the intersection of three distinct planes. Note that the converse is not true, i.e., we can construct point-plane configurations with no $K_{2,3}$, but with (many) triples of collinear planes. Thus, Edelsbrunner *et al.* have implicitly proved a slightly stronger statement, whose proof follows the one in [EGS] almost verbatim.

Theorem 5.2 *Let $P \subset \mathbb{R}^3$ be a set of m points and let Π be a set of n planes in $\mathbb{R}^3$, such that $G(P, \Pi)$ does not contain a $K_{2,3}$ subgraph. Then the number of incidences is bounded by*

$$I(P, \Pi) = O(m^{4/5}n^{3/5} + m + n).$$

The symmetric bound $I(P, \Pi) = O(m^{3/5}n^{4/5} + m + n)$ holds in the dual case, where $G(P, \Pi)$ does not contain a $K_{3,2}$ subgraph.

Recently, Brass and Knauer [BK] constructed an example that effectively shows that these bounds are worst-case tight. For the sake of completeness, we repeat (and slightly modify) their construction here. It relies on the following result.

Theorem 5.3 (Bárány *et al.* [BHPT]) *Let Q be a subset of the integer lattice in $\mathbb{R}^3$ contained in the ball of radius r centered at the origin. Assume further that every three distinct vectors of Q are linearly independent, and that Q is a maximal set satisfying this property. Then*

$$|Q| = \Theta\left(r^{3/2}\right).$$

Theorem 5.4 (Brass and Knauer [BK]) *For any m and n, such that $m = O(n^3)$, there exist a set P of m points and a set Π of n planes in $\mathbb{R}^3$, with no three collinear planes, such that*

$$I(P, \Pi) = \Omega(m^{4/5}n^{3/5}).$$

Proof: Let $P = \{1, \ldots, m^{1/3}\}^3$ be an $m^{1/3} \times m^{1/3} \times m^{1/3}$ lattice section. Put $r = \Theta(n^{2/5}m^{-2/15})$, and let Q be a maximal lattice subset of the ball of radius r about the origin that satisfies the property in Theorem 5.3, i.e., every three vectors of Q are linearly independent, and $|Q| = \Theta(r^{3/2})$. Note that for our assumed range of m and n, we have $r > 1$, with an appropriate choice of the constants of proportionality. For each point $p \in P$, and for each vector $q \in Q$, construct a plane through p normal to q; its equation is $x \cdot q = p \cdot q$. Let Π denote the resulting set of planes. Since each coordinate of p is an integer $\leq m^{1/3}$, and each coordinate of q is an integer $\leq r$, there are $O(m^{1/3}r)$ distinct values of $p \cdot q$, and the number of planes is thus $|\Pi| = |Q| \cdot O(m^{1/3}r) = O(n)$. The number of incidences between P and Π is $I(P, \Pi) = |P| \cdot |Q| = \Theta(mr^{3/2}) = \Theta(m^{4/5}n^{3/5})$, and no three planes are collinear. Indeed, suppose there were three collinear planes in Π with normals $q_1, q_2, q_3 \in Q$. These normals are all distinct, and lie in the plane through the origin normal to the intersection line of the three planes, and are thus linearly dependent — a contradiction. □

Interestingly, this construction, when transformed to dual space, does not have the dual property that no three points are collinear. This is because the duals of three parallel planes are three collinear points, and the construction does contain many triples of parallel planes. Thus, the problem of obtaining a tight bound on the number of incidences between m points, no three of which are collinear, and n planes in $\mathbb{R}^3$, remains open. Nevertheless, the following somewhat weaker result, which follows from the dual construction, holds:

Corollary 5.5 *The maximum number of incidences between m points and n planes in $\mathbb{R}^3$, such that no three collinear points lie in two or more common planes, is $\Theta(m^{3/5}n^{4/5} + m + n)$.*

In other words, both primal and dual versions of Theorem 5.2 yield bounds that are worst-case tight. In contrast, the bound in the primal version of Theorem 5.1 is worst-case tight, but the bound in the dual version is not known to be tight.