An Improved Bound for Weak Epsilon-Nets in the Plane*

Natan Rubin†

July 22, 2022

Abstract

We show that for any finite point set P in the plane and $\epsilon > 0$ there exist $O\left(\frac{1}{\epsilon^{3/2+\gamma}}\right)$ points in $\mathbb{R}^2$, for arbitrary small $\gamma > 0$, that pierce every convex set K with $|K \cap P| \geq \epsilon|P|$. This is the first improvement of the bound of $O\left(\frac{1}{\epsilon^2}\right)$ that was obtained in 1992 by Alon, Bárány, Füredi and Kleitman for general point sets in the plane.

*The project leading to this application has received funding from European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme under grant agreement No. 678765. An extended abstract of this work has appeared in Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science, October 7-9, 2018, Paris, France.

†Email: rubinnat.ac@gmail.com. Ben Gurion University of the Negev, Beer-Sheba. Also supported by grant 1452/15 from Israel Science Foundation and by grant 2014384 from the U.S.-Israeli Binational Science Foundation.

1 Introduction

Transversals and ϵ-nets. Given a family $\mathcal{K}$ of geometric ranges in $\mathbb{R}^d$ (e.g., lines, triangles, or convex sets), we say that $Q \subset \mathbb{R}^d$ is a transversal to $\mathcal{K}$ (or Q pierces $\mathcal{K}$) if each $K \in \mathcal{K}$ is pierced by at least one point of Q. Given an underlying set P of n points, we say that a range $K \in \mathcal{K}$ is *ϵ-heavy* if $|P \cap K| \geq \epsilon n$. We say that Q is an *ϵ-net* for $\mathcal{K}$ if it pierces every ϵ-heavy range in $\mathcal{K}$. We say that an ϵ-net for $\mathcal{K}$ is a *strong ϵ-net* if $Q \subset P$, that is, the points of the net are drawn from the underlying point set P. Otherwise (i.e., if Q includes additional points outside P), we say that Q is a *weak ϵ-net*.

The study of ϵ-nets was initiated by Vapnik and Chervonenkis [49], in the context of statistical learning theory. Following a seminal paper of Haussler and Welzl [30], ϵ-nets play a central role in discrete and computational geometry [35]. For example, bounds on ϵ-nets determine the performance of the best-known algorithms for minimum hitting set/set cover problem in geometric hypergraphs [8, 13, 23, 24], and the transversal numbers of families of convex sets [4, 5, 7, 33].

Informally, the cardinality of the smallest possible ϵ-net for the range set $\mathcal{K}$ determines the integrality gap of the corresponding transversal problem – the ratio between (1) the size of the smallest possible transversal Q to $\mathcal{K}$ and (2) the weight of the "lightest" possible fractional transversal to $\mathcal{K}$ [7, 4, 24].

Haussler and Welzl [30] proved in 1986 the existence of strong ϵ-nets of cardinality $O\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ for families of so called semi-algebraic ranges of bounded description complexity in d-space[1], for a fixed $d > 0$ (e.g., lines, boxes, spheres, halfspaces, or simplices), by observing that their induced hypergraphs have a bounded Vapnik-Chervonenkis dimension (so called *VC-dimension*).[2] While the bound is generally tight for set systems with a bounded VC-dimension [34], better constructions are known for several special families of geometric ranges, including tight bounds for discs in $\mathbb{R}^2$, halfplanes in $\mathbb{R}^2$ and halfspaces in $\mathbb{R}^3$ [23, 34, 39], and rectangles in $\mathbb{R}^2$ and boxes in $\mathbb{R}^3$ [8, 43]. We refer the reader to a recent state-of-the-art survey [42] for the best known bounds.

In particular, it had long been conjectured that all the "natural" geometric instances, that involve simply-shaped geometric ranges in a fixed-dimensional Euclidean space $\mathbb{R}^d$, admit a strong ϵ-net of cardinality $O(1/\epsilon)$ (where the constant of proportionality depends on the VC-dimension). The conjecture was refuted in 2010 by Alon [3] who used a density version of Hales-Jewett Theorem [25] to show that some families of geometric ranges (e.g., lines in the Euclidean plane) require an ϵ-net whose cardinality is larger than $O(1/\epsilon)$. Pach and Tardos [43] subsequently demonstrated that the multiplicative term $\Theta(\log 1/\epsilon)$ is necessary for strong ϵ-nets with respect to halfspace ranges in dimension higher than 3.

Weak ϵ-nets for convex sets. In sharp contrast to the case of simply-shaped ranges, no constructions of small-size strong ϵ-nets exist for general families of convex sets in $\mathbb{R}^d$, for $d \geq 2$. For example, given an underlying set of n points in convex position in $\mathbb{R}^2$, any strong ϵ-net with respect to convex ranges must include at least $n - \epsilon n$ of the points. Informally, this phenomenon can be attributed to the fact that the VC-dimension of a geometric set system is closely related to the *description complexity* of the underlying ranges, and it is unbounded for general convex sets. Nevertheless, Bárány, Füredi and Lovász [12] observed in 1990 that families of convex sets in $\mathbb{R}^2$ still admit weak ϵ-nets of cardinality $O(\epsilon^{-1026})$. Alon, Bárány, Füredi, and Kleitman [2] were the first to show in 1992 that families of convex sets in any dimension $d \geq 1$ admit weak ϵ-nets

[1] Namely, each of these sets is the locus of all the points that satisfy a given Boolean combination of a bounded number of algebraic equalities and inequalities of bounded degree in the Euclidean coordinates $x_1, \ldots, x_d$ [1].

[2] The constant hidden within the $O(\cdot)$-notation is specific to the family of geometric ranges under consideration, and is proportional to the VC-dimension of the induced hypergraph.

whose cardinality is bounded in terms of $1/\epsilon$ and d. The subsequent study and application of weak ϵ-nets bear strong relations to convex geometry, including Helly-type, Centerpoint and Selection Theorems; see [36, Sections 8 – 10] for a comprehensive introduction.

Weak ϵ-nets and the Hadwiger-Debrunner Problem. Alon and Kleitman [7] used the boundedness of weak ϵ-nets to confirm a long-standing (p,q)*-conjecture* by Hadwiger and Debrunner [27]. To this end, we say that a family $\mathcal{K}$ of convex sets satisfies the (p,q)*-property* if any its p-size subfamily $\mathcal{K}' \subset \mathcal{K}$ contains a q-size subset $\mathcal{K}'' \subset \mathcal{K}$ with a non-empty common intersection $\bigcap \mathcal{K}'' \neq \emptyset$. Hadwiger and Debrunner conjectured (and Alon and Kleitman showed) that for every positive integers p,q and d that satisfy $p \geq q \geq d+1$, there exists an integer $C_d(p,q) < \infty$ so that the following statement holds: Any family $\mathcal{K}$ of convex sets in $\mathbb{R}^d$ with the (p,q)-property admits a transversal by at most $C_d(p,q)$ points.[3] Showing good quantitative estimates for the Hadwiger-Debrunner numbers $C_d(p,q)$ requires tight asymptotic bounds for weak ϵ-nets; see the latest study by Keller, Smorodinsky and Tardos [33], and the concluding discussion in Section 4.

Very recently, lower bounds for several of the above questions – including strong and weak ϵ-nets with respect to line ranges [11], and the 2-dimensional Hadwiger-Debrunner numbers $C_2(p,q)$ [32] – were improved using the novel combinatorial machinery of hypergraph containers [10, 45].

Bounds on weak ϵ-nets. For any $\epsilon > 0$ and $d \geq 0$, let $f_d(\epsilon)$ be the smallest number $f > 0$ so that, for any underlying finite point set P, one can pierce all the ϵ-heavy convex sets using only f points in $\mathbb{R}^d$. It is an outstanding open problem in Discrete and Computational geometry to determine the true asymptotic behaviour of $f_d(\epsilon)$ in dimensions $d \geq 2$. As Alon, Kalai, Matoušek, and Meshulam noted in 2001: "*Finding the correct estimates for weak ϵ-nets is, in our opinion, one of the truly important open problems in combinatorial geometry*" [5].

Alon *et al.* [2] (see also [7]) used Tverberg-type results to show that $f_d(\epsilon) = O(1/\epsilon^{d+1-1/\beta_d})$ (where $0 < \beta_d < 1$ is a selection ratio which is fixed for every d), and $f_2(\epsilon) = O\left(1/\epsilon^2\right)$.[4] The bound in higher dimensions $d \geq 3$ has been subsequently improved in 1993 by Chazelle *et al.* [19] to roughly $\tilde{O}\left(\frac{1}{\epsilon^d}\right)$ (where $\tilde{O}(\cdot)$-notation hides multiplicative factors that are polylogarithmic in $\log 1/\epsilon$). Though the latter construction was somewhat simplified in 2004 by Matoušek and Wagner [38] using simplicial partitions with low hyperplane-crossing number [37], no improvements in the upper bound for general families of convex sets and arbitrary finite point sets occurred for the last 25 years, in any dimension $d \geq 2$.

In view of the best known lower bound of $\Omega\left(\frac{1}{\epsilon}\log^{d-1}\left(\frac{1}{\epsilon}\right)\right)$ for $f_d(\epsilon)$ due to Bukh, Matoušek and Nivasch [14], it still remains to settle whether the asymptotic behaviour of this quantity substantially deviates from the long-known "almost-$(1/\epsilon)$" bounds on strong ϵ-nets (e.g., for lines and triangles in $\mathbb{R}^2$ or simplices in $\mathbb{R}^d$)?

The only interesting instances in which the gap has been essentially closed, involve special point sets [19, 15, 6]. For example, Alon *et al.* [6] showed in 2008 that any finite point set in *a convex position in* $\mathbb{R}^2$ allows for a weak ϵ-net of cardinality $O\left(\alpha(\epsilon)/\epsilon\right)$ with respect to convex sets, where $\alpha(\cdot)$ denotes the inverse Ackerman function.

Our result and organization. We provide the first improvement[5] of the general bound in $\mathbb{R}^2$.

[3] The celebrated Helly Theorem yields a transversal by a *single* point in the case $p = q = d+1$ whenever $|\mathcal{K}| \geq d+1$.

[4] An outline of the planar $f_2(\epsilon) = O\left(1/\epsilon^2\right)$ bound can be found in a popular textbook by Chazelle [18].

[5] A follow-up study of the author [44] establishes a similar improvement in all dimensions $d \geq 3$. Its remarkable that the higher-dimensional argument does not yield Theorem 1.1 for $d = 2$; see Section 4 for a brief comparison.

Theorem 1.1. *We have*

$$f_2(\epsilon) = O\left(\frac{1}{\epsilon^{3/2+\gamma}}\right),$$

for any $\gamma > 0$.

That is, for any underlying set P of n points in $\mathbb{R}^2$, and any $\epsilon > 0$, one can construct a weak ϵ-net with respect to convex sets whose cardinality is $O\left(\frac{1}{\epsilon^{3/2+\gamma}}\right)$; here $\gamma > 0$ is an arbitrary small constant which does not depend on ϵ.[6]

The main underlying idea of our proof of Theorem 1.1 is that one can find a small auxiliary net Q', of $o\left(1/\epsilon^2\right)$ points, with the property that every ϵ-heavy convex set K that is missed by Q' is "line-like" – the subset $P \cap K$ is largely determined by an edge that connects some pair of its points; in the sequel, these sets are called narrow. (As a matter of fact, the number of such "proxy" edges for a narrow set K is close to $\binom{\epsilon n}{2}$.) These narrow sets are pierced by a careful adaptation of the decomposition paradigm that was previously used to bound the number of point-line incidences in the plane [48, 22].

The rest of the paper is organized as follows:

In Section 2 we provide a comprehensive overview of our approach, lay down the recursive framework, and establish several basic properties that are used throughout the proof of Theorem 1.1.

In Section 3 we use the recursive framework of Section 2 to give a constructive proof of Theorem 1.1. The eventual net combines the following elementary ingredients: (1) vertices of certain trapezoidal decompositions of $\mathbb{R}^2$, (2) 1-dimensional $\hat{\epsilon}$-nets, for $\hat{\epsilon} = \omega(\epsilon^2)$, which are constructed within few vertical lines with respect to carefully chosen point sets, and (3) strong $\hat{\epsilon}$-nets with respect to triangles in $\mathbb{R}^2$.

In Section 4 we briefly summarize the properties of our construction and survey the future lines of work.

2 Preliminaries

2.1 Proof outline

We briefly outline the main ideas behind our proof of Theorem 1.1. We begin by sketching the $O\left(1/\epsilon^2\right)$ planar construction of Alon *et al.* [2] (or, rather, its more comprehensive presentation by Chazelle [18]).

The quadratic construction. Refer to Figure 1 (left). We split the underlying point set P by a vertical median line L into subsets P^- and P^+ (of cardinality $n/2$ each), and recursively construct a weak $(4\epsilon/3)$-net with respect to each of these sets. Let K be an ϵ-heavy convex set; in what follows, we assign to each such set K a unique subset $P_K \subseteq P \cap K$ of exactly $\lceil \epsilon n \rceil$ points. If at least $3\epsilon n/4$ points of P_K lie to the same side of L, we pierce K by one of the auxiliary $(4\epsilon/3)$-nets. Otherwise, the points of P_K span at least $\epsilon^2 n^2/16$ edges that cross $K \cap L$, so we can pierce K by adding to our net each $(\epsilon^2 n^2/16)$-th crossing point of L with the edges of $\binom{P}{2}$.[7]

[6] We do not seek to optimize the implicit constant of proportionality within $O(\cdot)$, which heavily depends on $\gamma > 0$. The particular recurrence scheme that we establish for $f_2(\epsilon)$ in Section 3.6 results in a constant that is super-exponential in $1/\gamma$.

[7] In the sequel we use $\binom{A}{2}$ to denote the complete set of edges spanned by a (finite) point set $A \subset \mathbb{R}^2$.

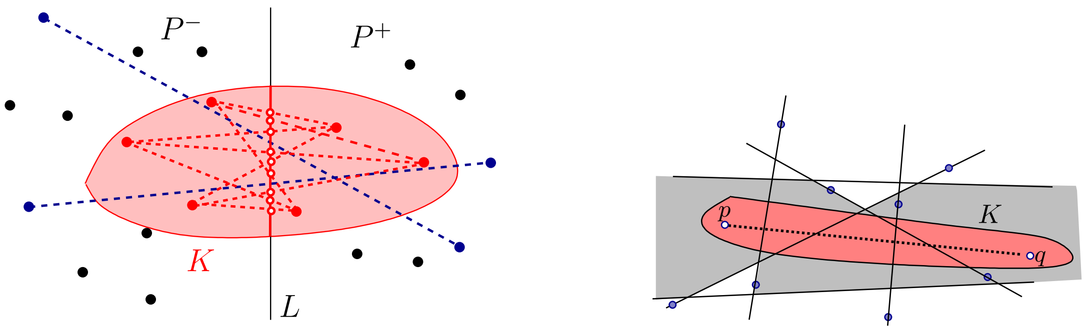

Figure 1: Left: Constructing the net of cardinality $O(1/\epsilon^2)$. If the points of P_K are well distributed between P^- and P^+, the intercept $K \cap L$ is crossed by $\Theta(\epsilon^2 n^2)$ edges of $\binom{P_K}{2}$. Notice that the intercept $K \cap L$ can be crossed by many edges outside $\binom{P_K}{2}$. Right: Our decomposition of $\mathbb{R}^2$ uses cells of the arrangement of certain lines which are sampled from among the lines spanned by P. The depicted set K is narrow – its zone is also the zone of the principal edge pq.

The above argument yields a recurrence of the form $f_2(\epsilon) \leq 2f_2(4\epsilon/3) + 16/\epsilon^2$ which bottoms out when ϵ surpasses 1 (in which case we use the trivial bound $f(\epsilon) \leq 1$ for all $\epsilon \geq 1$).

Notice that the above approach immediately yields a net of size $o(1/\epsilon^2)$ for sets K that fall into one of the following favourable categories:

1. The interval $K \cap L$ is crossed by more than $\Theta(\epsilon^2 n^2)$ edges of $\binom{P}{2}$, with either one or both of their endpoints lying outside K.

   For example, we need only $1/\delta = o\left(1/\epsilon^2\right)$ points to pierce such sets K whose cross-sections $K \cap L$ contain at least $\delta n^2 = \omega\left(\epsilon^2 n^2\right)$ intersection points of L with the edges of $\binom{P}{2}$.

2. At least a fixed fraction of the $\Omega\left(\epsilon^2 n^2\right)$ edges spanned by P_K belong to a relatively sparse subset $\Pi \subset \binom{P}{2}$ of cardinality $m = o(n^2)$. This subset Π is carefully constructed in advance and does not depend on the choice of K.

   This too leads to a net of size $O\left(m/(\epsilon^2 n^2)\right) = o\left(1/\epsilon^2\right)$ provided that a large fraction of these edges of Π end up crossing L. (In other words, the endpoints of these edges must be sufficiently spread between the halfplanes of $\mathbb{R}^2 \setminus L$.)

Decomposing $\mathbb{R}^2$. To force at least one of the above favourable scenarios, we devise a randomized decomposition of $\mathbb{R}^2$ and P. Rather than using a single line to split $\mathbb{R}^2$ into halfplanes, we use a subset $\mathcal{R}$ of $r = o(1/\epsilon)$ lines that are chosen at random from among the lines that support the edges of $\binom{P}{2}$, and consider their entire *arrangement* $\mathcal{A}(\mathcal{R})$ – the decomposition of $\mathbb{R}^2 \setminus \bigcup \mathcal{R}$ into open 2-dimensional faces. (See Section 2.3 for the precise definition of an arrangement, and its essential properties.) We use the $\binom{r}{2} = o\left(1/\epsilon^2\right)$ vertices of $\mathcal{A}(\mathcal{R})$ to construct a small-size point set Q with the following property: Every convex set K that is *not* pierced by Q must demonstrate a "line-like" behaviour with respect to $\mathcal{A}(\mathcal{R})$ – its zone (namely, the 2-faces intersected by K) must be contained, to a large extent, in the zone of a single edge $pq \in \binom{P_K}{2}$. See Figure 1 (right). In what follows, we will refer to such convex sets as *narrow*. Though such a "proxy" edge pq can be selected from $\binom{P_K}{2}$ in (almost) $\Theta(\epsilon^2 n^2)$ ways, we will assign a unique "proxy" edge $pq \in \binom{P_K}{2}$ to every narrow convex set K, and refer to that edge pq as the *principal edge* of K.

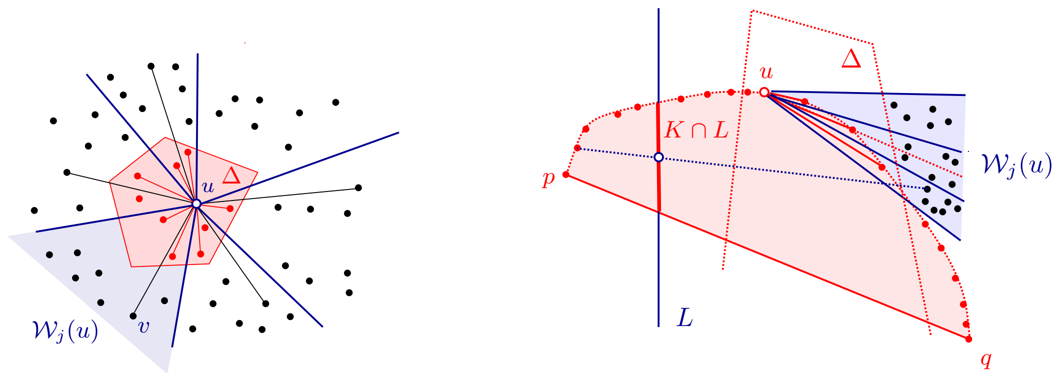

Figure 2: Left: We partition the plane into $z = O(1/\epsilon)$ sectors $\mathcal{W}_j(u)$, each containing roughly ϵn outgoing edges uv, and an average amount of $O(\epsilon n/r^2)$ outgoing short edges. Right: The point u with the outgoing short edges that are "parallel" to the principal edge pq, and whose supporting lines are roughly tangent to K. In Case 1, the $\Omega(\epsilon n/r)$ outgoing short edges of u within $\Delta \cap K$ occupy multiple sectors $\mathcal{W}_j(u)$ which are almost tangent to K. This yields $\omega(\epsilon^2 n^2)$ segments that cross the intercept $K \cap L$.

Representing narrow convex sets by edges. The fundamental difficulty of representing and manipulating convex sets (as opposed to lines, segments, simplices, and other simply-shaped geometric objects) is that they can cut the underlying point set P into exponentially many subsets $P \cap K$, so the standard divide-and-conquer schemes [22] hardly apply in this setting. Fortunately, every narrow convex set K can be largely described by its principal edge $pq \in \binom{P_K}{2}$. (For example, K cannot include points outside the respective zone of pq.)

From narrowness to expansion. The main geometric phenomenon behind our choice of the sparse (i.e., non-dense) subset $\Pi \subset \binom{P}{2}$ is that the "expected" rate of expansion of P_K within the arrangement $\mathcal{A}(\mathcal{R})$ from a point $u \in P_K$, for a narrow convex set K, is generally lower than that of the entire set P from that same point.[8]

To illustrate this behaviour, assume first that the points of P are evenly distributed among the cells of $\mathcal{A}(\mathcal{R})$, so each cell contains roughly n/r^2 points. We say that an edge $uv \in \binom{P}{2}$ is *short* if both of its endpoints lie in the same cell of $\mathcal{A}(\mathcal{R})$.

For each point u of P we partition the plane into $z = \Theta\left(\frac{1}{\epsilon}\right)$ sectors $\mathcal{W}_1(u), \mathcal{W}_2(u), \ldots, \mathcal{W}_z(u)$ so that each sector encompasses $\Theta(\epsilon n)$ outgoing edges $uv \in \binom{P}{2}$; see Figure 2 (left). To pierce a narrow convex set K whose zone in $\mathcal{A}(\mathcal{R})$ is traced by an edge $pq \in \binom{P_K}{2}$, we combine the following key observations:

i. For an *average point* u in P_K, its cell Δ contains at least $\epsilon n/(r+1)$ points of P_K, which are connected to u by short edges (because K crosses at most $r+1$ cells of $\mathcal{A}(\mathcal{R})$).

ii. For an *average edge* $uv \in \binom{P}{2}$, the respective sector $\mathcal{W}_j(u)$ contains only $O\left(\epsilon n/r^2\right)$ short edges.

We further guarantee that the points of P_K are in a sufficiently convex position, and are substantially distributed in the zone of K: The former property is enforced by using a strong $\hat{\epsilon}$-net [30],

[8] To this end, we define the pseudo-distance between a pair of points $u, v \in \mathbb{R}^2$ as the number of lines in $\mathcal{R}$ that are crossed by the open segment uv; see [18, Section 2.8] and [17]. For a finite set $A \subset \mathbb{R}^2$, and a point $u \in A$, we examine the "expected" order of magnitude of the volume $|A \cap D(u, \theta)|$ of the disc $D(u, \theta)$ as a function of $\theta \geq 0$. Clearly, this informal notion is related to the more standard concepts of doubling dimension [21] and graph expansion [31].

with $\hat{\epsilon} = \Theta(\epsilon/r)$, to eliminate the forbidden convex sets K, whereas the latter condition is enforced using a suitably amplified version of the prior line-splitting argument. Thus, for $\Omega(\epsilon n)$ choices of $u \in P_K$, we can assume that both endpoints of the principal edge $pq \in \binom{P}{2}$ of K lie outside the cell Δ of u, and at least half of the $\Omega(\epsilon n/r)$ points $v \in P_K \setminus \{p\}$ within Δ (which exist by property (i)) lie to the same side of pq as u. By the near convex position of P_K, most lines spanned by such short edges uv within Δ are roughly tangent to the convex hull of P_K; see Figure 2 (right). (In particular, the four points p, u, v, q form a convex quadrilateral.)

Assume with no loss of generality that at least half of the above short edges uv are parallel to pq, in the sense that the four points p, u, v, q appear in this order along their convex hull. Since an average sector $\mathcal{W}_j(u)$ contains only $O(\epsilon n/r)$ such edges, we interpolate between the following scenarios.

Case 1. The wedge spanned by the above $\Omega(\epsilon n/r)$ short edges $uv \in \binom{P_K}{2}$ (along with uq) occupies r "average" sectors $\mathcal{W}_j(u), \mathcal{W}_{j+1}(u), \ldots, \mathcal{W}_{j+r}(u)$, which are almost tangent to K. We show that the points of P within $\mathcal{W}_j(u) \cup \mathcal{W}_{j+1}(u) \cup \ldots \cup \mathcal{W}_{j+r}(u)$ yield $r\epsilon^2 n^2$ adjacent edges that cross the intercept $K \cap L$ of K with the "middle" vertical line L that we use to split the points of P. (Again, see Figure 2 (right).) Hence, the intersection of $K \cap L$ is relatively "thick", so we can pierce such sets using $O\left(1/(r\epsilon^2)\right)$ points.

Case 2. The previous scenario does not occur. Using the near-convexity of P_K, we find $\Omega(\epsilon n)$ outgoing edges of u within $\binom{P_K}{2}$ that are parallel to pq in the above sense and occupy a constant number of *rich* sectors $\mathcal{W}_j(u)$ with at least $\Omega(\epsilon n/r)$ short edges.

Property (ii) implies that there exist $O\left(1/(r\epsilon)\right)$ rich sectors $\mathcal{W}_j(u)$, which encompass a total of $O(n/r)$ edges that emanate from u. To pierce such convex sets K that fall into Case 2, we define our sparse set $\Pi \subset \binom{P}{2}$ as the set of edges uv which lie in rich sectors $\mathcal{W}_j(u)$, $\mathcal{W}_{j'}(v)$ (for at least one of the respective endpoints u or v). Note that the construction does not depend on the convex set K. It is easy to check that P_K spans at least $\Omega(\epsilon^2 n^2)$ such edges within Π, and sufficiently many of these edges must cross L. Hence, K falls into the second favourable case.

The vertical decomposition. Since the actual distribution of P in $\mathcal{A}(\mathcal{R})$ is not necessarily uniform, we subdivide the cells of $\mathcal{A}(\mathcal{R})$ into a total of $O(r^2)$ more homogeneous trapezoidal cells, so that each cell contains at most n/r^2 points of P.[9] To adapt the preceding expansion argument to the faces of the resulting decomposition Σ, we extend the notion of narrowness to Σ and guarantee that every narrow convex set K crosses only a small fraction of the faces in Σ. More specifically, the (strong) Epsilon Net Theorem implies that any trapezoidal cell τ is crossed by $O(n^2 \log r/r)$ of the lines that support the edges of $\binom{P}{2}$,[10] so an average edge of $\binom{P}{2}$ crosses only $O(r \log r)$ trapezoidal cells of Σ. As a result, a "typical" narrow convex set K (whose zone in Σ can "read off" from its principal edge pq within $\binom{P_K}{2}$) crosses relatively few faces of Σ; in other words, K has a *low crossing number* with respect to Σ.

The "exceptional" convex sets K, which cross too many faces of Σ, are dispatched separately using that, for $\Theta(\epsilon^2 n^2)$ of their edges, their supporting lines cross too many cells Σ and, thereby, belong to another sparse subset of $\binom{P}{2}$.

[9] A similar decomposition was used, e.g., by Clarkson *et al.* [22] to tackle the closely related problem of bounding generalized point-line incidences (e.g., incidences between points and unit circles, or incidences between lines and certain cells of their arrangement); the relation between the two problems is briefly discussed in the concluding Section 4.

[10] In other words, Σ is a $\Theta(r/\log r)$-cutting [20] of $\mathbb{R}^2$ with respect to these lines.

Discussion. Our trapezoidal decomposition Σ of $\mathbb{R}^2$ overly resembles the first step of the proof of the Simplicial Partition Theorem of Matoušek [37] in dimension $d = 2$, which provides $s = O(r^2)$ triangles $\Delta_1, \ldots, \Delta_s$ so that each triangle Δ_i contains $\Theta(n/s)$ points, and any line in $\mathbb{R}^2$ crosses $O(\sqrt{s}) = O(r)$ of these triangles.

It is instructive to compare our approach to the partition-based technique of Matoušek and Wagner [38], which directly uses the above theorem in $\mathbb{R}^d$ to re-establish the near-$1/\epsilon^d$ bounds of Alon *et al.* [2] (in $\mathbb{R}^2$) and Chazelle *et al.* [19] (in any dimension $d \geq 2$), via a simple recursion on the point set and the parameter ϵ.

Notice that the triangles Δ_i in the Simplicial Partition Theorem, for $1 \leq i \leq s$, cannot be related to particular cells of any single arrangement of lines. To enforce a low crossing number among the convex sets K with respect to the partition $\{\Delta_1, \ldots, \Delta_s\}$, Matoušek and Wagner pick a point p_i in each triangle Δ_i and pierce the "exceptional" sets by an auxiliary net of $O\left(s^{O(d^2)}\right)$ centerpoints which are obtained via Rado's Centerpoint Theorem [36] for all the possible subsets of $\{p_i \mid 1 \leq i \leq s\}$. Hence, the cardinality of their net heavily depends on the size s of the partition. Unfortunately, their simplicial partition does not quite suit our analysis, which needs a relatively large number of cells to achieve a substantial improvement over the $O\left(1/\epsilon^2\right)$ bound.

2.2 The recursive framework.

We refine the notation of Section 1 and lay down the formal framework in which our construction and its analysis are cast.

Definition. For a finite point set P in $\mathbb{R}^2$ and $\epsilon > 0$, let $\mathcal{K}(P, \epsilon)$ denote the family of all the ϵ-heavy convex sets with respect to P. We then say that $Q \subset \mathbb{R}^2$ is *a weak ϵ-net* for a family of convex sets $\mathcal{G}$ in $\mathbb{R}^2$ if it pierces every set in $\mathcal{G} \cap \mathcal{K}(P, \epsilon)$.

If the parameter ϵ is fixed, we can assume that each set in $\mathcal{K}$ is ϵ-heavy, so Q is simply a point transversal to $\mathcal{K}$. Note also that every weak ϵ-net with respect to P is, in particular, a weak ϵ-net with respect to any subfamily $\mathcal{K}$ of convex sets in $\mathbb{R}^2$.

Notice that the previous constructions [6, 19, 38] employed recurrence schemes in which every problem instance (P, ϵ) was defined over a finite point set P, and sought to pierce each ϵ-heavy convex set $K \in \mathcal{K}(P, \epsilon)$ using the smallest possible number of points. This goal was achieved in a divide-and-conquer fashion, by tackling a number of simpler sub-instances (P', ϵ') with a smaller point set $P' \subset P$ and a larger parameter $\epsilon' > \epsilon$.

To amplify our sub-quadratic bound on $f_2(\epsilon)$, we employ a somewhat more refined framework: each recursive instance is now endowed not only with the underlying point set P, but also with a certain subset of edges $\Pi \subset \binom{P}{2}$ which contains a large fraction of the edges spanned by the points of $P \cap K$. Thus, our recurrence can advance not only by increasing the parameter ϵ, but also by restricting the convex sets to "include" $\Theta\left(\epsilon^2 n^2\right)$ edges of the progressively sparser subset Π.

Definition. Let $\Pi \subset \binom{P}{2}$ be a subset of edges spanned by the underlying n-point set P. Let $\sigma > 0$. We say that a convex set K is (ϵ, σ)*-restricted* to the graph (P, Π) if $P \cap K$ contains a subset P_K of $\lceil \epsilon n \rceil$ points so that the induced subgraph $\Pi_K = \binom{P_K}{2} \cap \Pi$ contains at least $\sigma\binom{\lceil \epsilon n \rceil}{2}$ edges. (In particular, each (ϵ, σ)-restricted set K must be ϵ-heavy with respect to P.)

Notice that the choice of the set P_K may not be unique, and that K may enclose additional points of P. To simplify the presentation, in the sequel we select a unique witness set P_K for every convex set K that is (ϵ, σ)-restricted to (P, Π).

At each recursive step we construct a weak ϵ-net Q for a certain family $\mathcal{K} = \mathcal{K}(P, \Pi, \epsilon, \sigma)$ of convex sets which is determined by $\epsilon > 0$, a ground set $P \subset \mathbb{R}^2$ of n points, a set of edges $\Pi \subseteq \binom{P}{2}$, and a threshold $0 < \sigma \leq 1$. This family $\mathcal{K}$ consists of all the convex sets K that are (ϵ, σ)-restricted to (P, Π). In what follows, we refer to (P, Π) (or simply to Π) as the *restriction graph*, and to σ as the *restriction threshold* of the recursive instance.

The topmost instance of our recurrence involves $\Pi = \binom{P}{2}$ and $\sigma = 1$. Each subsequent instance $\mathcal{K}' = \mathcal{K}(P', \Pi', \epsilon', \sigma')$ involves a larger ϵ' and/or a *much sparser* restriction graph (P', Π'). Each such increase in ϵ or decrease in the density $|\Pi|/\binom{n}{2}$ is accompanied only by a comparatively mild decrease in the restriction threshold σ which, throughout the recurrence, is bounded from below by a certain positive constant.

Our weak ϵ-net construction bottoms out when either (i) the cardinality of P falls below a certain threshold $n_0(\epsilon)$ that is close to $1/\epsilon^{3/2}$, or (ii) ϵ surpasses a certain (suitably small) constant $0 < \tilde{\epsilon} < 1$, or (iii) the density $|\Pi|/\binom{n}{2}$ of the restriction graph falls below ϵ. In the first case, P comprises the desired weak ϵ-net. In the second case, we can use the $O\left((1/\tilde{\epsilon})^2\right) = O(1)$ bound of Alon *et al.* [2]. In the third case, the discussion in Section 2.1, which we formalize in Lemma 2.5, yields a simpler "near-linear" recurrence in $1/\epsilon$. [11]

Deriving a recurrence formula for $f_2(\epsilon)$. In the course of our analysis we stick with the following notation. We use $f(\epsilon, \lambda, \sigma)$ to denote the smallest number f so that for any finite point set P in $\mathbb{R}^2$, and any subset $\Pi \subseteq \binom{P}{2}$ of density $|\Pi|/\binom{n}{2} \leq \lambda$, there is a point transversal of size f to $\mathcal{K}(P, \Pi, \epsilon, \sigma)$. We set $f(\epsilon, \lambda, \sigma) = 1$ whenever $\epsilon \geq 1$. Since the underlying dimension $d = 2$ is fixed, for the sake of brevity we use $f(\epsilon)$ to denote the quantity $f_2(\epsilon) = f(\epsilon, 1, 1)$, and note that the trivial bound $f(\epsilon, \lambda, \sigma) \leq f(\epsilon)$ always holds.

In the sequel, we bound the quantity $f(\epsilon, \lambda, \sigma)$, for $\lambda > \epsilon$, by a recursive expression of the general form

$$f(\epsilon, \lambda, \sigma) \leq f\left(\epsilon, \lambda/r_0, \sigma/2\right) + O\left(\sum_{i=1}^{l} h_i^{1+\gamma'} \cdot f\left(\epsilon \cdot h_i\right) + 1/\epsilon^{3/2+\gamma'}\right), \tag{1}$$

where l is a constant that does not depend on the choice of γ in Theorem 1.1, γ' is a small constant that satisfies $0 < \gamma' < \gamma$, and the parameters r_0 and h_i, for $1 \leq i \leq l$, are very small (albeit fixed) degrees of $1/\epsilon$ that are bounded by $(1/\epsilon)^{\gamma'}$.

In particular, (1) will hold for $f(\epsilon) = f(\epsilon, 1, 1)$. As the recurrence in the density λ bottoms out for $\lambda \leq \epsilon$, applying $J = \log_{r_0}\lceil 1/\epsilon \rceil$ substitution steps to the first term in (1) while keeping ϵ fixed, and then using the bound in Lemma 2.5 (e.g., with $r = \Theta\left(1/\epsilon^{\gamma'}\right)$) when λ falls below ϵ, will result in the following bound.

$$f(\epsilon) = O\left(r \cdot f\left(\epsilon \cdot r\right) + \sum_{i=1}^{l} h_i^{1+\gamma'} \cdot f\left(\epsilon \cdot h_i\right) + 1/\epsilon^{3/2+\gamma'}\right). \tag{2}$$

Notice that in each of the intermediate substitutions, the restriction threshold σ in $f(\epsilon, \lambda, \sigma)$ remains bounded from below by $2^{-J} = \Theta(1)$.

As was previously mentioned, the recurrence in $0 < \epsilon < 1$ bottoms out when it bypasses a certain constant threshold $0 < \tilde{\epsilon} < 1$. With a sufficiently small (albeit, constant) choice of $\tilde{\epsilon}$ which too depends on γ, the standard and fairly general induction argument (as presented, e.g., in [38, 46] and

[11] As a matter of fact, we have $f(\epsilon, \lambda, \sigma) = o\left(1/\epsilon^2\right)$ once the maximum density λ falls substantially below 1 (given that the restriction threshold σ is a constant). Hence, our recurrence over Π is used to merely amplify this gain.

[47, Section 7.3.2]) shows that recurrences of the general form of (2) solve to $f(\epsilon) = O\left(1/\epsilon^{\frac{3}{2}+\gamma}\right)$, where the constant of proportionality is super-exponential in $1/\gamma$.

2.3 Geometric essentials: Arrangements and strong ϵ-nets

Strong ϵ-nets. Let X be a (finite) set of elements and $\mathcal{F} \subset 2^X$ be a set of hyperedges spanned by X. A *strong ϵ-net* for the hypergraph $(X, \mathcal{F})$ is a subset $Y \subset X$ of elements so that $F \cap Y \neq \emptyset$ is satisfied for all hyperedges $F \in \mathcal{F}$ with $|F| \geq \epsilon n$.

Definition. Let X be a set of n elements, and $r > 0$ be an integer. An *r-sample* of X is a subset $Y \subset X$ of r elements chosen at random from X, so that each such subset $Y \in \binom{X}{r}$ is selected with uniform probability $1/\binom{n}{r}$.

The Epsilon-Net Theorem of Haussler and Welzl [30] states that any such hypergraph $(X, \mathcal{F})$, that is drawn from a so called range space of a bounded VC-dimension $D > 0$, admits a strong ϵ-net Y of cardinality $r = O\left(\frac{D}{\epsilon} \log \frac{D}{\epsilon}\right)$. Moreover, such a net Y can obtained, with probability at least $1/2$, by choosing an r-sample of X.

In particular, this implies the following result.

Theorem 2.1. *Let P be a finite set of points in $\mathbb{R}^2$, then one can pierce all the ϵ-heavy triangles with respect to P using a net $Q^{\triangle}(P, \epsilon)$ of cardinality $O\left(\frac{1}{\epsilon} \log \frac{1}{\epsilon}\right)$.*

Cuttings of families of lines. An important corollary of the Epsilon-Net Theorem is the existence of *ϑ-cuttings* of finite families $\mathcal{L}$ of lines in $\mathbb{R}^2$ – decompositions of $\mathbb{R}^2$ into $O^*(\vartheta^{-2})$ interior-disjoint (and possibly unbounded) convex polygonal regions, which are called *cells*, so that any cell is delimited by $O(1)$ edges and its interior is crossed by at most $\vartheta|\mathcal{L}|$ lines. A two-stage such construction of worst-case size $O\left(\vartheta^{-2}\right)$ was obtained by Chazelle and Friedman [20]; it uses the so called Exponential Decay Lemma to control the number of cells that arise in the secondary subdivision.

In what follows, we lay out a simpler cutting which is based on trapezoidal subdivisions of random line arrangements, and loosely corresponds to the first stage in the optimal cutting of Chazelle and Friedman.

Arrangements of lines in $\mathbb{R}^2$. Our divide-and-conquer approach uses cells in the arrangement of lines that are sampled at random from among the lines spanned by the edges of our restriction graph (P, Π).

To simplify the exposition, we can assume that the points of P are in a general position. In particular, no three of them are collinear, and no two of them span a vertical line.[12]

Definition. Any finite family $\mathcal{L}$ of m lines in $\mathbb{R}^2$ induces the *arrangement* $\mathcal{A}(\mathcal{L})$ – the partition of $\mathbb{R}^2$ into 2-dimensional *cells*, or *2-faces* – maximal connected regions of $\mathbb{R}^2 \setminus (\bigcup \mathcal{L})$. Each of these cells is a convex polygon whose boundary is composed of *edges* – portions of the lines of $\mathcal{L}$, which connect *vertices* – crossings among the lines of $\mathcal{L}$. The *complexity* of a cell is the total number of edges and vertices that lie on its boundary.

[12] To construct a weak ϵ-net for a degenerate point set P, we perform a routine symbolic perturbation of P into a general position. A weak ϵ-net with respect to the perturbed set would immediately yield such a net with respect to the original set.

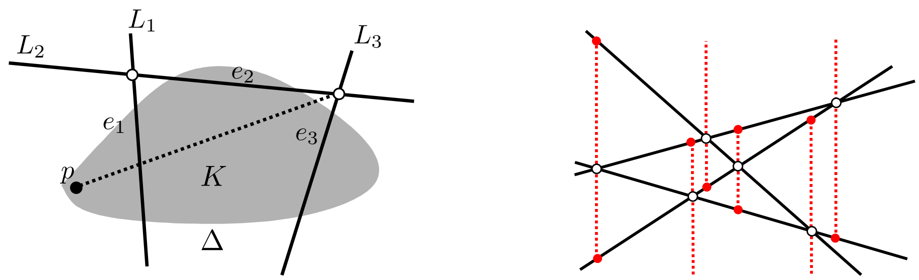

Figure 3: Left: Lemma 2.2 – the set K meets the three boundary edges e_1, e_2, e_3 of the cell $\Delta = L_1^- \cap L_2^- \cap L_3^-$. The point $p \in K$ lies in $L_1^+ \cap L_2^- \cap L_3^-$. The segment between p and $L_2 \cap L_3$ crosses $K \cap L$. Right: The trapezoidal decomposition $\Sigma(\mathcal{L})$.

Lemma 2.2. *Let L_1, L_2 and L_3 be three lines in $\mathbb{R}^2$, and $\Delta \subset \mathbb{R}^2 \setminus (L_1 \cup L_2 \cup L_3)$ be a cell in their arrangement. For each $1 \leq i \leq 3$, let L_i^- and L_i^+ be the two halfplanes of $\mathbb{R}^2 \setminus L_i$ so that $\Delta \subset L_i^-$ (see Figure 3 (left)). Suppose that each line L_i contains a bounary edge e_i of Δ so that the three edges e_1, e_2, and e_3, appear in this clockwise order along the boundary of Δ. Then for any convex set K that meets all the three sides e_1, e_2, e_3 of Δ, and any point $p \in K \cap L_1^+ \cap L_2^- \cap L_3^-$, the segment between $L_2 \cap L_3$ and p must cross L_1 within the interval $K \cap L_1$.*[13]

The trapezoidal decomposition. We further subdivide each cell Δ of the above arrangement $\mathcal{A}(\mathcal{L})$ by raising a vertical wall from every boundary vertex of Δ that is not x-extremal (i.e., if the vertical line through the vertex enters the interior of Δ); see Figure 3 (right). As is easy to check, the resulting decomposition $\Sigma(\mathcal{L})$ is composed of $O\left(m^2\right)$ open trapezoidal cells. The boundary of each cell μ in $\Sigma(\mathcal{L})$ consists of at most 4 edges[14], including at most 2 vertical edges, and the at most 2 other edges that are contained in non-vertical lines of $\mathcal{L}$.

Theorem 2.3. *Let $\mathcal{L}$ be a family of m lines in $\mathbb{R}^2$, and $0 < r \leq m$ integer. Then, with probability at least $1/2$, an r-sample $\mathcal{R} \in \binom{\mathcal{L}}{r}$ of $\mathcal{L}$ crosses every segment in $\mathbb{R}^2$ that is intersected by at least $C(m/r)\log r$ lines of $\mathcal{L}$. Here $C > 0$ is a sufficiently large constant that does not depend on m or r.*

The proof of Theorem 2.3 can be found, e.g., in [18]. It is established by applying the Epsilon Net Theorem to the range space in which every vertex set is a finite family $\mathcal{L}$ of lines in $\mathbb{R}^2$, and each hyperedge consists of all the lines in $\mathcal{L}$ that are crossed by some segment in $\mathbb{R}^2$.

An easy consequence of Theorem 2.3 is that, with probability at least $1/2$, every (open) trapezoidal cell of the induced vertical decomposition $\Sigma(\mathcal{R})$ is crossed by at most $4C(m/r)\log r$ lines of $\mathcal{L}$. In other words, it serves as a $\left(\frac{4C\log r}{r}\right)$-cutting of $\mathcal{L}$. Despite a marginally sub-optimal bound on the number of cells (in the terms of $\vartheta := 4C\log r/r$), the simplicity of $\Sigma(\mathcal{R})$ will prove beneficial for our ad-hoc argument in Section 3.3.

The zone. Let Σ be a family of open cells in $\mathbb{R}^2$ (e.g., the above arrangement $\mathcal{A}(\mathcal{L})$ or its refinement $\Sigma(\mathcal{L})$). The *zone* of a convex set $K \subset \mathbb{R}^2$ in Σ is the subset of all the cells in Σ that intersect K.

[13] In the sequel, we apply the lemma only in the special case where L_1 and L_3 are vertical lines.

[14] Some of the trapezoidal cells can be triangles, or unbounded.

The *crossing number* of a convex set K with respect to Σ is the cardinality of its zone within Σ, that is, the number of the cells in Σ that are intersected by K.

Definition. For every pair $p, q \in \mathbb{R}^2$ let $L_{p,q}$ denote the line through p and q. We say that the line $L_{p,q}$ is *spanned* by the segment pq, and that pq is *supported* by $L_{p,q}$. Given an n-point set P with an edge set $\Pi \subset \binom{P}{2}$, let

$$\mathcal{L}(\Pi) := \{L_{p,q} \mid \{p,q\} \in \Pi\}$$

be the set of all the lines spanned by the edges of Π. If the underlying restriction graph (P, Π) is clear from the context, we resort to a simpler notation $\mathcal{L} := \mathcal{L}(\Pi)$.

Decomposing $\mathbb{R}^2$ into vertical slabs. For any n-point set P, and any integer $r > 0$, we fix a collection $\mathcal{Y}(P, r)$ of r vertical lines so that every vertical slab of the arrangement $\mathcal{A}(\mathcal{Y}(P, r))$ contains between $\lfloor n/(r+1) \rfloor$ to $\lceil n/(r+1) \rceil$ points of P, and no line of $\mathcal{Y}(P, r)$ passes through a point of P. (The two extremal slabs of $\mathcal{A}(\mathcal{Y}(P, r))$ are halfplanes, and each of them is delimited by a single line of $\mathcal{Y}(P, r)$.)

In what follows, we use $\Lambda(P, r)$ to denote the above slab decomposition $\mathcal{A}(\mathcal{Y}(P, r))$.

We say that a segment $pq \subset \mathbb{R}^2$ crosses a slab $\tau \in \Lambda(P, r)$ *transversally* if pq intersects the interior of τ, and none of its endpoints p, q lies in Δ; see Figure 4.

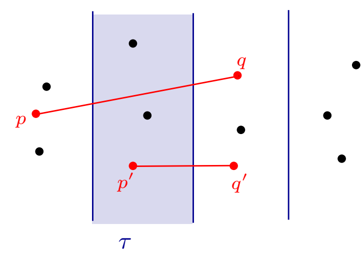

Figure 4: The vertical lines of $\mathcal{Y}(P, r)$ determine a decomposition $\Lambda(P, r)$ of $\mathbb{R}^2$ into $r+1$ vertical slabs. (In the depicted scenario, we have $r = 3$.) The segment pq crosses the slab $\tau \in \Lambda(P, r)$ transversally, while $p'q'$ does not.

We say that a convex set K is ϵ'*-crowded* in $\Lambda(P, r)$ if there a slab in $\Lambda(P, r)$ that contains at least $\epsilon' n$ points of $P \cap K$; otherwise, we say that K is ϵ'*-spread* in $\Lambda(P, r)$.[15]

The following main property of the decompositions $\Lambda(P, r)$ is used throughout our proof of Theorem 1.1.

Lemma 2.4. *Let P be an underlying set of n points in $\mathbb{R}^2$ and $r > 0$ be an integer. For each $\epsilon' \geq 0$ there is a set $Q(P, r, \epsilon')$ of $O\left(r \cdot f\left(\epsilon' \cdot r\right)\right)$ points that pierce every convex set K that is ϵ'-crowded in $\Lambda(P, r)$.*[16]

Notice that the recursive term in Lemma 2.4 is essentially linear in ϵ for ϵ' close enough to ϵ; see, e.g., [19, Section 3] for a similar recurrence.

[15] We emphasize that the ϵ'-crowdedness of a convex set K depends not only on the slabs of $\Lambda(P, r)$ but also its underlying point set P.

[16] To simplify the presentation, we routinely omit the constant factors within the recursive terms of the form $f(\epsilon \cdot hr)$ as long as these constants are much larger than $1/h$. A suitably small choice of the constant $\tilde{\epsilon} > 0$ (and, thereby, $\epsilon < \tilde{\epsilon}$) guarantees that ϵ indeed increases with each invocation of the recurrence.

Proof of Lemma 2.4. Assume with no loss of generality that $r < 2n$, for otherwise our net consists of P. Recall that each slab $\tau \in \Lambda(P,r)$ cuts out a subset $P_\tau := P \cap \tau$ of cardinality $n_\tau := |P_\tau| \leq \lceil n/(r+1) \rceil = \Theta(n/r)$.

The crucial observation is that each ϵ'-crowded convex set K must belong to the family $\mathcal{K}(P_\tau, \epsilon' n/n_\tau)$ for some slab τ in $\Lambda(P,r)$. (In particular, we can further assume that $\epsilon' = O(1/r)$.) For each slab $\tau \in \Lambda(P,r)$ we recursively construct the net Q_τ for the above instance $\mathcal{K}(P_\tau, \epsilon' n/n_\tau)$. Using the definition of the function $f(\cdot)$, and that $n_\tau = \Theta(n/r)$, it is easy to check that the total cardinality of the union $Q(P,r,\epsilon') := \bigcup_{\tau \in \Lambda(P,r)} Q_\tau$ is indeed $O(r \cdot f(r \cdot \epsilon'))$. □

The following lemma implies that the recursive instance $\mathcal{K} = (P, \Pi, \epsilon, \sigma)$ admits a net of size $o(1/\epsilon^2)$ given that the underlying restriction graph (P, Π) is not dense (and that the restriction threshold σ is sufficiently close 1).

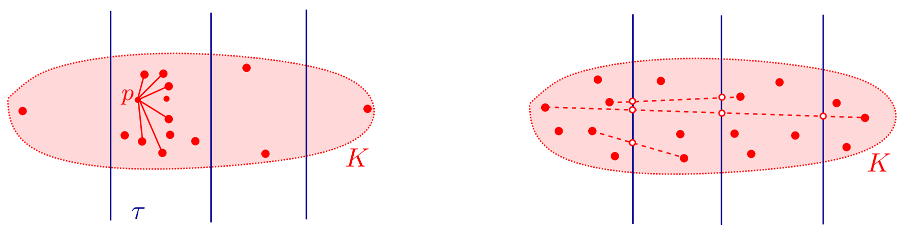

Figure 5: Proof of Lemma 2.5. Left: Most edges of Π_K do not cross any line of $\mathcal{Y}(P,r)$. Hence, there is a slab that contains $\Omega(\sigma\epsilon n)$ points of P_K. Right: At least half of the edges of Π_K cross one or more lines of $\mathcal{Y}(P,r)$, so K must be pierced by one of the nets Q_L.

Lemma 2.5. *Let $r \geq 1$ be an integer. Then any family $\mathcal{K} \subset \mathcal{K}(P, \Pi, \epsilon, \sigma)$ admits a point transversal of size*

$$O\left(r \cdot f\left(\epsilon \cdot \sigma \cdot r\right) + \frac{r^2|\Pi|}{\sigma\epsilon^2 n^2}\right).$$

Proof. Assume with no loss of generality that $|P| \geq 2r$, for otherwise the claim follows trivially. We consider the slab decomposition $\Lambda(P,r)$ and apply Lemma 2.4 with $\epsilon' = \sigma\epsilon/4$ to obtain a net $Q(P,r,\epsilon')$ of size $O\left(r \cdot f\left(\epsilon \cdot \sigma \cdot r\right)\right)$ that pierces every set $K \in \mathcal{K}$ that is ϵ'-crowded in $\Lambda(P,r)$.

In addition, for each vertical line $L \in \mathcal{Y}(P,r)$ we construct an auxiliary net Q_L by choosing every $\lceil \sigma \binom{\lceil \epsilon n \rceil}{2} / (2r) \rceil$-th crossing point of L with the edges of Π. Notice that

$$\sum_{L \in \mathcal{Y}(P,r)} |Q_L| = O\left(\frac{r^2|\Pi|}{\sigma\epsilon^2 n^2}\right)$$

It suffices to check that every convex set $K \in \mathcal{K}$ is stabbed by at least one of the above nets. To this end, we distinguish between two cases.

1. If at least half of the segments of $\Pi_K = \binom{P_K}{2} \cap \Pi$ do not cross any line of $\mathcal{Y}(P,r)$, we find a point $p \in P_K$ so that at least $2\sigma\binom{\lceil \epsilon n \rceil}{2} / \lceil \epsilon n \rceil \geq \sigma\epsilon n/4$ of its neighbors in the graph (P_K, Π_K) lie in the same slab $\tau \in \Lambda(P,r)$ that contains p. Hence, K is ϵ'-crowded in $\Lambda(P,r)$ and, therefore, pierced by a point of Q'. See Figure 5 (left).

2. At least half of the segments of Π_K cross a line of $\mathcal{Y}(P,r)$. Since there are at least $(\sigma/2)\binom{\lceil \epsilon n \rceil}{2}$ intersection points between the edges of Π_K and the lines of $\mathcal{Y}(P,r)$, there must be a line

$L \in \mathcal{Y}(P, r)$ which contains at least $\sigma\binom{\lceil \epsilon n \rceil}{2}/(2r)$ of these intersections. Hence, K is hit by the corresponding net Q_L. See Figure 5 (right).

□

As mentioned in Section 2.2, throughout our analysis σ remains bounded from below by a certain positive constant, and we apply Lemma 2.5 with r that is a very small (albeit, fixed) constant power $1/\epsilon$. (In particular, r is much larger than $1/\sigma$.) Notice that this yields the following bound

$$f(\epsilon, \lambda, \sigma) = O\left(r \cdot f(\epsilon \cdot \sigma \cdot r) + \frac{r^2\lambda}{\sigma\epsilon^2}\right) \tag{3}$$

in which the recursive term on the right side is essentially linear in $1/\epsilon$, and the constants of proportionality that are hidden by the $O(\cdot)$-notation do not depend on ϵ, σ, and λ. Moreover, the non-recursive term is $o(1/\epsilon^2)$ provided that the density λ is substantially smaller than 1, and it is close to $1/\epsilon$ if $\lambda \leq \epsilon$. A standard inductive approach to solving recurrences of this kind is presented, e.g., in [28] and [47, Section 7.3.2].

3 Proof of Theorem 1.1

To establish Theorem 1.1, we fix $\gamma > 0$ and show that $f(\epsilon) = O\left(1/\epsilon^{3/2+\gamma}\right)$. To this end, we first derive a recurrence formula of the general form (1) for the quantity $f(\epsilon, \lambda, \sigma)$, where $\lambda > \epsilon$. As mentioned in Section 2.2, the final recurrence for $f(\epsilon)$, of the form (2), will follow by iterating this recurrence for $f(\epsilon, \lambda, \sigma)$ with $\sigma = \Theta(1)$, and then plugging in the bound of Lemma 2.5 for $\lambda \leq \epsilon$. [17]

To obtain the desired recurrence for $f(\epsilon, \lambda, \sigma)$ for $\lambda > \epsilon$, we bound the piercing number of the family $\mathcal{K} := \mathcal{K}(P, \Pi, \epsilon, \sigma)$ for an arbitrary choice of the finite point set $P \subset \mathbb{R}^2$ in general position, the parameters $0 \leq \epsilon, \sigma \leq 1$, and the edge set $\Pi \subseteq \binom{P}{2}$ that satisfies

$$|\Pi| / \binom{|P|}{2} \leq \lambda. \tag{4}$$

In the course of our construction, we introduce three auxiliary parameters r_0, s_0, and r_1. The first two parameters are set to be very small degrees of $1/\epsilon$ that depend on γ. To this end, we set $\eta := \gamma/100$ and $r_0 := \left\lceil 1/\epsilon^{\eta^2} \right\rceil$, and choose $s_0 = \Theta\left(1/\epsilon^{\eta}\right)$ with the property that $s_0 + 1$ is the smallest multiple of $r_0 + 1$ that is larger than $r_0^{1/\eta}$. In addition, we will set $r_1 := \left\lceil \sqrt{1/\epsilon} \right\rceil$.

In what follows, we can assume that ϵ is bounded from above by a sufficiently small absolute constant $\tilde{\epsilon} > 0$ which, in particular, guarantees that the following inequality holds[18]

$$10^5 \log \frac{1}{\epsilon} < \frac{1}{\epsilon^{\eta}}. \tag{5}$$

Otherwise, if $\epsilon \geq \tilde{\epsilon}$, the previous $O\left(1/\tilde{\epsilon}^2\right) = O(1)$ bound applies [2]. We can also assume that the number of points $n = |P|$ is larger than a certain threshold $n_0(\epsilon)$, namely,

$$n > n_0(\epsilon) := \left\lceil \frac{3200 r_0 r_1 \log r_1}{\sigma\epsilon} \right\rceil. \tag{6}$$

[17]Though we have $\sigma = \Theta(1)$ in all the subsequent applications of our recurrence for $f(\epsilon, \lambda, \sigma)$, starting with $f(\epsilon) = f(\epsilon, 1, 1)$, the dependence on the restriction threshold σ will be spelled out throughout our analysis.

[18]In the sequel $\log x$ denotes the binary logarithm $\log_2 x$.

Otherwise, if $|P| \leq n_0(\epsilon)$, our transversal consists of $|P| \leq n_0(\epsilon) = o\left(1/\epsilon^{3/2+\gamma}\right)$ points.

For $\epsilon < \min\{\lambda, \tilde{\epsilon}\}$, which satisfy (5), and the sets P and $\Pi \subseteq \binom{P}{2}$ that satisfy (4) and $|P| > n_0(\epsilon)$, the piercing number of $\mathcal{K}$ will be bounded in the terms of the quantities $f(\epsilon, \lambda/r_0, \sigma/2)$ and $f(\epsilon')$, for $\epsilon' > \epsilon$. To this end, we gradually construct a net Q which pierces every ϵ-heavy set $K \in \mathcal{K}$. Our construction begins with an empty net $Q = \emptyset$ and proceeds through several stages. At each stage we add a small number of points to the net Q and immediately eliminate the already pierced convex sets from the family $\mathcal{K}$. The surviving sets $K \in \mathcal{K}$, which have yet not been pierced by Q, satisfy additional restrictions which facilitate their treatment at the subsequent stages.

Our main decomposition $\Sigma = \Sigma(r_1)$ of $\mathbb{R}^2$ in Section 3.2 is based on cells in the arrangement of an r_1-sample $\mathcal{R}_1$ of $\mathcal{L} = \mathcal{L}(\Pi)$, for a fairly large value $r_1 = \left\lceil \sqrt{1/\epsilon} \right\rceil$. Informally, the lines of $\mathcal{R}_1$ are sampled from $\mathcal{L}$ so as to control the crossing number (i.e., size of the respective zone in $\Sigma(r_1)$) of an average edge pq of Π. This bound readily extends to the narrow convex sets K whose zones are traced by such edges pq. Recall that our main argument (which was sketched in Section 2.1) requires that the points P_K of each set $K \in \mathcal{K}$ are in a "sufficiently convex" position, and are substantially spread within the zone of K in $\mathcal{A}(\mathcal{R}_1)$. To this end, we employ an auxiliary slab decomposition $\Lambda(P, r_0)$ of Lemma 2.4 in combination with The Epsilon Net Theorem 2.1.

The roadmap. The rest of this section is organized as follows.

In Section 3.1 we construct an auxiliary slab decomposition $\Lambda(P, r_0)$, and use Lemma 2.4 to guarantee that the points of our convex sets K are sufficiently spread among the slabs of $\Lambda(P, r_0)$. This is achieved at expense of adding to Q a small-size auxiliary net Q_0 which is provided by Lemma 2.4.

In Section 3.2 we use the larger sample $\mathcal{R}_1$ of r_1 lines from $\mathcal{L} = \mathcal{L}(\Pi)$ to define the finer main decomposition $\Sigma(r_1)$ of $\mathbb{R}^2$. As mentioned in Section 2.1, $\Sigma(r_1)$ is obtained by vertically subdividing the cells of $\mathcal{A}(\mathcal{R}_1)$ into trapezoidal sub-cells. By the properties of $\Sigma(r_1)$ as a $\Theta(\log r_1/r_1)$-cutting for $\mathcal{L}$ [20], an average line of $\mathcal{L}$ crosses only $O(r_1 \log r_1)$ cells of $\Sigma(r_1)$. We further "normalize" Π by omitting a relatively small fraction of its edges whose supporting lines in $\mathcal{L}$ cross too many of the cells of $\Sigma(r_1)$. We then remove from $\mathcal{K}$ every convex set that is not $(\epsilon, \sigma/2)$-restricted to the surviving graph (P, Π). To that end, we add to Q another auxiliary net Q_1 which is obtained by solving a simpler recursive instance $\mathcal{K}(P, \Pi', \epsilon, \sigma/2)$, with a much sparser restriction graph Π'.

In Section 3.3 make sure that every remaining set $K \in \mathcal{K}$ is narrow in the decomposition $\Sigma(r_1)$ (in the sense described in Section 2.1) and, therefore, it crosses roughly $O(r_1 \log r_1)$ of the decomposition cells.[19] The leftover convex sets, that are not sufficiently narrow in $\Sigma(r_1)$, are pierced by an auxiliary net Q_2 whose size is close to r_1/ϵ.

In Section 3.4 we use the properties of $\Sigma(r_1)$ to construct the final net Q_3 which pierces all the remaining sets $K \in \mathcal{K}$ (missed by the auxiliary nets Q_i of the previous stages $0 \leq i \leq 2$). This is achieved through a skillful combination of the two paradigms sketched in Section 2.1. Thus, the eventual net Q for our family $\mathcal{K}$ is given by the union $\bigcup_{i=0}^{3} Q_i$.

In Section 3.6 combine the bounds of the preceding Sections 3.1 – 3.4 to bound the cardinality of the complete net Q, and then derive the final recurrences for the quantities $f(\epsilon, \lambda, \sigma)$ and $f(\epsilon)$.

3.1 Stage 0: The slab decomposition $\Lambda(P, r_0)$

At this stage we construct an auxiliary, almost constant-size slab decomposition $\Lambda(P, r_0)$, and use Lemma 2.4 to guarantee for each convex set $K \in \mathcal{K}$ that the points of P_K are sufficiently spread

[19] More precisely, we "clip" every set K to a carefully chosen slab $\tau \in \Lambda(P, r_0)$, and apply a similar restriction to $\Sigma(r_1)$.

among the slabs of $\Lambda(P, r_0)$. This is achieved at the expense of adding to Q a certain auxiliary net Q_0, and immediately removing from $\mathcal{K}$ all the sets already pierced by Q_0.

To this end, we select a set $\mathcal{Y}(P, r_0)$ of vertical lines, as detailed in Section 2.3; each slab τ of the resulting arrangement $\Lambda(P, r_0)$ contains between $\lfloor n/(r_0+1) \rfloor$ and $\lceil n/(r_0+1) \rceil$ points of P. (As previously noted, the integer parameter $r_0 = \Theta\left(1/\epsilon^{\eta^2}\right)$ is set to a very small degree of $1/\epsilon$.)

The net Q_0. By Lemma 2.4, we can pierce (and subsequently remove from $\mathcal{K}$) every $(\sigma\epsilon/100)$-crowded convex set K using an auxiliary net

$$Q_0 := Q\left(P, r_0, \sigma\epsilon/100\right) \tag{7}$$

whose cardinality satisfies

$$|Q_0| = O\left(r_0 \cdot f\left(\epsilon \cdot \sigma \cdot r_0\right)\right). \tag{8}$$

In what follows, we can assume that every remaining set $K \in \mathcal{K}$ is $(\sigma\epsilon/100)$-spread in the slab decomposition $\Lambda(P, r_0)$.

3.2 Stage 1: The main decomposition of $\mathbb{R}^2$

At this stage we construct the main decomposition $\Sigma(r_1)$ of $\mathbb{R}^2$ into $O(r_1^2)$ cells. Since $\Sigma(r_1)$ is a refinement of the auxiliary slab decomposition $\Lambda(P, r_0)$, we can use the properties of $\Lambda(P, r_0)$ to show that the points of P_K are sufficiently spread in the finer decomposition $\Sigma(r_1)$. In particular, a so called middle slab $\tau \in \Lambda(P, r_0)$ can be obtained for every remaining convex set $K \in \mathcal{K}$, which will play a quintessential role in the analysis of Sections 3.3 and 3.4.

The decomposition $\Sigma(r_1)$. We sample a subset $\mathcal{R}_1$ of $r_1 = \left\lceil \sqrt{1/\epsilon} \right\rceil$ lines from $\mathcal{L} = \mathcal{L}(\Pi)$. We can assume with no loss of generality that no line of $\mathcal{Y}(P, r_0)$ passes through a vertex of $\mathcal{A}(\mathcal{R}_1)$.[20] To simplify the exposition, we add the vertical lines of $\mathcal{Y}(P, r_0)$ to $\mathcal{R}_1$, so the arrangement $\mathcal{A}(\mathcal{R}_1)$ is a refinement of $\Lambda(P, r_0)$.

We then construct the trapezoidal decomposition $\Sigma(\mathcal{R}_1)$ of $\mathcal{A}(\mathcal{R}_1)$ which was described in Section 2.3; see Figure 6. We further subdivide each cell $\hat{\mu} \in \Sigma(\mathcal{R}_1)$ (where necessary) into sub-trapezoids μ so that $|P \cap \mu| \leq n/r_1^2$; this can be achieved using $O\left(\lfloor r_1^2 |P \cap \hat{\mu}|/n \rfloor\right)$ additional vertical walls. Furthermore, we can assume that none of these walls coincides with a point of P.

A standard calculation (see, e.g., [16]) shows that the resulting finer partition $\Sigma(r_1)$ encompasses a total of $O(r_1^2)$ trapezoids. Since $\Sigma(r_1)$ is a refinement of $\Sigma(\mathcal{R}_1)$, the relative interior of each of its cells is still crossed by $O\left((m \log r_1)/r_1\right)$ lines of $\mathcal{L}$, where m denotes the cardinality of Π and $\mathcal{L} = \mathcal{L}(\Pi)$. As explained in Section 2.3, $\Sigma(\mathcal{R})$ along with its refinement $\Sigma(r_1)$ serve as simple $O\left(\frac{\log r_1}{r_1}\right)$-cuttings with respect to $\mathcal{L}$.

Refining the restriction graph Π. Since every trapezoidal cell of $\Sigma(r_1)$ is crossed by $O((m \log r_1)/r_1)$ lines of $\mathcal{L}$, the zone of an "average" line in $\mathcal{L} = \mathcal{L}(\Pi)$ consists of $O(r_1 \log r_1)$ cells of $\Sigma(r_1)$.

To eliminate the edges of Π whose supporting lines in $\mathcal{L}$ deviate "too far" from the average behaviour with respect to our decomposition $\Sigma(r_1)$, we set $t := r_0 r_1 \log r_1$ and use $\mathcal{L}_{>t} \subset \mathcal{L}$ to denote the subset of all the lines in $\mathcal{L}$ that cross more than t (open) cells of $\Sigma(r_1)$.

[20] If $m < r_1$ then we obtain the desired decomposition by choosing $\mathcal{R}_1 = \mathcal{L}$. Note that the lines of $\mathcal{R}_1$ are not necessarily in a general position: many of them can pass through the same point of P. Nevertheless, there exist at most $2r_1$ such points in P that lie on one or more lines of $\mathcal{R}_1$.

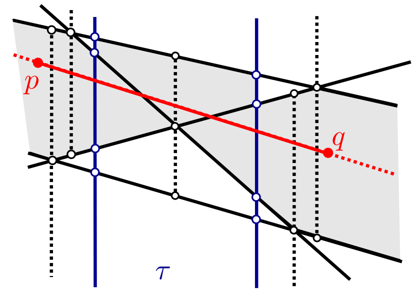

Figure 6: Our vertical decomposition $\Sigma(r_1)$ which incorporates the lines of $\mathcal{Y}(P, r_0)$. The zone of a line $L_{p,q} \in \mathcal{L}$ is shaded. Following the removal of $\Pi_{>t}$, every remaining edge $pq \in \Pi$ crosses at most $t = r_0 r_1 \log r_1$ cells of $\Sigma(r_1)$.

Proposition 3.1. *We have that*

$$|\mathcal{L}_{>t}| = O\left(\frac{m}{r_0}\right).$$

For the sake of completeness, we spell out the fairly standard proof of Proposition 3.1.

Proof. Since any trapezoidal cell μ in $\Sigma(r_1)$ is crossed by $O\left((m \log r_1)/r_1\right)$ lines of $\mathcal{L}$, the bipartite graph of pairwise intersections between the lines of $\mathcal{L}$ and the cells of $\Sigma(r_1)$ contains

$$O\left(r_1^2 \cdot \frac{m \log r_1}{r_1}\right) = O(r_1 m \log r_1)$$

edges. Since every line of $\mathcal{L}_{>t}$ contributes at least $t = r_0 r_1 \log r_1$ intersections, the claim now follows by applying the pigeonhole principle (or Markov's inequality). □

The net Q_1. Let $\Pi_{>t}$ be the set of edges that span the lines of $\mathcal{L}_{>t}$. Consider the recursive instance

$$\mathcal{K}_{>t} := \mathcal{K}(P, \Pi_{>t}, \epsilon, \sigma/2).$$

Using the bound of Proposition 3.1 on $|\Pi_{>t}| = |\mathcal{L}_{>t}|$, we can pierce the sets of $\mathcal{K}_{>t}$ by an auxiliary net Q_1 of size[21]

$$f\left(\epsilon, \frac{|\Pi_{>t}|}{\binom{n}{2}}, \frac{\sigma}{2}\right) \leq f\left(\epsilon, \frac{m}{r_0\binom{n}{2}}, \frac{\sigma}{2}\right) \leq f\left(\epsilon, \frac{\lambda}{r_0}, \frac{\sigma}{2}\right). \tag{9}$$

We immediately add the points of Q_1 to our net Q, and remove the sets of $\mathcal{K}_{>t}$ from our family $\mathcal{K}$. Note that choosing r_0 to be a very small (albeit, constant) positive power of $1/\epsilon$ guarantees that the recurrence (9) in the maximum density $\lambda \geq m/\binom{n}{2}$ is invoked only a fixed number of times before λ falls below ϵ; thus, σ remains bounded from below by a sufficiently small constant.

Notice that every remaining set $K \in \mathcal{K}$ belongs to the family $\mathcal{K}\left(P, \Pi \setminus \Pi_{>t}, \epsilon, \sigma/2\right)$. We thus remove the edges of $\Pi_{>t}$ from Π and, accordingly, remove the edges of $\Pi_{>t}$ from the subsets $\Pi_K := \Pi \cap \binom{P_K}{2}$ induced by all the convex sets $K \in \mathcal{K}$. In doing so, we stick with the same remaining family $\mathcal{K}$ even if some of its sets $K \in \mathcal{K}$ are only $(\epsilon, \sigma/2)$-restricted with respect to the refined graph (P, Π).

[21] For the sake of brevity, in our asymptotic analysis we omit the multiplicative constants within the arguments λ' of the recursive terms $f(\epsilon, \lambda', \sigma')$.

Choosing a middle slab in $\Lambda(P, r_0)$. Denote

$$\epsilon_0 := \frac{\sigma\epsilon}{100 r_0}. \tag{10}$$

Definition. Let $K \in \mathcal{K}$ be a convex set. We say that a slab τ in the auxiliary decomposition $\Lambda(P, r_0)$ is a *middle slab* with respect to K if the following conditions are satisfied (see Figure 7):

(M1) $\epsilon_0 n \leq |P_K \cap \tau| \leq \sigma\epsilon n/100$, and

(M2) At least $\sigma\binom{\lceil \epsilon n \rceil}{2}/(8r_0) = \Omega\left(\sigma\epsilon^2 n^2/r_0\right)$ of the edges of Π_K cross τ transversally. [22] (In particular, there is *at least one* such edge $pq \in \Pi_K$ that crosses τ transversally.)

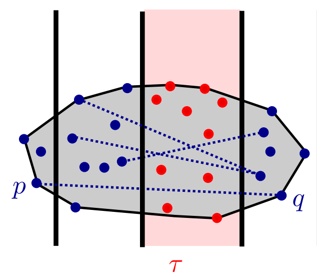

Figure 7: The slab $\tau \in \Lambda(P, r_0)$ is a middle slab for K. The depicted edge $pq \in \Pi_K$ crosses τ transversally.

Informally, the second property (M2) yields $\Omega\left(\epsilon^2 n^2/r_0\right)$ edges in Π_K that can be used to trace the zone of K within the restriction of $\Sigma(r_1)$ to τ. In what follows, this facilitates the study of generalized incidences between such convex sets K and the points of $P_\tau = P \cap \tau$.

Proposition 3.2. *For each remaining convex set $K \in \mathcal{K}$, there is at least one middle slab in $\Lambda(P, r_0)$.*

Proof. Fix a set $K \in \mathcal{K}$. By definition, any convex set $K \in \mathcal{K}$ with at least $\sigma\epsilon n/100$ points in a single slab $\tau \in \Lambda(P, r_0)$ is $(\sigma\epsilon/100)$-crowded and, therefore, already pierced by the net $Q_0 = Q\left(P, r_0, \sigma\epsilon/100\right)$ of Section 3.1. Hence, the second inequality in (M1) holds for *any* slab $\tau \in \Lambda(P, r_0)$.

Let Λ_K be the set of all the slabs τ in $\Lambda(P, r_0)$ that are intersected by K and satisfy $|P_K \cap \tau| \geq \epsilon_0 n = \sigma\epsilon n/(100 r_0)$. Notice that every slab of Λ_K satisfies condition (M1), and the points in the slabs of $\Lambda(P, r_0) \setminus \Lambda_K$ are involved in a total of at most

$$\frac{\sigma\epsilon n}{100 r_0} \cdot (r_0 + 1) \cdot \lceil \epsilon n \rceil \leq \frac{\sigma}{4}\binom{\lceil \epsilon n \rceil}{2}$$

adjacencies with the edges of Π_K. Using that K is $(\epsilon, \sigma/2)$-restricted with respect to the refined graph (P, Π), so that $|\Pi_K| = |\binom{P_K}{2} \cap \Pi| \geq \frac{\sigma}{2}\binom{\lceil \epsilon n \rceil}{2}$, we obtain a subset $\Pi'_K \subseteq \Pi_K$ of at least $\frac{\sigma}{4}\binom{\lceil \epsilon n \rceil}{2}$ edges so that both of their endpoints lie in the slabs of Λ_K.

If no cell in Λ_K satisfies condition (M2), we obtain at least $|\Pi'_K|/2 > \frac{\sigma}{8}\binom{\lceil \epsilon n \rceil}{2}$ edges of Π'_K so that none of them has a transversal crossing with a slab of Λ_K. Thus, by the pigeonhole principle, there must be a slab $\tau \in \Lambda_K$ and a point $p \in P_K \cap \tau$ so that at least $\sigma\binom{\lceil \epsilon n \rceil}{2}/(4\lceil \epsilon n \rceil)$ of its neighbors in the graph Π'_K lie either in τ or in one of its (at most) two neighboring slabs within Λ_K. (Notice

[22] See Section 2.3 for the definition. Here Π_K denotes the induced sub-graph $\Pi \cap \binom{P_K}{2}$ *after* removing the edges of $\Pi_{>t}$.

that these slabs need not be consecutive in $\Lambda(P, r_0)$ or Λ_K.) Since one of these three slabs of Λ_K must then contain at least $\sigma\binom{\lceil \epsilon n \rceil}{2}/(12\lceil \epsilon n \rceil)$ neighbors of p in Π'_K (and we have $n > n_0(\epsilon)$, as defined in (6)), the convex set K is $(\sigma\epsilon/100)$-crowded in $\Lambda(P, r_0)$. Hence, K must have been pierced by the net $Q_0 = Q\,(P, r_0, \sigma\epsilon/100)$ of Stage 0, and already removed from $\mathcal{K}$. This contradiction establishes the claim. □

To recap, for every remaining convex set $K \in \mathcal{K}$ (which is missed by the combination $Q_0 \cup Q_1$) there is at least one middle slab $\tau \in \Lambda(P, r_0)$. Furthermore, each of the edges $pq \in \Pi_K$ (at least $(\sigma/2)\binom{\lceil \epsilon n \rceil}{2}$ in number) that cross τ transversally by property (M2), meets the interiors of at most $t = r_0 r_1 \log r_1$ cells of the decomposition $\Sigma(r_1)$.

In the following Section 3.3 we use these two properties to guarantee that every set $K \in \mathcal{K}$ intersects at most $t = r_0 r_1 \log r_1$ cells of $\Sigma(r_1)$ within some middle slab τ of K. As before, this is achieved at expense of adding an additional small-size auxiliary net to Q.

3.3 Stage 2: Controlling the crossing number in $\Sigma(r_1)$

Definition. To simplify our exposition, for each $K \in \mathcal{K}$ we fix a middle slab $\tau \in \Lambda(P, r)$ with an edge $pq \in \Pi_K$ that crosses τ transversally. (By condition (M2), such an edge pq exists and can be chosen in $\Omega(\sigma\epsilon^2 n^2/r_0)$ possible ways.) In what follows, we refer to τ as the *principal middle slab*, and to pq as the *principal edge*, of K.

For each slab $\tau \in \Lambda(P, r_0)$ we consider the subfamily $\mathcal{K}_\tau \subset \mathcal{K}$ of all the convex sets $K \in \mathcal{K}$ so that τ is their principal middle slab. By Proposition 3.2, we have $\mathcal{K} = \biguplus_{\tau \in \Lambda(P, r_0)} \mathcal{K}_\tau$.

In Section 3.4, we will use the decomposition $\Sigma(r_1)$ to construct a small-size net Q_τ for each sub-family $\mathcal{K}_\tau$. To this end, for every slab $\tau \in \Lambda(P, r_0)$ we consider the restriction

$$\Sigma_\tau := \{\mu \in \Sigma(r_1) \mid \mu \subset \tau\}.$$

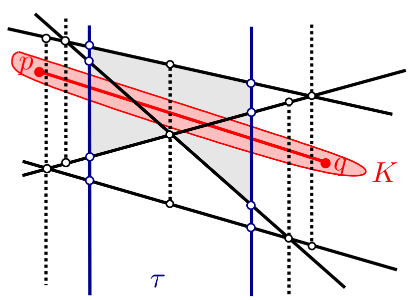

Figure 8: The set K is narrow in Σ_τ because $K \cap \tau$ is contained in the zone of the principal edge $pq \in \Pi_K$ which crosses τ transversally. (The cells of the zone of $K \cap \tau$ within Σ_τ are shaded.)

Definition. Let $K \in \mathcal{K}$ be a convex set, and let τ be the principal middle slab of K. We say that K is *narrow* (in Σ_τ) if the restriction $K \cap \tau$ is contained in the zone of its principal edge $pq \in \Pi_K$ of K within Σ_τ. (By definition, the cells of this zone lie in the zone of pq within the arrangement $\mathcal{A}(\mathcal{R}_1)$.) See Figure 8.

Informally, the narrowness of $K \in \mathcal{K}_\tau$ means that its behaviour is "line-like" in Σ_τ, so the zone of K in Σ_τ can be completely "read off" from its principal edge $pq \in \Pi_K$.

Proposition 3.3. *Let τ be a slab of $\Lambda(P, r_0)$, and let $K \in \mathcal{K}_\tau$ be a narrow convex set. Then K intersects at most $r_0 r_1 \log r_1$ cells of Σ_τ.*

Proof. Let $pq \in \Pi_K$ be the principal edge of K. Since the pq does not belong to the set $\Pi_{>t}$ which we removed at Stage 1, its zone in $\Sigma(r_1)$ (and, in particular, in the restriction Σ_τ of $\Sigma(r_1)$ to the principal middle slab τ) consists of at most $r_0 r_1 \log r_1$ cells. By the narrowness of K, these cells form the zone of $K \cap \tau$ within Σ_τ. $\square$

The net Q_2. We now get rid of the sets $K \in \mathcal{K}$ that are not narrow.

Proposition 3.4. *There is a set $Q_2 \subset \mathbb{R}^2$ of cardinality $O\left(r_0^2 r_1/\epsilon\right)$ that, for each slab $\tau \in \Lambda(P, r_0)$, pierces every convex set $K \in \mathcal{K}_\tau$ that is not narrow in Σ_τ.*

Proof. We first add to Q_2 all the $O(r_1^2)$ vertices of the trapezoids of $\Sigma(r_1)$. We then add to Q_2 the set X of the $r_0 r_1$ intersection points of the r_0 vertical lines of $\mathcal{Y}(P, r_0)$ with the lines of $\mathcal{R}_1$, and construct an even larger family $Y \subset \bigcup \mathcal{Y}(P, r_0)$ by intersecting each line of $\mathcal{Y}(P, r_0)$ with the edges of $P \times X$. Notice that the resulting point set has cardinality at most $O(r_0^2 r_1 n)$, as each line of $\mathcal{Y}(P, r_0)$ contains at most $r_0 r_1 n$ crossing points. Let $C_2 > 0$ be a sufficiently small constant that will be determined in the sequel. For each line $L \in \mathcal{Y}(P, r_0)$ we add to Q_2 every $\lceil C_2 \epsilon n \rceil$-th point of $L \cap Y$, for a total of $O(r_0^2 r_1/\epsilon)$ such points.

Since $r_1 = \Theta\left(\sqrt{1/\epsilon}\right)$, the overall cardinality of our auxiliary net Q_2 is bounded by $O\left(r_1^2 + r_0^2 r_1/\epsilon\right) = O\left(r_0^2 r_1/\epsilon\right)$. It, therefore, suffices to check that Q_2 satisfies the asserted properties with a suitably small choice of the constant $C_2 > 0$. To this end, we fix a slab $\tau \in \Lambda(P, r_0)$ and a convex set $K \in \mathcal{K}_\tau$ that is missed by Q_2.

Let $pq \in \Pi_K$ be the principal edge of K. Since K is missed by the points of Q_2, and pq crosses both of the lines of $\mathcal{Y}(P, r_0)$ that delimit τ, the edge pq is not contained in a line of $\mathcal{R}_1$, and it cannot pass through a vertex of $\Sigma(r_1)$. Assume with no loss of generality that q lies to the right of p.

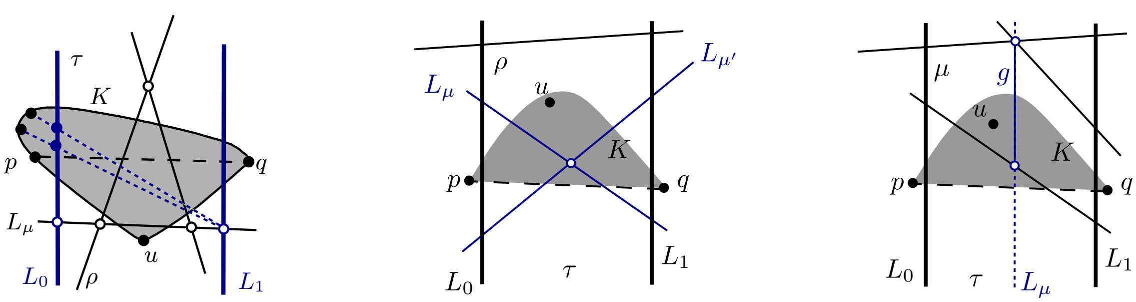

Figure 9: Proof of Proposition 3.4. Showing that every set $K \in \mathcal{K}_\tau$, that is not narrow in Σ_τ, is pierced by a point of Q_2. Left: In the first case, the supercell ρ of u in $\mathcal{A}(\mathcal{R}_1)$ is separated from $pq \cap \tau$ by a line $L_\mu \in \mathcal{R}_1$. The interval $K \cap L_0$ is crossed by every edge that connects a point of $P_K \cap L_0^-$ to the vertex $L_\mu \cap L_1 \in X$. Center: In the second case, ρ lies in the only wedge of $\mathbb{R}^2 \setminus (L_\mu \cup L_{\mu'})$ that is missed by $pq \cap \tau$. The vertex $L_\mu \cap L_{\mu'}$ belongs to K. Right: The cell μ is separated from $pq \cap \rho$ by a vertical wall g, so K contains at least one of the endpoints of g.

Assume for a contradiction that there is a point $u \in K \cap \tau$ that lies in a cell $\mu \in \Sigma_\tau$ outside the zone of pq. Let ρ be the parent cell of μ in the arrangement $\mathcal{A}(\mathcal{R}_1)$. Let L_0 (resp., L_1) be the vertical line adjacent to τ from the left (resp., right). Let L_0^- (resp., L_1^+) denote the halfplane of $\mathbb{R}^2 \setminus L_0$ (resp., $\mathbb{R}^2 \setminus L_1$) containing p (resp., q). We distinguish between three possible cases as illustrated in Figure 9.

1. Both cells μ and ρ are separated from $pq \cap \tau$ by a line $L_\mu \in \mathcal{R}_1$ that misses $pq \cap \tau$.

   Since K is not pierced by X, $K \cap L_0^-$ must lie to the same side of L_μ as p (or, else, K would contain the point $L_0 \cap L_\mu \in Q_2$), and a symmetric property must hold for $K \cap L_1^+$. Since $|P_\tau \cap K| \leq \epsilon n/4$ (by the property (M1) of the middle slab τ with respect to K), at least one of the subsets $P_K \cap L_0^-, P_K \cap L_1^+$, let it be the former set, must contain more than $\epsilon n/4$ points of P_K. Applying Lemma 2.2 to the cell $\Delta \subset \mathbb{R}^2 \setminus (L_0 \cup L_1 \cup L_\mu)$ that contains $pq \cap \tau$ readily implies that $L_0 \cap K$ must be crossed by all the edges connecting the vertex $L_1 \cap L_\mu \in X$ and the $\Omega(\epsilon n)$ points of $P_K \cap L_0^-$. Given a sufficiently small choice of C_1, the intercept $K \cap L$ must contain a point of Q_2.

2. There exist lines L_μ, L'_μ, each crossing $pq \cap \tau$, so that both ρ and $\mu \subset \rho$ lie in the only wedge of $\mathbb{R}^2 \setminus (L_\mu \cup L'_\mu)$ that does not meet pq. In this case, K must be pierced by the vertex $L_\mu \cap L'_\mu \in X \subset Q_2$.

3. The principal edge pq crosses ρ but the cell μ is separated from $pq \cap \rho$ by a vertical line L_μ which supports a vertical wall g on the boundary of μ. (In particular, K must cross g.) Since pq crosses τ transversally, it must cross the line L_μ which is "sandwiched" within τ, and this crossing must happen outside g. Therefore, and due to its convexity, K must contain at least one of the endpoints of g, which again belong to Q_2 as the vertices of $\Sigma(r_1)$.

We conclude that, in either of the above three cases, K must contain a point of Q_2. This contradiction confirms that K is indeed narrow in Σ_τ. □

Remark. A careful look at the proof of Proposition 3.4 indicates that, for each convex set $K \in \mathcal{K}_\tau$ that is missed by Q_2, its portion $K \cap \tau$ is contained in the zone (within Σ_τ) of *any* segment $pq \subset K$ that crosses τ transversally.

We immediately add the points of Q_2 to our net Q, and remove from $\mathcal{K}$ (and, thus, from each subset $\mathcal{K}_\tau$) every set that is pierced by Q_2. As a result, for every $\tau \in \Lambda(P, r_0)$, every remaining set of $\mathcal{K}_\tau$ is narrow in the restriction Σ_τ of $\Sigma(r_1)$ to the principal middle slab τ of K.

Combing the bound $|Q_2| = O\left(r_0^2 r_1/\epsilon\right)$ of Proposition 3.4 with the bounds (8) and (9) on the auxiliary nets Q_0 and Q_1 that were constructed at the previous Stages 0 and 1, so far we have added a total of

$$f\left(\epsilon, \frac{\lambda}{r_0}, \frac{\sigma}{2}\right) + O\left(r_0 \cdot f(\epsilon \cdot \sigma \cdot r_0) + \frac{r_0^2 r_1}{\epsilon}\right) \tag{11}$$

points to the net Q. As previously mentioned, choosing r_0 to be a very small (albeit, constant) positive power of $1/\epsilon$ guarantees that our recurrence (9) in λ has only constant depth; thus, σ remains bounded from below by a certain positive constant. Hence, the second recursive term is essentially linear in $1/\epsilon$. Therefore, the contribution of (11) to the cardinality of Q is effectively dominated by the non-recursive term, which is roughly bounded by $1/\epsilon^{3/2}$ for $r_0 \ll r_1 = \Theta(\sqrt{1/\epsilon})$.[23]

Discussion. Note that a more economical construction of the sets Y_L, for $L \in \mathcal{Y}(P, r_0)$, would have resulted in an auxiliary net of size $O(r_0 r_1/\epsilon)$, and with exactly same properties as argued in Proposition 3.4. However, the actual polynomial dependence on r_0 is immaterial for the eventual recurrence that we derive for $f(\epsilon)$ in Section 3.6.

[23] For $x, y \geq 1$, the notation $x \ll y$ means that $x = O\left(y^{O(\eta)}\right)$. (For $0 < x, y \leq 1$, the notation $x \ll y$ means that $1/y \ll 1/x$.)

3.4 Stage 3: The set P_K – from the low crossing number to expansion in $\Sigma(r_1)$

At this stage we complete the construction of the net Q for $\mathcal{K}(P, \Pi, \epsilon, \sigma)$. As each convex set $K \in \mathcal{K}$ is equipped with the principal middle slab τ and, therefore, assigned to the respective subfamily $\mathcal{K}_\tau \subseteq \mathcal{K}$, it suffices to construct a "local" net Q_τ for each subfamily $\mathcal{K}_\tau$. To this end, we implement the paradigm of Section 2.1 within the restriction Σ_τ of $\Sigma(r_1)$ to τ, each of whose trapezoidal cells contains at most n/r_1^2 points of $P_\tau = P \cap \tau$.

Definition. We fix a sufficiently small constant $0 < \hat{C} \leq 1/120$ that will be determined in the sequel, and denote

$$\epsilon_1 := \frac{\epsilon_0}{40 \log 1/\epsilon} \text{ and } \hat{\epsilon} := \frac{\epsilon_0}{8 r_0 r_1 \log r_1}$$

The auxiliary nets $Q(P, s_0, \epsilon_1/4)$ and $Q^{\triangle}\left(P, \hat{C}\hat{\epsilon}\right)$. We first guarantee that the points of P_K are in a sufficiently convex position, and that they are sufficiently spread within Σ_τ. (The latter property is essential for guessing the splitting line L, whose intercept $K \cap L$ is crossed by many edges of $\binom{P_\tau}{2}$.) To this end, we introduce two auxiliary nets.

1. We construct a finer slab decomposition $\Lambda(P, s_0)$ where $s_0 = \Theta\left(r_0^{1/\eta}\right)$ is yet another small constant power of $1/\epsilon$ that was mentioned in the beginning of Section 3. Since $s_0 + 1$ is a multiple of $r_0 + 1$, we can assume with no loss of generality that $\Lambda(P, s_0)$ is a refinement of $\Lambda(P, r_0)$, that is, we have $\mathcal{Y}(P, s_0) \supset \mathcal{Y}(P, r_0)$. Furthermore, since $s_0 \ll r_1 = \Theta\left(\sqrt{1/\epsilon}\right)$, we can add the lines of $\mathcal{Y}(P, s_0)$ to the sample $\mathcal{R}_1$ with no affect on the asymptotic properties of $\mathcal{A}(\mathcal{R}_1)$ and its vertical decomposition $\Sigma(r_1)$.

   We then apply Lemma 2.4 to construct an auxiliary net $Q(P, s_0, \epsilon_1/4)$ that pierces every convex set that is $(\epsilon_1/4)$-crowded in $\Lambda(P, s_0)$. Notice that

   $$|Q(P, s_0, \epsilon_1/4)| = O\left(s_0 \cdot f\left(\epsilon_1 \cdot s_0\right)\right) = O\left(s_0 \cdot f\left(\epsilon \cdot \frac{s_0 \cdot \sigma}{r_0 \log 1/\epsilon}\right)\right), \tag{12}$$

   where the last estimate uses the definition (10) of ϵ_0 in Section 3.1.

   Upon adding $Q(P, s_0, \epsilon_1/4)$ to Q, we can assume that each remaining convex set K is $(\epsilon_1/4)$-spread in $\Lambda(P, s_0)$.

2. We invoke Theorem 2.1 to construct a strong $\left(\hat{C}\hat{\epsilon}\right)$-net $Q^{\triangle}\left(P, \hat{C}\hat{\epsilon}\right)$ over the set P with respect to triangles, and add its points to the nets Q and Q_τ.

   Notice that this step increases the cardinality of Q by

   $$\left|Q^{\triangle}\left(P, \hat{C}\hat{\epsilon}\right)\right| = O\left(\frac{1}{\hat{\epsilon}} \log \frac{1}{\hat{\epsilon}}\right) = O\left(\frac{r_0^2 r_1}{\epsilon\sigma} \log^2 \frac{1}{\epsilon}\right). \tag{13}$$

   Accordingly, we remove from $\mathcal{K}$ and $\mathcal{K}_\tau$ every convex set that contains a triangle whose interior encloses at least $\hat{C}\hat{\epsilon}n$ points of P.

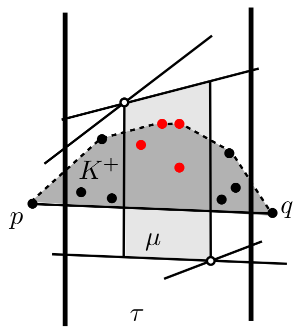

Figure 10: The principal middle slab τ of K, the principal edge $pq \in \Pi_K$, the set $K^+ = \mathsf{conv}(P_K^+ \cup \{p, q\})$, and a cell $\mu \in \Sigma_\tau$. (The points of $P_K(\mu) = P_K^+ \cap \mu$ are colored red.)

We now establish the key properties of the remaining sets $K \in \mathcal{K}$, which are missed by the combination $Q(P, s_0, \epsilon_1/4) \cup Q^{\triangle}\left(P, \hat{C}\hat{\epsilon}\right)$. For each $K \in \mathcal{K}$, we study the distribution of the points of $P_K \subseteq P \cap K$ in the decomposition $\Sigma_\tau \subset \Sigma(r_1)$ of the respective principal middle slab τ of K.

The setup. Fix $K \in \mathcal{K}_\tau$. Since τ is a middle slab for K, it satisfies the criteria (M1) and (M2) detailed in Section 3.2. Namely, we have $|P_K \cap \tau| \geq \epsilon_0 n$ by condition (M1) and the graph Π_K contains $\Omega\left(\frac{\sigma}{r_0}\binom{\lceil \epsilon n \rceil}{2}\right)$ edges that cross τ transversally by condition (M2); these include the unique principal edge $pq \in \Pi_K$ of K. We also assume that K is narrow in Σ_τ. Therefore, by Proposition 3.3, the zone of K in Σ_τ is comprised of the at most $r_0 r_1 \log r_1$ trapezoidal cells that are crossed by pq.

In what follows, we can assume for each $K \in \mathcal{K}_\tau$ that at least $\epsilon_0 n/2 - 2$ of the points of $P_K \cap \tau$ lie above the line $L_{p,q}$ from p to q. Otherwise, K is treated in a fully symmetric manner by reversing the direction of the y-axis. In addition, we can assume that p lies to the left of q. (Recall that pq is not contained in a line of $\mathcal{R}_1$, and it cannot pass through a vertex of Σ_τ.)

Let P_K^+ denote the portion of $P_K \cap \tau$ above the line $L_{p,q}$. Denote $K^+ := \mathsf{conv}(P_K^+ \cup \{p, q\})$. Notice that K^+ is supported by the line $L_{p,q}$ at its boundary edge pq; see Figure 10. Note also that K^+ too is narrow in Σ_τ.

Definition. For each (open) cell $\mu \in \Sigma_\tau$ we denote $P_K(\mu) := P_K^+ \cap \mu$ and $k_\mu := |P_K(\mu)|$. We say that a cell $\mu \in \Sigma_\tau$ is *full* with respect to K^+ if $k_\mu \geq \hat{\epsilon} n \geq 100$, where the second inequality always holds due to the lower bound $n \geq n_0(\epsilon)$ in (6). Let Σ_K denote the sub-collection of all the full cells in Σ_τ.

Proposition 3.5. *At least $\epsilon_0 n/5$ points of P_K^+ lie in the cells of Σ_K.*

Proof. Since K intersects at most $r_0 r_1 \log r_1$ cells of Σ_τ, the non-full cells of Σ_τ contain a total of at most $\epsilon_0 n/8$ points of P_K^+. Since at most $2r_1$ points of P are contained in the lines of $\mathcal{R}_1$, at least $\epsilon_0 n/2 - 2 - \epsilon_0 n/8 - 2r_1$ points of P_K^+ lie in the cells of Σ_K. The claim now follows from the lower bound $n > n_0(\epsilon)$ in (6), and our choice (10) of ϵ_0. $\square$

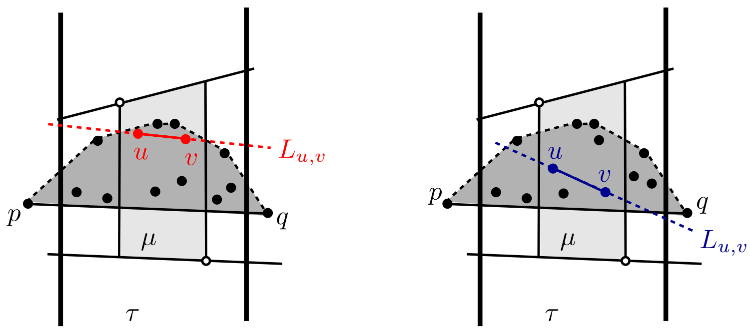

Figure 11: Left: The short edge uv is good for K because all the points of $P_K^+ \cup \{p, q\}$ that lie outside μ are to the same side of $L_{u,v}$. Right: The short edge uv is bad for K.

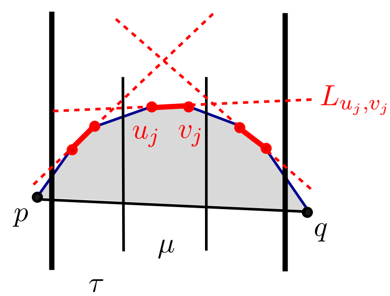

Figure 12: Proposition 3.6 (i) – The good edges $u_j v_j$ with supporting lines L_{u_j,v_j} are depicted. Since these edges lie in distinct cells, they bound a convex polygon (together with the principal edge pq).

Definition. We say that an edge $uv \in \binom{P_\tau}{2}$ is *short* if its endpoints lie in the same cell $\mu \in \Sigma_\tau$.

Notice that the set K^+ contains $\binom{k_\mu}{2} = \Omega(\hat{\varepsilon}^2 n^2)$ short edges within every full cell $\mu \in \Sigma_K$. Let uv be such a short edge whose endpoints belong to $P_K(\mu)$, for some cell μ of Σ_K. We say that uv is *good* for K^+ if all the points of $P_K^+ \cup \{p, q\}$ outside μ lie to the same side of the line $L_{u,v}$, and otherwise we say that uv is *bad* for K^+; see Figure 11.

Informally, the good edges span lines that are nearly tangent to K.[24] In particular, for every good edge uv the corresponding line $L_{u,v}$ must miss the principal edge pq. Since uv lies above pq, the edges uv and pq are boundary edges of a convex quadrilateral.

Proposition 3.6. *(i) Let $u_1v_1, u_2v_2, \ldots, \ldots u_kv_k$ be good edges with respect to K so that no two of these edges lie in the same cell of Σ_K. Then the $k+1$ edges of $\{u_jv_j \mid 1 \le j \le k\} \cup \{pq\}$ lie on the boundary of the same convex $(2k+2)$-gon; see Figure 12.*

(ii) Let $\mu \in \Sigma_K$ be a full cell. Then the points of $P_K(\mu)$ determine at least $\frac{3}{4}\binom{k_\mu}{2}$ good edges.

Proof. The first part readily follows from the definition of a good edge. To see the second part, we consider the set E_{bad} of all the bad edges that are spanned by the points of $P_K(\mu)$. To bound its cardinality, we represent E_{bad} as the union of the following subsets:

[24] We emphasize that the definition of a short edge is independent of K whereas the notion of a good edge assumes both the underlying convex set K, and the principal edge $pq \in \Pi_K$ which crosses τ transversally.

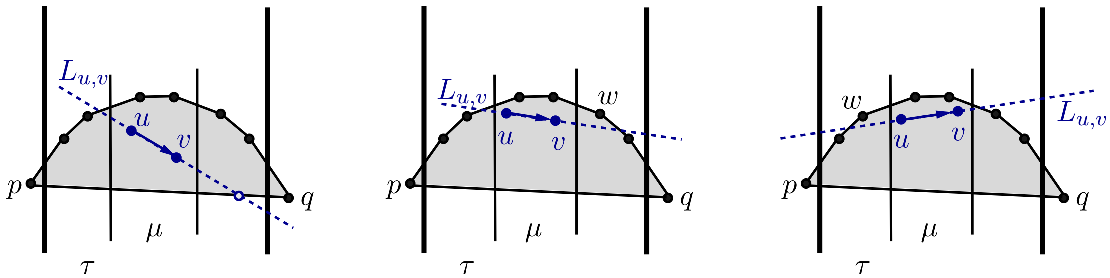

Figure 13: Proof of Proposition 3.6 (ii). The bad edges of E_1, E_2, and E_3 are depicted (resp., left, center, and right). Notice that every edge $uv \in E_1$ is directed towards the intersection $L_{u,v} \cap pq$, whereas the edges of $E_2 \cup E_3$ are directed rightwards.

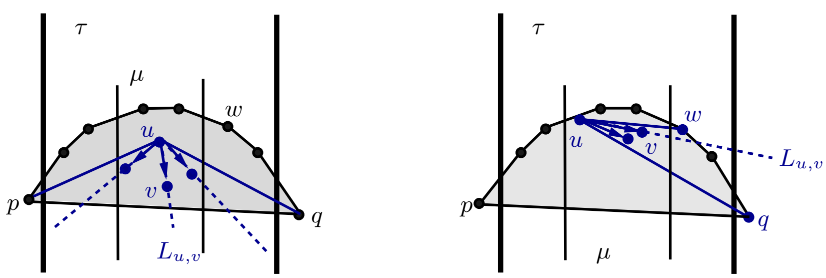

Figure 14: Proof of Proposition 3.6 (ii). The point u and its outgoing bad edges of E_1 and E_2 (resp., left and right). In each case, the other endpoints of the outgoing edges lie inside a triangle $T \subset K$ with apex u.

- E_1 consists of all the bad edges uv whose supporting lines $L_{u,v}$ cross the principal edge pq of K; see Figure 13 (left).
- E_2 (resp., E_3) consists of all the bad edges $uv \in E_{bad} \setminus E_1$ for which there is a point $w \in P_K^+$ that lies in a cell $\mu' \in \Sigma_\tau \setminus \{\mu\}$ crossed by pq to the right (resp., left) of $\mu \cap pq$, so that w is separated by $L_{u,v}$ from pq. See Figure 13 (center and right).

Provided that $\hat{C} < 1/120$, it suffices to show that each of these sets E_1, E_2, E_3 has cardinality at most $10\hat{C}\binom{k_\mu}{2}$. (Notice that E_2 and E_3 may overlap, and for every edge $uv \in E_2 \cup E_3$ the respective line $L_{u,v}$ misses pq.)

Bounding $|E_1|$. Assume for a contradiction that $|E_1| \geq 10\hat{C}\binom{k_\mu}{2}$. We direct every edge $uv \in E_1$ from u to v if the intersection of $L_{u,v}$ with the principal edge pq is closer to v than to u (and otherwise we direct the edge from v to u). By the pigeonhole principle, and since $k_\mu \geq 100$, there is a vertex $u \in P_K(\mu)$ whose out-degree is at least $\hat{C}\hat{\epsilon}n$. Hence, the triangle $T = \triangle pqu \subset K^+$ contains at least $\hat{C}\hat{\epsilon}n$ points of P, so K must have been previously pierced by the auxiliary net $Q^\triangle\left(P, \hat{C}\hat{\epsilon}\right)$, and subsequently removed from $\mathcal{K}_\tau$ and $\mathcal{K}$. See Figure 14 (left).

Bounding $|E_2|$ *and* $|E_3|$. Since the definitions of E_2 and E_3 are fully symmetrical (up to reversing the x-axis), we bound only the cardinality of the former set. We direct every edge $uv \in E_2$ rightwards; see Figure 13 (center).[25]

Once again, we assume for a contradiction that $|E_2| \geq 10\hat{C}\binom{k_\mu}{2}$, so there is a vertex u whose out-degree $d(u)$ is at least $\hat{C}\hat{\epsilon}n$. As in the previous case, we will find a triangle $T \subset K$ which contains all the $d(u) \geq \hat{C}\hat{\epsilon}n$ endpoints of the edges of E_2 that emanate from u; see Figure 14 (right).

Indeed, let $uv_1, uv_2, \ldots, uv_{d(u)} = uv^*$ be the counterclockwise sequence of all the outgoing edges of u in E_2 (so that the occupied sector of $\mathbb{R}^2$ to the right of u does not contain any of the points p, q). By the definition of $E_2 \subset E \setminus E_1$, we can choose a point w in a cell $\mu' \in \Sigma_K$ that is separated by L_{u,v^*} from the edge pq, and is crossed by pq to the right of $\mu \cap pq$. Our analysis is assisted by the following property.

Claim 3.7. *There is a line L' that crosses both edges uw and uq and so that the entire segment $\triangle uqw \cap L$ lies outside the interior of μ.*

Proof of Claim 3.7. If μ and μ' are separated by a line $L' \in \mathcal{R}_1$, then this line must also cross uq; see Figure 15. Indeed, the principal edge pq meets μ and μ' in this horizontal order, so L' crosses pq in-between the intersections $pq \cap \mu$ and $pq \cap \mu'$. Thus, q must lie to the same side of L' as μ'.

On the other hand, if u and w lie within the same cell ρ in $\mathcal{A}(\mathcal{R}_1)$, then their respective cells $\mu \subset \rho$ and $\mu' \subset \rho$ in $\Sigma(r_1)$ must be separated by a vertical wall. Since q lies to the right of τ, and pq crosses μ' to the right of its intersection with μ, both uw and uq must cross the vertical line L' which supports that wall (as depicted in Figure 14 (right)).

In either case, the intersection $\triangle uwq \cap L'$ lies outside the interior of μ by the definition of $\mathcal{A}(\mathcal{R}_1)$ and $\Sigma(r_1)$. □

Let a and b be the respective L'-intercepts of uw and uq as depicted in Figure 15. Claim 3.7, along with the convexity of μ, imply that the triangle $T := \triangle uab$ indeed contains the $d(u) \geq \hat{C}\hat{\epsilon}n$ points $v_1, \ldots, v_{d(u)}$ within μ. As before, this is contrary to the assumption that K is missed by the strong $(\hat{C}\hat{\epsilon})$-net $Q^\triangle\left(P, \hat{C}\hat{\epsilon}\right)$. This contradiction completes the proof of Proposition 3.6. □

[25] Specifically, if this edge is directed from u to v then pu and vq are edges of the convex quadrilateral $\mathsf{conv}(p, q, u, v)$.

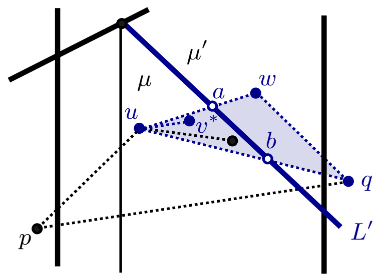

Figure 15: Proof of Claim 3.7. If u and w are separated by a line L' of $\mathcal{R}_1$, this line must also cross the edge uq, and the intersection $\triangle uwq \cap L'$ must lie outside the interior of μ.

Definition. Let μ be a full cell of Σ_K. We orient every good edge within μ from the left to the right. We say that a point $u \in P_K(\mu)$ is *good* if it is adjacent to at least $k_\mu/10$ outgoing good edges (where, as before, k_μ denotes $|P_K(\mu)|$).

By combining Proposition 3.6 (ii) with Proposition 3.5, and recalling that every full cell $\mu \in \Sigma_K$ satisfies $k_\mu \geq 100$, we obtain the following property:

Proposition 3.8. *Every full cell $\mu \in \Sigma_K$ contains at least $k_\mu/4$ good points of P_K^+, for a total of at least $\epsilon_0 n/20$ such points.*

Definition. Let $u \in P_K^+$ be a good point. The *characteristic wedge $\mathcal{W}_K(u)$ at u* is the smallest planar wedge with apex u that lies in the hyperplane to the right of u, and contains uq along with all the outgoing good edges uv of u (but does not contain up); see Figure 16 (left). Note that $\mathcal{W}_K(u)$ lies entirely in the halfplane to the right of u.

Denote

$$D(u) := |\,(P_\tau \cap \mathcal{W}_K(u)) \setminus \{u\}|.$$

That is, $D(u)$ the number of the edges in $uw \in \binom{P_\tau}{2}$ that are adjacent to u and lie within the triangle $\mathcal{W}_K(u) \cap \tau$.[26] Since the characteristic wedge $\mathcal{W}_K(u)$ encompasses all the outgoing good edges of $u \in P_K(\mu)$, we trivially have $D(u) \geq k_\mu/10 \geq \hat{\epsilon}n/10$, and $D(u)$ can be much larger than ϵn due to the additional points of $P_\tau \setminus P_K$ that potentially lie within $\mathcal{W}_K(u) \cap \tau$.

To interpolate between the two favourable scenarios sketched in Section 2.1, we subdivide the good points $u \in P_K^+$ into $O(\log 1/\epsilon)$ classes according to their respective degrees $D(u)$.

Definition. Let i be an integer. We denote $\delta_i := 2^i\epsilon_1/4 = 2^i\epsilon_0/(160 \log 1/\epsilon)$. We say that a good point $u \in P_K^+$ is of *type* i if $\delta_i n \leq D(u) < \delta_{i+1}n$. We use $P_K(i)$ to denote the (possibly empty) subset of all the good points of i-type in P_K^+.

Proposition 3.9. *There is an integer $\lfloor \log(2\hat{\epsilon}/5\epsilon_1) \rfloor \leq i \leq \log 1/\epsilon_1$ so that $|P_K(i)| \geq \epsilon_1 n$.*

Proof. Since the degrees $D(u)$ of the good points $u \in P_K^+$ satisfy $\hat{\epsilon}n/10 \leq D(u) \leq n/\lceil (r_0+1) \rceil$, each point $u \in P_K^+$ is of some type i in the asserted range, where the second inequality $i \leq \log 1/\epsilon_1$ uses (5) (which holds for all $\epsilon < \tilde{\epsilon}$). Combining (5) with the definition of ϵ_1 and $\hat{\epsilon}$ in the beginning of Section 3.4, we conclude that there exist at most $2 \log 1/\epsilon$ types with $|P_K(i)| > 0$. By Proposition 3.8 there

[26] Recall that P_τ denotes the set $P \cap \tau$. Notice that many of these points $w \in (P_\tau \cap \mathcal{W}_K(u)) \setminus \{u\}$, which contribute to the count $D(u)$, may not belong to P_K or even to $P \cap K$.

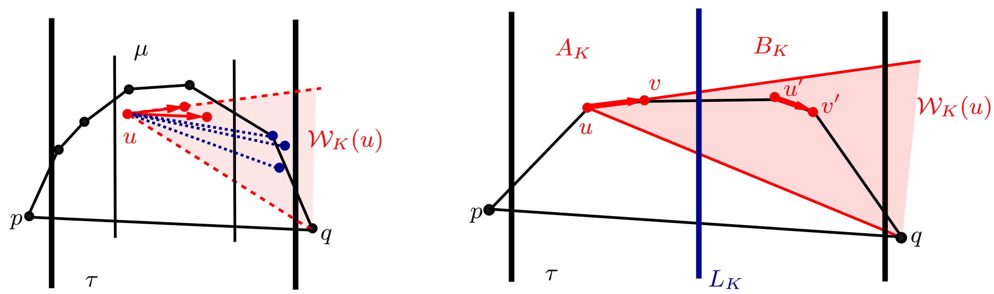

Figure 16: Left: The characteristic wedge $\mathcal{W}_K(u)$ to the right of u encompasses uq and all the outgoing good edges of u. $D(u)$ denotes the overall number of the edges in $\binom{P_\tau}{2}$ that are adjacent to u and lie within the triangle $\mathcal{W}_K(u) \cap \tau$. Right: For each set K of type i we use the principal vertical line $L_K \in \mathcal{Y}(P, s_0)$ to split the subset $P_K(i)$ of the good points of type i into subsets A_K and B_K, of cardinality at least $\epsilon_1 n/4$ each. For each point $u \in A_K$, the characteristic wedge $\mathcal{W}_K(u)$ contains all the points of B_K.

exist at least $\epsilon_0 n/20$ good points in P_K^+, and each of them belongs to one of these types i. Hence, by the pigeonhole principle, there must be an integer i so that $|P_K(i)| \geq \epsilon_0 n/(40 \log 1/\epsilon) = \epsilon_1 n$. $\square$

Definition. We say that a convex set $K \in \mathcal{K}$ is *of type* i if $i := \min\{j \in \mathbb{Z} \mid |P_K(j)| \geq \epsilon_1 n\}$. According to Proposition 3.9, each convex set is of exactly one well-defined type $i \in \mathbb{Z}$.

Let $K \in \mathcal{K}$ be a convex set of type i. Since K is $(\epsilon_1/4)$-spread in the secondary slab decomposition $\Lambda(P, s_0)$, there must be a line $L_K \in \mathcal{Y}(P, s_0)$, within the principal middle slab τ of K, so that at least $\epsilon_1 n/4$ good points in $P_K(i)$ lie to each side of L_K. In what follows, we refer to L_K as the *principal vertical line* of K.

Remark. For the success of our strategy sketched in Section 2.1, it is quintessential that the (remaining) convex sets $K \in \mathcal{K}$ can be split using at most s_0 principal vertical lines $L_K \in \mathcal{Y}(P, s_0)$.

Definition. Let $K \in \mathcal{K}$ be a convex set of type i. We use A_K (resp., B_K) to denote the subset of the good points in $P_K(i)$ that lie to the left (resp., right) of L_K; see Figure 16 (right). Note that $|A_K|, |B_K| \geq \epsilon_1 n/4$.

Proposition 3.10. *Let $K \in \mathcal{K}$ be a convex set, and u be a good point in A_K. Then the characteristic wedge $\mathcal{W}_K(u)$ at u contains at least $k_\mu/10 \geq \hat{\epsilon} n/10$ outgoing good edges uv, and all the points of B_K.*

Since we have $|B_K| \geq \epsilon_1 n/4$ for every remaining convex set $K \in \mathcal{K}$ of type i, the proposition implies the following property.

Corollary 3.11. *Every remaining set $K \in \mathcal{K}$ has type $0 \leq i \leq \log 1/\epsilon_1$.*

Proof of Proposition 3.10. The desired number of the good edges in the characteristic wedge $\mathcal{W}_K(u)$ at any point $u \in A_K$ follows from the construction of $\mathcal{W}_K(u)$ (and because each point in A_K is good and, therefore, satisfies, $D(u) \geq k_\mu/10$).

To show that $\mathcal{W}_K(u)$ also contains all the points of B_K, let $\mu \in \Sigma_K$ be the full cell that contains u. Since every point $u' \in B_K$ lies in a cell $\mu' \in \Sigma_K$ to the right of μ and L_K, the desired property follows from the first part of Proposition 3.6. Indeed, since we have $\min\{D(u), D(u')\} \geq \min\{k_\mu/10, k_{\mu'}/10\} \geq 10$, good edges uv and $u'v'$ can be chosen within, respectively, μ and μ' so that the vertices u, v, u', v' appear in this horizontal order. Since p, u, v, u', v', q form a convex chain by Proposition 3.6, and the characteristic wedge $\mathcal{W}_K(u)$ contains the edges uv and uq, it must also contain u', as depicted in Figure 16 (right). $\square$

To pierce the remaining sets $K \in \mathcal{K}_\tau$ for each $\tau \in \Lambda(P, r_0)$, we combine the following two properties whose somewhat technical proofs are relegated to Section 3.5.

Lemma 3.12. *Let τ be a slab in $\Lambda(P, r_0)$, and let $K \in \mathcal{K}_\tau$ be a convex set of type $0 \leq i \leq \log 1/\epsilon_1$. For each $u \in \mathcal{A}_K$, its respective characteristic wedge $\mathcal{W}_K(u)$ contains $\Omega(\delta_i n)$ edges of $\binom{P_\tau}{2}$ that are adjacent to u and cross the principal vertical line L_K of K within the interval $K \cap L_K$ (for a total of $\Omega(\epsilon_1 \delta_i n^2)$ such edges that cross $K \cap L_K$). See Figure 17.*

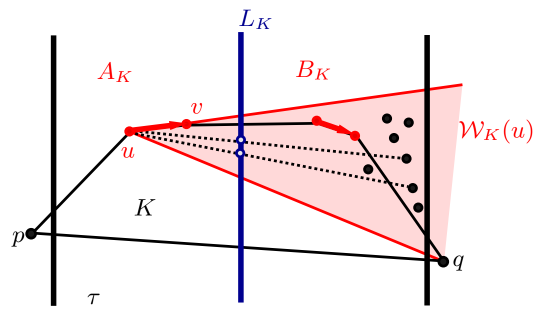

Figure 17: Lemma 3.12 – a schematic illustration. For each $u \in \mathcal{A}_K$, its characteristic wedge $\mathcal{W}_K(u)$ contains $\Omega(\delta_i n)$ edges that are adjacent to u and cross the interval $K \cap L_K$.

Lemma 3.13. *For each $\tau \in \Lambda(P, r_0)$, and each $0 \leq i \leq \log 1/\epsilon_1$, there is a subset $\Pi(\tau, i) \subset \binom{P_\tau}{2}$ with the following properties:*

i. We have that

$$|\Pi(\tau, i)| = O\left(\frac{\delta_i n^2}{r_1^2 \hat{\epsilon}}\right). \tag{14}$$

ii. For each convex set $K \in \mathcal{K}_\tau$ of the type i, and each point $u \in A_K$, the set $\Pi(\tau, i)$ contains all the edges $uw \in \binom{P_\tau}{2}$ that are adjacent to u and lie within the characteristic wedge $\mathcal{W}_K(u)$ at u.

Notice that the density of the graph $\Pi(\tau, i)$ is proportional to δ_i, giving rise to the following tradeoff:

1. If δ_i exceeds $r_1 \epsilon_1$ then we are in the first favourable scenario of Section 2.1 – combining Lemma 3.12 and Lemma 3.13 (ii) for each $u \in A_K$ yields that the intercept $K \cap L_K$ is crossed by roughly $(r_1 \epsilon_1 n) \cdot (\epsilon_1 n) \simeq r_1 \epsilon^2 n$ edges.
2. On the other hand, as δ_i approaches ϵ, the set $\Pi(\tau, i)$ contains roughly n^2/r_1 edges, which gives rise to the second favourable scenario of Section 2.1 (e.g., via Lemma 2.5, or through a direct application of Lemma 3.12).

Our net $Q(\tau, i)$ for the convex sets $K \in \mathcal{K}_\tau$ of type i interpolates between these two extreme cases.

The nets Q_τ and Q_3. For every $0 \le i \le \log 1/\epsilon_1$, and every line $L \in \mathcal{Y}(P, s_0)$ within τ, we add every intersection of L with an edge of $\Pi(\tau, i)$ to the set $X(L, i)$. We then select every $\lceil C_3 \epsilon_1 \delta_i n^2 \rceil$-th point of $X(L, i)$ into our net $Q(L, i)$, for a sufficiently small constant $C_3 > 0$.

We then define

$$Q(\tau, i) := \bigcup \{Q(L, i) \mid L \in \Lambda(P, s_0), L \subset \tau\}.$$

and

$$Q_\tau := \bigcup_{0 \le i \le \log 1/\epsilon_1} Q(\tau, i)$$

The complete net Q_3 at Stage 3 is then given by

$$Q_3 := Q(P, s_0, \epsilon_1/4) \cup Q^{\triangle}\left(P, \hat{C}\hat{\epsilon}\right) \cup \bigcup_{\tau \in \Lambda(P, r_0)} Q_\tau.$$

Theorem 3.14. *With a suitably small constant $C_3 > 0$, the net Q_3 pierces every remaining convex set in $\mathcal{K}$ that is missed by the combination $Q_0 \cup Q_1 \cup Q_2$. Furthermore, we have that*

$$|Q_3| = O\left(s_0 \cdot f\left(\epsilon \cdot \frac{s_0 \cdot \sigma}{r_0 \log 1/\epsilon}\right) + \frac{r_0^2 r_1}{\epsilon\sigma} \log^2 \frac{1}{\epsilon} + \frac{s_0 r_0^3 \log^3 1/\epsilon}{\sigma^2 r_1 \epsilon^2}\right). \tag{15}$$

Proof. Using the definition of Q_3, we argue for each slab $\tau \in \Lambda(P, r_0)$ that the respective net Q_τ pierces all the sets $K \in \mathcal{K}_\tau$ that were missed by the previous nets $Q_0, Q_1, Q_2, Q(P, s_0, \epsilon_1/4)$ and $Q^{\triangle}\left(P, \hat{C}\hat{\epsilon}\right)$. It suffices to check, for all $\tau \in \Lambda(P, r_0)$, and all $0 \le i \le \log 1/\epsilon_1$, that every convex set $K \in \mathcal{K}_\tau$ of type i is pierced by one of the nets $Q(L, i) \subset Q_\tau$ whose vertical lines $L \in \mathcal{Y}(P, s_0)$ lie within τ.

Indeed, according to Lemma 3.12, every point $u \in A_K$ gives rise to $\Omega(\delta_i n)$ outgoing edges that cross the intercept $K \cap L_K$ with the principal vertical line $L_K \in \mathcal{Y}(P, s_0)$ which separates A_K and B_K, for a total of $\Omega(\epsilon_1 \delta_i n^2)$ such edges. Hence, choosing a small enough constant $C_3 > 0$ guarantees that K is pierced by $Q(L_K, i)$.

For every type $1 \le i \le 1/\epsilon_1$, and every line $L \in \mathcal{Y}(P, s_0)$ within τ, the cardinality of $Q(L, i)$ is bounded by

$$O\left(\frac{|X(L, i)|}{\delta_i \epsilon_1 n^2}\right) = O\left(\frac{|\Pi(\tau, i)|}{\delta_i \epsilon_1 n^2}\right) = O\left(\frac{\delta_i n^2}{r_1^2 \hat{\epsilon} \delta_i \epsilon_1 n^2}\right) = O\left(\frac{r_0 \log r_1 \log 1/\epsilon}{r_1 \epsilon_0^2}\right).$$

where the second equality uses the bound of Lemma 3.13 (ii), and the third one uses the definitions of ϵ_1 and $\hat{\epsilon}$.

Recall that $\Lambda(P, s_0)$ is a refinement of $\Lambda(P, r_0)$, every slab $\tau \in \Lambda(P, r_0)$ contains $O(s_0/r_0)$ lines of $\mathcal{Y}(P, s_0)$. Using this and the definition (10), we can bound the cardinality of Q_τ by

$$O\left(\frac{s_0}{r_0} \log(1/\epsilon) \cdot \frac{r_0 \log r_1 \log 1/\epsilon}{r_1 \epsilon_0^2}\right) = O\left(\frac{s_0 r_0^2 \log^3 1/\epsilon}{\sigma^2 r_1 \epsilon^2}\right).$$

Repeating this bound for each slab $\tau \in \Lambda(P, r_0)$ and combining it with the prior bounds (12) and (13) on the cardinalities of the nets $Q(P, s_0, \epsilon_1/4)$ and $Q^{\triangle}\left(P, \hat{C}\hat{\epsilon}\right)$, we conclude that the overall cardinality of Q_3 indeed satisfies the bound (15). $\square$

Remark. Since s_0 and r_0 are very small (albeit, fixed) positive power of $1/\epsilon$ that satisfy $s_0 \gg r_0 \gg 1/\sigma$, the recursive term in (15) is again near-linear in $1/\epsilon$. Furthermore, the two non-recursive terms sum up to roughly $r_1/\epsilon + 1/(r_1\epsilon^2)$, so choosing $r_1 = \Theta(\sqrt{1/\epsilon})$ renders them close to $1/\epsilon^{3/2}$.

In Section 3.6 we combine (15) with the bounds on the sizes of the auxiliary nets Q_0, Q_1, and Q_2 of the previous Stages $0-2$ to derive a recurrence for $f(\epsilon)$ whose solution is close to $1/\epsilon^{3/2}$.

3.5 Proofs of Lemmas 3.12 and 3.13

Proof of Lemma 3.12. Refer to Figure 18. Fix a point $u \in A_K$, and let uv be the good edge that delimits from above its characteristic wedge $\mathcal{W}_K(u)$. (In other words, uv attains the largest slope among the good edges that emanate from u to the right.)

By Proposition 3.10, the wedge $\mathcal{W}_K(u)$ contains all the points of B_K. We fix any of these points $u' \in B_K$ together with the edge $u'v'$ which delimits from above the respective wedge $\mathcal{W}_K(u')$. Since the point u' too is of type i, the characteristic wedge $\mathcal{W}_K(u')$ contains at least $\delta_i n$ points $w \in P_\tau$. Since the edges uv and $u'v'$ are good, Proposition 3.6 implies that the three edges $uv, u'v'$ and pq form a convex 6-gon G. It, therefore, suffices to show that all the resulting edges uw cross the intercept $G \cap L_K \subset K \cap L_K$, where L_K denotes the principal vertical line of K.

Let $L_0 \in \mathcal{Y}(P, r_0)$ (resp., $L_1 \in \mathcal{Y}(P, r_0)$) be the line that supports τ from the left (resp., right). Notice that L_0 is crossed by the edges pq and pu, and L_1 is crossed by the edges pq and $v'q$, and none of the remaining edges uv, vu', and $u'v'$, of G crosses L_0 or L_1. Thus, the intersection $G_\tau := G \cap \tau$ is a convex 8-gon. The claim now follows since (1) $\mathcal{W}_K(u')$ is separated from u by L_K, and (2) every point $w \in \mathcal{W}_K(u) \cap P_\tau$ lies either inside $G_\tau \subset K$, or in the triangular "ear" that is adjacent to the edge $v'q \cap \tau$ of G_τ and delimited by L_1 and $L_{u'v'}$. $\square$

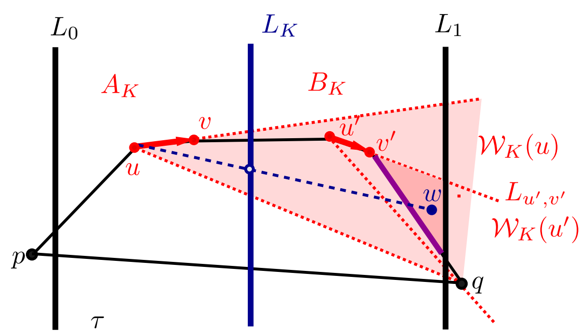

Figure 18: Proof of Lemma 3.12. The convex 6-gon $G = \mathsf{conv}(p, q, u, v, u', v')$ is depicted. Every point $w \in \mathcal{W}_K(u')$ is separated from u by L_K. It lies either inside K, or in the triangular "ear" that is adjacent to the edge $v'q \cap \tau$ of $G_\tau = G \cap \tau$ and delimited by L_1 and $L_{u',v'}$.

Proof of Lemma 3.13. We first describe the sparse subgraph $\Pi(\tau, i) \subset \binom{P_\tau}{2}$ for all $\tau \in \Lambda(P, r_0)$ and $0 \le i \le \log 1/\epsilon_1$.

The graph $\Pi(\tau, i)$. Denote $P_\tau := P \cap \tau$ and $n_\tau := |P \cap \tau|$. For each $u \in P_\tau$ which lies in some cell $\mu \in \Sigma_\tau$ we partition the $n_\tau - 1$ adjacent edges $uv_1, \ldots, uv_{n_\tau - 1} \in \binom{P_\tau}{2}$ (which appear in this clockwise order around u) into $z = O(1/\delta_i)$ blocks $\mathcal{E}_j$, for $0 \le j \le z-1$, so that every block but the last one contains $\lceil 2\delta_i n \rceil$ edges, and the last block contains at most $\lceil 2\delta_i n \rceil$ edges. If $z \ge 4$, we define z canonical sectors with apex u, where each sector $\mathcal{W}_j(u)$ encompasses three consecutive

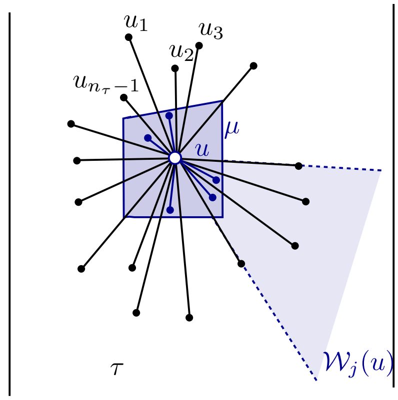

Figure 19: Proof of Lemma 3.13 – defining the sparse graph $\Pi(\tau, i) \subset \binom{P_\tau}{2}$. We define $z = O(1/\delta_i)$ sectors $\mathcal{W}_j(u)$ with apex p. In each sector, the number of the edges of $\binom{P_\tau}{2}$ that are adjacent to p ranges between $2\lceil 2\delta_i n \rceil$ and $3\lceil 2\delta_i n \rceil$. We add the edges of $\mathcal{W}_j(u)$ to $\Pi(\tau, i)$ only if this sector is *rich* and encompasses at least $\hat{\epsilon}n/10$ short edges.

blocks $\mathcal{E}_j, \mathcal{E}_{j+1}, \mathcal{E}_{j+3}$ of edges, and the indexing is modulo z. See Figure 19. Otherwise (i.e., if $\lceil 2\delta_i n \rceil > (n-1)/3$), we define only one sector $\mathcal{W}_0(u) = \mathbb{R}^2$.

Notice that, given that $z \geq 4$, the neighboring sectors overlap, each sector $\mathcal{W}_j(u)$ satisfies $2\lceil 2\delta_i n \rceil \leq |(\mathcal{W}_j(u) \cap P_\tau) \setminus \{u\}| \leq 3\lceil 2\delta_i n \rceil$, and each edge uv lies in exactly three of the sectors of u.

We say that the sector $\mathcal{W}_j(u)$ is *rich* if $|(\mathcal{W}_j(u) \cap P_\mu) \setminus \{u\}| \geq \hat{\epsilon}n/10$. In other words, the sector $\mathcal{W}_j(u)$ must contain at least $\hat{\epsilon}n/10$ short edges uv.

We add to $\Pi(\tau, i)$ every edge uv that lies a rich sector of at least one of its endpoints u or v.

Analysis. To see the first property of $\Pi(\tau, i)$, it is sufficient to show that any point $u \in P_\tau$ contributes $O\left(\frac{\delta_i n}{r_1^2 \hat{\epsilon}}\right)$ edges to the set $\Pi(\tau, i)$.

Indeed, recall that for each cell $\mu \in \Sigma_\tau$ we have that $n_\mu = |P_\mu| \leq n/r_1^2$. Therefore, the pigeonhole principle implies that for each $u \in P_\mu$ there can be only $O\left(\frac{n}{r_1^2 \hat{\epsilon} n}\right) = O\left(\frac{1}{r_1^2 \hat{\epsilon}}\right)$ rich sectors $\mathcal{W}_j(u)$, which satisfy $|\mathcal{W}_j(u) \cap P_\mu \setminus \{u\}| \geq \hat{\epsilon}n/10$, and any such sector contributes $O(\delta_i n)$ edges to $\Pi(\tau, i)$.

For the second property, we recall that, for every good point $u \in A_K$ that lies in some full cell $\mu \in \Sigma_K$, the respective wedge $\mathcal{W}_K(u)$ contains at most $2\delta_i n$ outgoing edges uw within τ and, therefore, is contained in (at least) one of the sectors $\mathcal{W}_j(u)$. Proposition 3.10 now implies that this sector $\mathcal{W}_j(u)$ is rich, for it contains at least $k_\mu/10 \geq \hat{\epsilon}n/10$ outgoing short edges uv (where, as before, k_μ denotes $|P_K(\mu)|$). $\square$

3.6 The final recurrence

In this section we derive the complete recurrence of the general form (1) for the quantity $f(\epsilon, \lambda, \sigma)$. As mentioned in Section 2.2, this will yield a simpler recurrence of the form (2) for the quantity $f_2(\epsilon) = f(\epsilon)$. This ultimately solves to $f(\epsilon) = O\left(\frac{1}{\epsilon^{3/2+\gamma}}\right)$, where $\gamma > 0$ is an arbitrary small constant that has been fixed in the beginning of this proof.

As mentioned in the beginning of this section, our recurrence will involve the auxiliary parameters $r_0 \ll s_0$ which are very small (albeit fixed) degrees of $1/\epsilon$ that depend on γ and satisfy $r_0 = \Theta(s_0^\eta)$ and $s_0 = \Theta\left(1/\epsilon^\eta\right)$, where $\eta := \gamma/100$.

As mentioned in Section 2, we fix a suitably small constant $0 < \tilde{\epsilon} < 1$ that satisfies (5) and use the old bound $f(\epsilon) = O\left(\frac{1}{\tilde{\epsilon}^2}\right) = O(1)$ of Alon *et al.* [2] whenever $\epsilon \geq \tilde{\epsilon}$. (The choice of $\tilde{\epsilon}$ will affect the multiplicative constant in the eventual asymptotic bound on $f(\epsilon)$.) Assume then that $\epsilon < \tilde{\epsilon}$.

Bounding $f(\epsilon, \lambda, \sigma)$. To obtain a bound of the general form (1) for $f(\epsilon, \lambda, \sigma)$, where $\epsilon < \tilde{\epsilon}$ and $\lambda > \epsilon$, we fix a family $\mathcal{K} = \mathcal{K}(P, \Pi, \epsilon, \sigma)$ that satisfies $|\Pi|/\binom{P}{2} \leq \lambda$, and bound the overall cardinality of the point transversal Q for $\mathcal{K}$ that was constructed in Sections 3.1 through 3.4. As explained in the beginning of Section 3, we can also assume that the cardinality of P is bounded from below by (6), so that $|P| = n > n_0(\epsilon)$.

Combining (11) and (15) yields the following bound on the overall cardinality of our net Q:

$$f\left(\epsilon, \frac{\lambda}{r_0}, \frac{\sigma}{2}\right) +$$

$$+O\left(r_0 \cdot f(\epsilon \cdot \sigma \cdot r_0) + \frac{r_0^2 r_1}{\epsilon} + s_0 \cdot f\left(\epsilon \cdot \frac{s_0 \cdot \sigma}{r_0 \log 1/\epsilon}\right) + \frac{r_0^2 r_1}{\epsilon\sigma} \log^2 \frac{1}{\epsilon} + \frac{s_0 r_0^3 \log^3 1/\epsilon}{\sigma^2 r_1 \epsilon^2}\right).$$

By substituting $r_1 = \Theta\left(\sqrt{1/\epsilon}\right)$ and rearranging the terms, we conclude for all $\epsilon < \tilde{\epsilon}$ and $\lambda > \epsilon$ that

$$f(\epsilon, \lambda, \sigma) \leq f\left(\epsilon, \frac{\lambda}{r_0}, \frac{\sigma}{2}\right) + O\left(\Psi(\epsilon, \sigma)\right), \tag{16}$$

where

$$\Psi(\epsilon, \sigma) := s_0 \cdot f\left(\epsilon \cdot s_0 \cdot \frac{\sigma}{r_0 \log 1/\epsilon}\right) + r_0 \cdot f(\epsilon \cdot r_0 \cdot \sigma) + \frac{r_0^2}{\epsilon^{3/2}} + \frac{r_0^2}{\sigma\epsilon^{3/2}} \log^2 \frac{1}{\epsilon} + \frac{s_0 r_0^3 \log^3 1/\epsilon}{\sigma^2 \epsilon^{3/2}}.$$

Our choice of r_0 and s_0, in combination with (5), yields

$$\Psi(\epsilon, \sigma) = O\left(s_0 \cdot f\left(\epsilon \cdot s_0^{1-2\eta} \cdot \sigma\right) + r_0 \cdot f(\epsilon \cdot r_0 \cdot \sigma) + \frac{1}{\sigma^2 \epsilon^{3/2+7\eta}}\right).$$

The recurrence for $f(\epsilon)$. We begin with $f(\epsilon) = f(\epsilon, 1, 1)$ and recursively apply the inequality (16) to the "leading" term, which involves the density λ, while keeping the parameters ϵ, r_0, s_0 and r_1 fixed. This recurrence in λ bottoms out when the value of λ falls below ϵ. Since r_0 is a fixed (though very small) positive power of $1/\epsilon$, the inequality (16) is applied $J = O\left(\log_{r_0} 1/\epsilon\right) = O(1)$ times. Hence, the value of the restriction threshold σ in the j-th application is bounded from below by $1/2^{j-1} = \Theta(1)$. Using the trivial property that $f(\epsilon, \lambda', \sigma') \geq f(\epsilon, \lambda, \sigma)$ and $\Psi(\epsilon, \sigma') \geq \Psi(\epsilon, \sigma)$ for all $0 \leq \lambda \leq \lambda' \leq 1$ and $0 < \sigma' \leq \sigma \leq 1$, and that $J = O(1)$, we conclude that

$$f(\epsilon) = f(\epsilon, 1, 1) \leq f\left(\epsilon, \epsilon, 2^{-J}\right) + \sum_{i=1}^{J} \Psi\left(\epsilon, 2^{-j+1}\right) = f\left(\epsilon, \epsilon, 2^{-J}\right) + O\left(\Psi\left(\epsilon, 2^{-J}\right)\right). \tag{17}$$

Note that

$$\Psi\left(\epsilon, 2^{-J}\right) = O\left(s_0 \cdot f\left(\epsilon \cdot s_0^{1-2\eta}\right) + r_0 \cdot f(\epsilon \cdot r_0) + \frac{1}{\epsilon^{3/2+7\eta}}\right). \tag{18}$$

To bound $f\left(\epsilon, \epsilon, 2^{-J}\right)$, we invoke Lemma 2.5 with $r := r_0$, which yields

$$f(\epsilon, \epsilon, 2^{-J}) = O\left(r_0 \cdot f\left(\epsilon \cdot r_0\right) + \frac{r_0^2}{\epsilon}\right) = O\left(r_0 \cdot f\left(\epsilon \cdot r_0\right) + \frac{1}{\epsilon^{1+2\eta}}\right). \tag{19}$$

Substituting (19) and (18) into (17) readily gives

$$f(\epsilon) = O\left(s_0 \cdot f\left(\epsilon \cdot s_0^{1-2\eta}\right) + r_0 \cdot f\left(\epsilon \cdot r_0\right) + \frac{1}{\epsilon^{3/2+7\eta}}\right), \tag{20}$$

where the implicit constants do not depend on the particular choice of the constant upper threshold $\tilde{\epsilon}$ for ϵ as long as the inequality (5) is satisfied. (As previously mentioned, we routinely omit the constant factors within the recursive terms of the form $f(\epsilon \cdot h)$.) A suitably small choice of $\tilde{\epsilon} > 0$ guarantees that ϵ indeed increases by a factor of at least 2 with each invocation of the recurrence.

Solving the recurrence for $f(\epsilon)$. This last recurrence (20) bottoms out when $\epsilon \geq \tilde{\epsilon}$, in which case we have $f(\epsilon) = O\left(1/\tilde{\epsilon}^2\right) = O(1)$. Since $\eta = \gamma/100$, fixing a suitably small constant threshold $\tilde{\epsilon} > 0$, and following the standard induction argument which applies to recurrences of this type (see, e.g., [38], and also [28, 46] and [47, Section 7.3.2]), yields

$$f(\epsilon) \leq \frac{F}{\epsilon^{3/2+\gamma}}, \tag{21}$$

for all $\epsilon > 0$, where $F \geq f\left(\tilde{\epsilon}\right) = O\left(1/\tilde{\epsilon}^2\right)$ [2] is a suitably large constant.

For the sake of completeness, we spell out the key details of this generic induction.[27] We first choose a constant threshold $\varepsilon(\gamma)$ so that the following conditions are satisfied for all $\epsilon \leq \varepsilon(\gamma)$: (i) the inequality (5) holds, and (ii) each of the recursive terms of the form $f(\epsilon')$ in (20) involves $\epsilon' > 2\epsilon$.

Since the inequality (21) trivially holds for $\tilde{\epsilon} \leq \epsilon \leq 1$ whenever $F \geq f\left(\tilde{\epsilon}\right)$, it suffices to choose the constant threshold $\tilde{\epsilon} \leq \varepsilon(\gamma)$ that would facilitate the induction step for all the smaller values $\epsilon < \tilde{\epsilon}$. To this end, we plug the desired induction assumption (with so far unknown $F > 0$) in (20) so as to replace the two recursive terms $f\left(\epsilon \cdot s_0^{1-2\eta}\right)$ and $f\left(\epsilon \cdot r_0\right)$. This substitution readily yields

$$f(\epsilon) \leq \frac{F}{\epsilon^{3/2+\gamma}} \cdot \Phi(\epsilon), \tag{22}$$

where

$$\Phi(\epsilon) \leq H_1 \cdot \frac{s_0}{s_0^{(1-2\eta)(3/2+\gamma)}} + H_2 \cdot \frac{r_0}{r_0^{3/2+\gamma}} + H_3 \cdot \epsilon^{\gamma-7\eta}. \tag{23}$$

Note that the positive constants H_1, H_2 and H_3 are determined by (20), and do not depend on the particular choice of $\tilde{\epsilon} > \varepsilon(\gamma)$ or $F > 0$. Using that $\eta = \gamma/100$, we at last choose the threshold $\tilde{\epsilon} \leq \varepsilon(\gamma)$ so that the right hand side of (23) remains smaller than 1 for all $\epsilon < \tilde{\epsilon}$. Invoking (20), and using the induction assumption with this particular choice of $\tilde{\epsilon}$ and $F = f\left(\tilde{\epsilon}\right)$, indeed confirms the inequality (21) for all $\epsilon < \tilde{\epsilon}$. This concludes the proof of Theorem 1.1. $\square$

[27] Since we trivially have $f(\epsilon) \leq f(\epsilon')$ for all $0 < \epsilon/2 \leq \epsilon' \leq \epsilon \leq 1$, it is sufficient to establish the asymptotic bound for the values ϵ of the form $\epsilon = 1/2^j$, where $j \in \mathbb{N}$.

4 Concluding remarks

- Our analysis is largely inspired by the partition-based proof [22] of the Szemerédi-Trotter Theorem [48] on the number of point-line incidences in the plane. In the case at hand, narrow convex sets are viewed as abstract lines, so a non-trivial incidence bound implies that a typical point $u \in P$ is involved in $o(1/\epsilon)$ such canonical sets (which are naturally associated with the surrounding sectors $\mathcal{W}_j(u)$). This gives rise to a sparse restriction graph Π.

  Therefore, it is no surprise that our main decomposition $\Sigma(r_1)$ overly repeats the one that was used by Clarkson *et al.* [22] in order to extend the Szemerédi-Trotter bound to more general settings (e.g., incidences between points and unit circles, and incidences between lines and certain cells in their arrangement).

- As previously mentioned, we did not seek to optimize the implicit constant factor in our asymptotic bound for $f(\epsilon)$. However, it is known to heavily depend on the choice of the constant $\gamma > 0$, and it is at least $2^{\Omega\left(1/\gamma^2\right)}$ for recurrences of the form (20); see the discussion in the end of [47, Section 7.3.2].

- Our proof of Theorem 1.1 is fully constructive, and the resulting net includes points of the following types:

  1. The vertices of the decompositions $\Sigma(r_1)$ which arise in the various recursive instances.
  2. 1-dimensional $\hat{\epsilon}$-nets within lines $L \in \mathcal{Y}(P, r_0)$, for $\hat{\epsilon} = \tilde{\Omega}(\epsilon^{3/2})$. In each net of this kind, the underlying point set is composed of the L-intercepts of the edges of $\binom{P}{2}$. These edges typically belong to one of the sparser graphs $\Pi_{>t}$ (in Section 3.3) or $\Pi(\tau, i)$ (in Section 3.4).
  3. 1-dimensional $\hat{\epsilon}$-nets within lines $L \in \mathcal{Y}(P, r_0)$, for $\hat{\epsilon} = \tilde{\Omega}(\epsilon^{3/2})$, where the underlying point sets are composed of the L-intercepts of the "mixed" edges, which connect the vertices of $\Sigma(r_1)$ to the points of P.
  4. 2-dimensional $\hat{\epsilon}$-nets of Theorem 2.1 with respect to triangles in $\mathbb{R}^2$.

- Our construction and its analysis combine classical elements of the 30-year old theory of linear arrangements in computational geometry (which generalize to any dimension) with a few ad-hoc arguments in $\mathbb{R}^2$ (which do not immediately extend to higher dimensions). A recent study [44] of the author combines the "incidences counting" strategy for the "narrow" convex sets (which loosely resemble hyperplanes) with novel selection-type results [2] in order to show that the comparable bound $f_d(\epsilon) = O\left(1/\epsilon^{d-1/2+\gamma}\right)$, for any $\gamma > 0$, holds in all dimensions $d \geq 3$. Its remarkable that, despite the superficial similarity with the bound of Theorem 1.1, applying the higher-dimensional analysis in the plane comes short of re-establishing Theorem 1.1.

  The author conjectures that the actual asymptotic behaviour of the functions $f_d(\epsilon)$ in any dimension $d \geq 1$ is close to $1/\epsilon$, as is indeed the case for their "strong" counterparts with respect to simply shaped objects in $\mathbb{R}^d$ [30]. Its worth mentioning that our main argument in Section 3.4 exploits the delicate interplay between the two notions of ϵ-nets which was partly explored by Mustafa and Ray [41] and, more recently, by Har-Peled and Jones [29].

- The recently improved analysis of the transversal numbers $C_d(p, q)$ that arise in the Hadwiger-Debrunner problem (see Section 1), due to Keller, Smorodinsky, and Tardos [33], implies that

$$C_d(p,q) \le f_d\left(\Omega\left(p^{-1-\frac{d-1}{q-d}}\right)\right). \qquad (24)$$

Here, as before, $f_d(\epsilon)$ denotes the smallest possible number f with the property that any n-point set $P \subset \mathbb{R}^d$ admits a weak ϵ-net of cardinality f with respect to convex sets.

Plugging the result of Theorem 1.1 into (24) yields an improved bound in dimension $d = 2$:

$$C_2(p,q) = O\left(p^{(3/2+\gamma)\left(1+\frac{1}{q-2}\right)}\right)$$

for any constant parameter $\gamma > 0$, and p larger than a certain constant threshold which depends on γ.

- As the primary focus of this study is on the combinatorial aspects of weak ϵ-nets, we did not seek to optimize the construction cost of our net Q.

  A straightforward implementation of the recursive construction of Q runs in time $\tilde{O}\left(n^2/\sqrt{\epsilon}\right)$. The construction of $\Sigma(\mathcal{R}_1)$ from the sample $\mathcal{R}_1 \subset \mathcal{L}(\Pi)$, the assignment of the points of P to the trapezoidal cells, and the zones of the lines of $\mathcal{L}(\Pi)$, can all be performed using the standard textbook algorithms [40, 47]. Most of the running time is spent on explicitly maintaining the restriction graphs Π along with the sparse graphs of Section 3.4, and tracing the zones of the lines of $\mathcal{L}(\Pi)$ in $\Sigma(r_1)$.

Acknowledgement. The author would like to thank János Pach, Micha Sharir and Gábor Tardos for their numerous invaluable comments on the early versions of this paper. In particular, he is indebted to Gábor Tardos for pointing out that Proposition 3.4 extends to vertical decompositions, which substantially simplified the subsequent exposition.

In addition, the author would like to thank an anonymous Journal of the ACM referee for suggestions which helped to improve the presentation.

References

[1] P. K. Agarwal, and J. Matousek, On range searching with semialgebraic sets,ffix *Discrete Comput. Geom.* 11 (1994), 393–418.

[2] N. Alon, I. Bárány, Z. Füredi and D. J. Kleitman, Point selections and weak ϵ-nets for convex hulls, *Comb. Prob. Comput.* 1 (1992), 189–200.

[3] N. Alon, A non-linear lower bound for planar epsilon-nets, *Discrete Comput. Geom.* 47 (2012), 235–244.

[4] N. Alon and G. Kalai, Bounding the Piercing Number, *Discrete Comput. Geom.* 13 (1995), 245–256.

[5] N. Alon, G. Kalai, R. Meshulam, and J. Matoušek. Transversal numbers for hypergraphs arising in geometry, *Adv. Appl. Math* 29 (2001), 79-–101.

[6] N. Alon, H. Kaplan, G. Nivasch, M. Sharir and S. Smorodinsky, Weak ϵ-nets and interval chains, *J. ACM* 55 (6) (2008), Article 28.

[7] N. Alon and D. J. Kleitman, Piercing convex sets and the Hadwiger-Debrunner (p,q)-problem, *Adv. Math.* 96 (1) (1992), 103–112.

[8] B. Aronov, E. Ezra and M. Sharir, Small-size epsilon nets for axis-parallel rectangles and boxes, *SIAM J. Comput.* 39 (2010), 3248–3282.

[9] B. Aronov, M. Pellegrini and M. Sharir, On the zone of a surface in a hyperplane arrangement, *Discrete Comput. Geom.* 9 (2) (1993), 177 – 186.

[10] J. Balogh, R. Morris and W. Samotij, Independent sets in hypergraphs, *J. AMS*, 28 (2015), 669 – 709.

[11] J. Balogh and J. Solymosi, On the number of points in general position in the plane, *Disc. Analysis* 16 (2018), 1 – 20.

[12] I. Bárány, Z. Füredi and L. Lovász, On the number of halving planes, *Discrete Comput. Geom.* 10 (2) (1990), 175 – 183.

[13] H. Brönnimann and M. T. Goodrich, Almost optimal set covers in finite VC-dimension, *Discrete Comput. Geom.* 14(4) (1995), 463–479.

[14] B. Bukh, J. Matousek, and G. Nivasch, Lower bounds for weak epsilon-nets and stair-convexity, *Israel J. Math.* 182 (2011), 199–228.

[15] Ph. G. Bradford and V. Capoyleas, Weak epsilon-nets for points on a hypersphere, *Discrete Comput. Geom.* 18(1) (1997), 83 – 91.

[16] T. M. Chan, Optimal Partition Trees, *Discrete Comput. Geom.* 47 (2012), 661-–690.

[17] B. Chazelle and E. Welzl, Quasi-optimal range searching in spaces of finite VC-dimension, *Discrete Comput. Geom.* 4 (1989), 467–489.

[18] B. Chazelle, The discrepancy method: randomness and complexity, Cambridge University Press, New York, NY, USA, 2000.

[19] B. Chazelle, H. Edelsbrunner, M. Grigni, L. J. Guibas, M. Sharir and E. Welzl, Improved bounds on weak epsilon-nets for convex sets, *Discrete Comput. Geom.* 13 (1995), 1–15. Also in *Proc. 25th ACM Sympos. Theory Comput. (STOC)*, 1993.

[20] B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica* 10 (3) (1990), 229–249.

[21] K. L. Clarkson, Nearest neighbor queries in metric spaces, *Discrete Comput. Geom.*, 22 (1) (1999), 63–93.

[22] K. L. Clarkson, H. Edelsbrunner, L. J. Guibas, M. Sharir and E. Welzl, Combinatorial complexity bounds for arrangement of curves and spheres, *Discrete Comput. Geom.* 5 (1990), 99–160.

[23] K. L. Clarkson and K. R. Varadarajan, Improved approximation algorithms for geometric set cover, *Discrete Comput. Geom.* 37(1) (2007), 430–58.

[24] G. Even, D. Rawitz and S. Shahar, Hitting sets when the VC-dimension is small, *Inf. Process. Lett.* 95(2) (2005), 358–362.

[25] H. Furstenberg and Y. Katznelson, A density version of the Hales-Jewett Theorem, *J. Analyse Math.*, 57 (1991), 64–119, 1991.

[26] J. E. Goodman, R. Pollack and R. Wenger, Geometric transversal theory, *New Trends in Discrete and Computational Geometry*, J. Pach (Ed.), Springer Verlag, Berlin, 1993, pp. 163–198.

[27] H. Hadwiger and H. Debrunner, Über eine variante zum Hellyschen satz, Archiv der Mathematik, 8(4) (1957), 309 – 313.

[28] D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions, with applications to visbility in terrains, *Discrete Comput. Geom.* 12 (1994), 313—326.

[29] S. Har-Peled and M. Jones, Journey to the center of the point set, in *Proc. 35th Int. Symposium Comput. Geom. (SOCG 2019)*, Article 41.

[30] D. Haussler and E. Welzl, ε-nets and simplex range queries, *Discrete Comput. Geom.* 2 (1987), 127–151.

[31] S. Hoory, N. Linial, and A. Wigderson, Expander graphs and their applications, *Bull. Amer. Math. Soc.* 43 (2006), 439–561.

[32] C. Keller and S. Smorodinsky, A New Lower Bound on Hadwiger-Debrunner Numbers in the Plane, in *Proc. 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA 2020)*, pp. 1155–1169.

[33] C. Keller, S. Smorodinsky and G. Tardos, Improved bounds for Hadwiger-Debrunner numbers, *Israel J. Math.* 225 (2) (2018), 925 – 945. (Also in *Proc. SODA 2017,* pp. 2254 – 2263.)

[34] J. Kolmós, J. Pach and G. J. Woeginger, Almost tight bounds for epsilon-Nets, *Discrete Comput. Geom.* 7 (1992), 163–173.

[35] J. Matoušek, Epsilon-Nets and Computational Geometry, in *New Trends in Discrete Computational Geometry*, J. Pach (Ed.), Algorithms and Combinatorics, Berlin, 1993, pp. 69–89.

[36] J. Matoušek, Lectures on Discrete Geometry. Springer-Verlag, New York, 2002.

[37] J. Matoušek, Efficient partition trees, *Discrete Comput. Geom.* 8 (3) (1992), 315–334.

[38] J. Matousek and U. Wagner, New constructions of weak epsilon-nets, *Discrete Comput. Geom.* 32 (2) (2004), 195–206.

[39] J. Matousek, R. Seidel and E. Welzl, How to net a lot with little: small epsilon-nets for disks and halfspaces, *Proc. 6th ACM Symp. Comput. Geom.*, 1990, pp. 16–22.

[40] K. Mulmuley, Computational Geometry: An Introduction Through Randomized Algorithms, 1st edition, Prentice Hall, 1993.

[41] N. H. Mustafa and S. Ray, Weak ε-nets have a basis of size $O(1/\varepsilon \log 1/\varepsilon)$, *Comput. Geom.* 40 (2008), 84 – 91.

[42] N. H. Mustafa and K. Varadarajan, Epsilon-approximations and epsilon-nets, Chapter 47 in Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (ed.), 3rd edition, CRC Press, Boca Raton, FL, 2017.

[43] J. Pach and G. Tardos, Tight lower bounds for the size of epsilon-nets, *J. AMS* 26 (2013), 645 – 658.

[44] N. Rubin, Stronger bounds for weak epsilon-nets, in *Proc. 53rd Annual ACM SIGACT Symp. Theory Comp. (STOC 2021)*, pp. 989 – 1002.

[45] D. Saxton and A. Thomason, Hypergraph containers, *Inventiones Mathematicae* 201 (3) (2015), 925 – 992.

[46] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* 12 (1994), 327–345.

[47] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

[48] E. Szemerédi and W. T. Trotter, Extremal problems in discrete geometry, Combinatorica 3 (3-4) (1983), 381–392.

[49] V. N. Vapnik and A. Y. Chervonenkis, On the uniform convergence of relative frequencies of events to their probabilities, *Theory Prob. Appls.* 16 (1971), 264–280.

[50] R. Wenger, Helly-type theorems and geometric transversals, in *Handbook of Discrete and Computational Geometry*, 2nd Edition (J.E. Goodman and J. O'Rourke, Eds.), Chapman & Hall/CRC Press, 2004, pp. 73–96.