2018 IEEE 59th Annual Symposium on Foundations of Computer Science

An Improved Bound for Weak Epsilon-Nets in the Plane

Natan Rubin
Department of Computer Science
Ben-Gurion University of the Negev
Beer-Sheva, Israel
Email: rubinnat.ac@gmail.com

***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, 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.**

***Keywords*-epsilon-nets; convex sets; piercing numbers; transversals; VC-dimension; arrangements of lines**

I. 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 [39], in the context of Statistical Learning Theory. Following a seminal paper of Haussler and Welzl [22], ϵ-nets play a central role in Discrete and Computational Geometry [27]. For example, bounds on ϵ-nets determine the performance of the best-known algorithms for Minimum Hitting Set/Set Cover Problem in geometric hypergraphs [6], [9], [18], [19], and the transversal numbers of families of convex sets [2], [3], [5], [24], [26].

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 $\mathcal{K}$ [5], [2], [19].

Haussler and Welzl [22] proved the existence of strong ϵ-nets of cardinality $O\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ for families of simply-shaped, or semi-algebraic geometric ranges in d-space, for a fixed $d > 0$ (e.g., boxes, spheres, halfspaces, simplices), by observing that their induced hypergraphs have a bounded Vapnik-Chervonenkis dimension (so called *VC-dimension*). While the bound is generally tight for a fixed VC-dimension [25], slightly better constructions are known for several special families of ranges, including tight bounds for discs in $\mathbb{R}^2$, halfplanes in $\mathbb{R}^2$ and halfspaces in $\mathbb{R}^3$ [18], [25], [31], and rectangles in $\mathbb{R}^2$ and boxes in $\mathbb{R}^3$ [6], [34]. We refer the reader to a recent state-of-the-art survey [33] for the best known bounds.

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. This phenomenon can be attributed to the fact that such families of ranges determine hypergraphs of unbounded VC-dimension. Nevertheless, Bárány, Füredi and Lovász [8] observed in 1990 that families of convex sets in $\mathbb{R}^2$ still admit weak ϵ-nets of cardinality $O(\epsilon^{-1026})$. Alon *et al.* [1] were the first to show in 1992 that families of convex sets in any dimension $d \geq 1$ admit weak ϵ-nets 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 [28, Sections 8 – 10] for a comprehensive introduction. For example, Alon and Kleitman [5] used the boundedness of weak ϵ-nets to prove Hadwiger-Debrunner (p, q)-conjecture, which concerns transversal numbers of convex sets in $\mathbb{R}^d$.

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 $d \geq 2$. Alon *et al.* [1] (see also [5]) used Tverberg-type results to show that $f_d(\epsilon) = O(1/\epsilon^{d+1-1/s_d})$ (where $0 < s_d < 1$ is a selection ratio which is fixed for every d), and $f_2(\epsilon) = O\left(1/\epsilon^2\right)$. The bound in higher dimensions $d \geq 3$ has been subsequently improved in 1993 by Chazelle *et al.* [14] to roughly $1/\epsilon^d$. Though the latter construction was somewhat simplified in 2004 by Matoušek and Wagner [30] using simplicial partitions with

2575-8454/18/$31.00 ©2018 IEEE
DOI 10.1109/FOCS.2018.00030

IEEE computer society

low hyperplane-crossing number [29], no improvements in the upper bound for general families of convex sets and arbitrary finite point sets occurred for the last 25 years.

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 [10], 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 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 [14], [11], [4]. For example, Alon *et al.* [4] showed in 2008 that any finite point set in a convex position in $\mathbb{R}^2$ allows for a weak ϵ-net of cardinality $O(\alpha(\epsilon)/\epsilon)$ with respect to convex sets.

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

Theorem I.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 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* $O\left(\epsilon^{-3/2-\gamma}\right)$*; here* $\gamma > 0$ *is an arbitrary small constant which does not depend on* ϵ*.*[1]

The rest of the paper is organized as follows. In Section II 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 I.1. In Section III we use the recursive framework of Section II to give a constructive proof of Theorem I.1. In Section IV we briefly summarize the properties of our construction and survey the future lines of work. Due to the lack of space in this extended abstract, many technical details are relegated to the full version [35].

II. Preliminaries

A. Proof outline

We briefly outline the main ideas behind our proof of Theorem I.1. We begin by sketching the $O\left(1/\epsilon^2\right)$ planar construction of Alon *et al.* [1] (or, rather, its more comprehensive presentation by Chazelle [13]).

The quadratic construction. Refer to Figure 1 (top). We split the underlying point set P by a vertical median line 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. If at least $3\epsilon n/4$ points of P lie to the same side of P, we pierce K by one of the auxiliary $(4\epsilon/3)$-nets. Otherwise, the points of $P_K := P \cap K$ span at least $\epsilon^2 n^2/16$ edges that cross $K \cap L$, so we can pierce P by adding to our net each $(\epsilon^2 n^2/16)$-th crossing point of L with the edges of $\binom{P}{2}$.[2]

[1] The constant of proportionality within $O(\cdot)$ may heavily depend on γ.

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

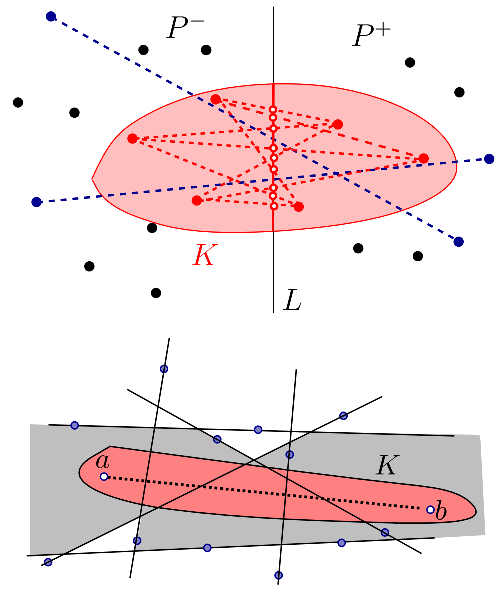

Figure 1. Top: Constructing the net of cardinality $O(1/\epsilon^2)$. If the points of $P_K = P \cap K$ are well distributed between P^- and P^+, the intercept $K \cap L$ is crossed by $\Omega(\epsilon^2 n^2)$ edges of $\binom{P_K}{2}$. Notice that the intercept $K \cap L$ can be crossed by many edges outside $\binom{P_K}{2}$. Bottom: 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 "proxy" edge ab.

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.

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* $\mathcal{A}(\mathcal{R})$

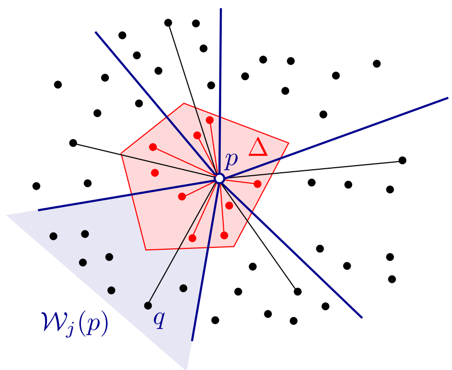

Figure 2. We partition the plane into $O(1/\epsilon)$ sectors $\mathcal{W}_j(p)$, each containing roughly ϵn outgoing edges pq, and an average amount of $O(\epsilon n/r^2)$ outgoing short edges.

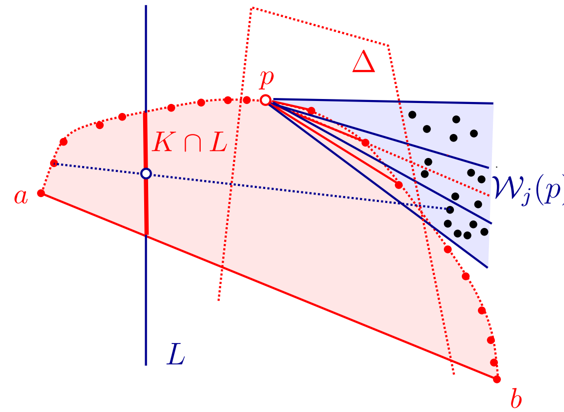

Figure 3. The point p with the outgoing short edges that are "parallel" to ab, and whose supporting lines are roughly tangent to K. In Case 1, the $\Omega(\epsilon n/r)$ outgoing short edges of p within $\Delta \cap K$ occupy multiple sectors $\mathcal{W}_j(p)$ which are almost tangent to K. This yields $\omega(\epsilon^2 n^2)$ segments that cross the intercept $K \cap L$.

– the decomposition of $\mathbb{R}^2 \setminus \bigcup \mathcal{R}$ into open 2-dimensional faces. (See Section II-C 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) is contained, to a large extent, in the zone of a single edge $ab \in \binom{P_K}{2}$; furthermore, there exist $\Omega(\epsilon^2 n^2)$ such "proxy" edges ab in $\binom{P_K}{2}$. See Figure 1 (bottom). In what follows, we refer to such sets as *narrow*.

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 out into exponentially many subsets P_K, so the standard divide-and-conquer schemes [17] hardly apply in this setting. Fortunately, every narrow convex set K can be largely described by its "proxy" edges $ab \in \binom{P_K}{2}$. (For example, K cannot include points outside the respective zones of these edges.)

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 $p \in P_K$, for a narrow convex set K, is generally lower than that of the entire set P from that same point.[3]

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 $pq \in \binom{P}{2}$ is *short* if both of its endpoints lie in the same cell of $\mathcal{A}(\mathcal{R})$.

For each point p of P we partition the surrounding plane into $z = \Theta\left(\frac{1}{\epsilon}\right)$ sectors $\mathcal{W}_1(p), \mathcal{W}_2(p), \ldots, \mathcal{W}_z(p)$ so that each sector encompasses $\Theta(\epsilon n)$ outgoing edges $pq \in \binom{P}{2}$; see Figure 2.

To pierce a narrow convex set K whose zone in $\mathcal{A}(\mathcal{R})$ is traced by an edge $ab \in \binom{P_K}{2}$, we combine the following key observations:

1. For an *average edge* $pq \in \binom{P}{2}$, the respective sector $\mathcal{W}_j(p)$ contains only $O\left(\epsilon n/r^2\right)$ short edges.
2. For an *average point* p in P_K, its cell Δ contains at least $\epsilon n/r$ points of P_K, which are connected to p by short edges (because K crosses at most $r+1$ cells of $\mathcal{A}(\mathcal{R})$).

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 [22], 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 $p \in P_K$, we can assume that both endpoints of the "proxy" edge $ab \in \binom{P}{2}$ of K lie outside the cell Δ of p, and at least half of the $\Omega(\epsilon n/r)$ points $q \in P_K \setminus \{p\}$ within Δ lie to the same side of ab as p. By the near convex position of P_K, most lines spanned by such short edges pq within Δ are roughly tangent to the convex hull of P_K; see Figure 3. (In particular, the four points a, p, q, b form a convex quadrilateral.)

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

Case 1. The wedge spanned by the above $\Omega(\epsilon n/r)$ short

[3] To this end, we define the pseudo-distance between a pair of points $p, q \in \mathbb{R}^2$ as the number of lines in $\mathcal{R}$ that are crossed by the open segment pq; see [13, Section 2.8] and [12]. For a finite set $A \subset \mathbb{R}^2$, and a point $p \in A$, we examine the "expected" order of magnitude of the volume $|A \cap B(p,t)|$ of the ball $B(p,t)$ as a function of t. Clearly, this informal notion is related to the more standard concepts of doubling dimension [16] and graph expansion [23].

edges $pq \in \binom{P_K}{2}$ (along with pb) occupies r "average" sectors $\mathcal{W}_j(p), \mathcal{W}_{j+1}(p), \ldots, \mathcal{W}_{j+r}(p)$, which are almost tangent to K. We show that the points P within $\mathcal{W}_j(p) \cup \mathcal{W}_{j+1}(p) \cup \ldots \cup \mathcal{W}_{j+r}(p)$ yield $r\epsilon^2 n^2$ 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 3.) 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 p within $\binom{P_K}{2}$ that are parallel to ab in the above sense and occupy a constant number of *rich* sectors $\mathcal{W}_j(p)$ with at least $\Omega(\epsilon n/r)$ short edges in each sector.

The first property implies that exist $O\left(1/(r\epsilon)\right)$ rich sectors $\mathcal{W}_j(p)$, which encompass a total of $O(n/r)$ edges that emanate from p. 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 pq which lie in rich sectors $\mathcal{W}_j(p)$, $\mathcal{W}_{j'}(q)$ (for at least one of the respective endpoints p or q). 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. 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, any trapezoidal cell τ is crossed by $O(n^2 \log r/r)$ of the lines that support the edges of $\binom{P}{2}$,[4] 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 any of its $\Omega(\epsilon^2 n^2)$ proxy edges) 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 $\Omega(\epsilon^2 n^2)$ of their edges, their supporting lines cross too many cells Σ and, thereby, belong to another sparse subset of $\binom{P}{2}$.

B. The recursive framework.

We refine the notation of Section I and lay down the formal framework in which our analysis is cast.

Definition. We then say that $Q \subset \mathbb{R}^2$ is *a weak ϵ-net* for a family $\mathcal{G}$ of convex sets $\mathcal{G}$ in $\mathbb{R}^2$ if it pierces every set in $\mathcal{G}$ that is ϵ-heavy with respect to P.

Notice that the previous constructions [4], [14], [30] 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 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" many edges of the (typically, sparse) 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 $\sigma \binom{\epsilon n}{2}$ edges. 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 with respect to a point set P and a certain family $\mathcal{K} = \mathcal{K}(P, \Pi, \epsilon, \sigma)$ of convex sets which is determined by $\epsilon > 0$, a finite point set $P \subset \mathbb{R}^2$, a set of edges $\Pi \subseteq \binom{P}{2}$, and a threshold $0 < \sigma$. 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 the recurrence involves $\Pi = \binom{P}{2}$ and $\sigma = 1$. Each sub-sequent sub-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 $\lambda := |\Pi|/\binom{n}{2}$ is accompanied only by a comparatively mild decrease in the restriction threshold σ, which is bounded from below by a certain positive constant throughout the recurrence. [5]

The above recurrence bottoms out when either ϵ surpasses a certain (suitably small) constant $0 < \tilde{\epsilon} < 1$, or the density λ of the restriction graph falls below ϵ. In the former case we can use the $O\left((1/\tilde{\epsilon})^2\right) = O(1)$ bound of Alon *et al.* [1], and in the latter we resort to a much simpler sub-recurrence which is effectively near-linear in $1/\epsilon$.

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 $\Pi \subset \binom{P}{2}$ of density $\lambda \leq |\Pi|/\binom{n}{2}$, 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

[4] In other words, Σ is a $\Theta(r/\log r)$-cutting [15] of $\mathbb{R}^2$ with respect to these lines.

[5] The preceding discussion in Section II-A, which we formalize below in Lemma II.4, implies then that $f(\epsilon, \sigma, \delta) = o\left(1/\epsilon^2\right)$ once the density λ falls substantially below 1 (given that the restriction threshold σ remains close to 1). Hence, our further recurrence over Π is used to merely amplify this gain.

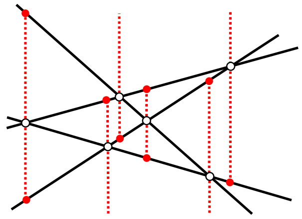

Figure 4. The trapezoidal decomposition $\Sigma(\mathcal{L})$.

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.

C. Geometric essentials: Arrangements and strong ϵ-nets

Strong ϵ-nets. 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 Strong Epsilon-Net Theorem of Haussler and Welzl [22] implies the following result.

Theorem II.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 only $O\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ points of P.*

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.[6]

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 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 4. 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, including at most 2 vertical edges, and the at most 2 other edges that are contained in non-vertical lines of $\mathcal{L}$.

[6]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.

Theorem II.2. *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 II.2 is based on the Epsilon-Net Theorem [22], and it can be found, e.g., in [13]. An easy consequence of Theorem II.2 is that, with probability at least $1/2$, every 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 an $\left(\frac{4C\log r}{r}\right)$-cutting with respect to $\mathcal{L}$.

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.

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. Given 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 each integer $r > 0$ we construct a collection $\mathcal{Y}(r)$ of r vertical (i.e., y-parallel) lines so that every vertical slab of the arrangement $\mathcal{A}(\mathcal{Y}(r))$ contains between $\lfloor n/(r+1)\rfloor$ to $\lceil n/(r+1)\rceil$ points of the underlying n-point set P, and no line of $\mathcal{Y}(r)$ passes through a point of P.

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

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

The following main property of the decompositions $\Lambda(r)$ is used throughout our proof of Theorem I.1.

Lemma II.3. *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 of $O\left(r \cdot f\left(\epsilon' \cdot r\right)\right)$ points that pierce every convex set K that is ϵ'-crowded in $\Lambda(r)$.*

Proof of Lemma II.3: Assume with no loss of generality that $r > 2n$, for otherwise our net consists of P. Recall that each slab $\tau \in \Lambda(r)$ cuts out a subset $P_\tau := P \cap \tau$ of cardinality $n_\tau := |P_\tau| \leq \lceil n/r \rceil = \Theta(n/r)$.

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 $\Lambda(r)$. (In particular, we can further assume that $\epsilon' =$

$O(1/r)$.) For each slab $\tau \in \mathcal{A}(\mathcal{R})$ 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 these nets Q_τ is indeed $O(r \cdot f(r \cdot \epsilon'))$.[7] ■

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).

Lemma II.4. *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*

$$O\left(r \cdot f(\epsilon \cdot \sigma \cdot r) + \frac{r^2 |\Pi|}{\sigma \epsilon^2 n^2}\right). \quad (1)$$

Proof: Assume with no loss of generality that $|P| \geq 2r$, for otherwise the claim follows trivially. We consider the slab decomposition $\Lambda(r)$ and apply Lemma II.3 with $\epsilon' = \sigma\epsilon/4$ to obtain a net Q' of size $O(r \cdot f(\epsilon \cdot \sigma \cdot r))$ that pierces every set $K \in \mathcal{K}$ that is ϵ'-crowded in $\Lambda(r)$.

In addition, for each vertical line $L \in \mathcal{Y}(r)$ we construct an auxiliary net Q_L by choosing every $\lfloor \sigma \binom{\epsilon n}{2}/(2r) \rfloor$-th crossing point of L with the edges of Π.[8] Notice that

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

It suffices to check that every convex set $K \in \mathcal{K}$ is stabbed by 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 $\mathcal{Y}(r)$, we find a point $p \in P_K$ so that at least $2\sigma\binom{\epsilon n}{2}/\lceil \epsilon n \rceil \geq \sigma\epsilon n/4$ of its neighbors in the graph (P_K, Π_K) lie in the same slab $\tau \in \Lambda(r)$ that contains p. Hence, K is ϵ'-crowded in $\Lambda(r)$ 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 $\mathcal{Y}(r)$. Since there are at least $(\sigma/2)\binom{\epsilon n}{2}$ intersection points between the edges of Π_K and the lines of $\mathcal{Y}(r)$, there must be a line $L \in \mathcal{Y}(r)$ which contains at least $\sigma\binom{\epsilon n}{2}/(2r)$ of these intersections. Hence, K is hit by the net Q_L. ■

In what follows, σ remains bounded from below by a certain positive constant, and we apply Lemma II.4 with r that is a very small (albeit, fixed) constant power $1/\epsilon$. (In particular, r is much larger than $1/\sigma$.) As a result, the recursive term on the right side of (1) is essentially linear in $1/\epsilon$. 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 $\lambda = O(\epsilon)$. A standard inductive approach to solving recurrences of this kind is presented, e.g., in [20] and [37, Section 7.3.2].

[7] To simplify the presentation, we routinely omit constant factors within the arguments of the recursive terms of the form $f(\epsilon \cdot r)$ and $f(\epsilon, \lambda/r, \sigma)$ as long as these constants do not depend on r.

[8] If $\lfloor \sigma\binom{\epsilon n}{2}/(2r) \rfloor = 0$, then no point is chosen from L.

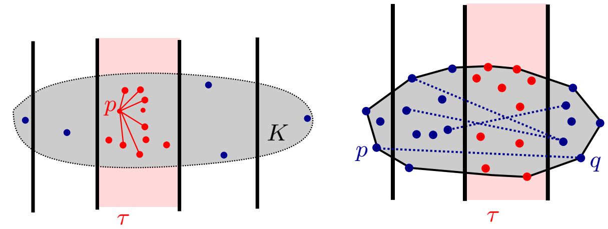

Figure 5. Left: Proof of Lemma II.4. If most edges of Π_K do not cross any line of $\mathcal{Y}(r)$, there is a slab that contains $\Omega(\sigma\epsilon n)$ points of P_K. Right: The slab $\tau \in \Lambda(r_0)$ is a middle slab for K. The depicted edge $pq \in \Pi_K$ crosses τ transversally.

III. Proof of Theorem I.1

To establish Theorem I.1, we derive a recursive bound for the quantity $f(\epsilon, \lambda, \sigma)$ (defined in Section II-B) which implies that $f(\epsilon) = f(\epsilon, 1, 1) = O\left(1/\epsilon^{3/2+\gamma}\right)$, for any $\gamma > 0$. To this end, we fix the family $\mathcal{K} := \mathcal{K}(P, \Pi, \epsilon, \sigma)$ for arbitrary $P \subset \mathbb{R}^2$, $0 \leq \epsilon, \sigma \leq 1$, and $\Pi \subset \binom{P}{2}$. We then bound the piercing number of $\mathcal{K}$ in terms of the simpler quantities $f(\epsilon, \lambda', \sigma')$ and $f(\epsilon)$, for $\lambda' < \lambda$ and $\epsilon' > \epsilon$.

Throughout our analysis, the restriction threshold σ is bounded from below by an absolute positive constant which does not depend on ϵ and λ. We can also assume that ϵ is bounded from above by a sufficiently small absolute constant $\tilde{\epsilon} > 0$; otherwise, we can use the previous $O\left(1/\tilde{\epsilon}^2\right) = O(1)$ bound [1]. In addition, we can assume that $|P| \geq 1/\epsilon$; otherwise our transversal consists of P. For most of this section we also assume that $|\Pi| \geq \epsilon\binom{n}{2} = \Omega(1/\epsilon)$ (or, else, Lemma II.4 would provide a much simpler recurrence (1), which is essentially linear in $1/\epsilon$).

To bound the piercing number of $\mathcal{K}$, 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 III-B 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 $r_1 = \Theta\left(\sqrt{1/\epsilon}\right)$. 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 II-A) 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 $\Sigma(r_1)$. To this end, we use the auxiliary slab decomposition $\Lambda(r_0)$ of Lemma II.3, with a suitable $r_0 \gg 1/\sigma$, in combination with Theorem II.1.[9]

A. Stage 0: The strip decomposition $\Lambda(r_0)$

At this stage we construct an auxiliary, almost constant-size slab decomposition $\Lambda(r_0)$ and use Lemma II.3 to guarantee for each convex set $K \in \mathcal{K}$ that the points of P_K are sufficiently spread among the slabs of $\Lambda(r_0)$. To this end, we fix r_0 to be an arbitrary small (albeit, fixed) positive degree of $1/\epsilon$, and select a set $\mathcal{Y}(r_0)$ of vertical lines as detailed in Section II-C.

Let $0 < C_0 < 1/4$ be a sufficiently small absolute constant which does not depend on σ. By Lemma II.3, we can pierce (and subsequently remove from $\mathcal{K}$) every ϵ'-crowded convex set K, for $\epsilon' = C_0\sigma\epsilon$, using an auxiliary net Q_0 of cardinality[10] $|Q_0| = O\left(r_0 \cdot f\left(\epsilon \cdot \sigma \cdot r_0\right)\right)$.

Definition. Let τ be a cell in an arrangement of lines. The edge pq crosses τ *transversally* if pq intersects the interior of τ, and none of p, q lies in τ; see Figure 5 (right).

Denote

$$\epsilon_0 := \sigma\epsilon/(100 r_0). \tag{2}$$

Let $K \in \mathcal{K}$ be a convex set. We say that a slab $\tau \in \Lambda(r_0)$ is a *middle slab* with respect to K if it satisfies the following conditions:

(M1) $\epsilon_0 n \leq |P_K \cap \tau| \leq C_0\sigma\epsilon n \leq \epsilon n/4$, and

(M2) $\Omega\left(\sigma\epsilon^2 n^2/r_0\right)$ of the edges of $\Pi_K = \Pi \cap \binom{P}{2}$ cross τ transversally.

Proposition III.1. *With the previous choice of* $0 < \epsilon < 1$, $0 < \sigma \leq 1$, *and* P, *and an arbitrary edge set* $\Pi \subseteq \binom{P}{2}$, *the following property holds: For every convex set* $K \in \mathcal{K}\left(P, \Pi, \epsilon, \sigma/2\right)$ *that is missed by the previously defined net* Q_0, *there is at least one middle slab in* $\Lambda(r_0)$.

Proof: Since K is not ϵ'-crowded, it easily follows that the majority of the edges of Π_K must cross transversally at least one slab of $\Lambda(r_0)$. In the full version [35], we use a simple pigeonhole argument to show that there is such a slab τ in $\Lambda(r_0)$ that satisfies both conditions (M1) and (M2). ∎

B. 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, for $r_1 \gg r_0$. Since $\Sigma(r_1)$ is a refinement of the auxiliary slab decomposition $\Lambda(r_0)$, we can use the properties of $\Lambda(r_0)$ to show that the points of P_K are sufficiently spread in the finer decomposition $\Sigma(r_1)$.

The decomposition $\Sigma(r_1)$. We select the parameter r_1 so that $1/\sigma \ll r_0 \ll r_1 = \Theta\left(\sqrt{1/\epsilon}\right)$ and sample a subset $\mathcal{R}_1$ of r_1 lines from $\mathcal{L} = \mathcal{L}(\Pi)$. We can assume with no loss of generality that no line of $\mathcal{Y}(r_0)$ passes through a vertex of $\mathcal{A}(\mathcal{R}_1)$.[11]

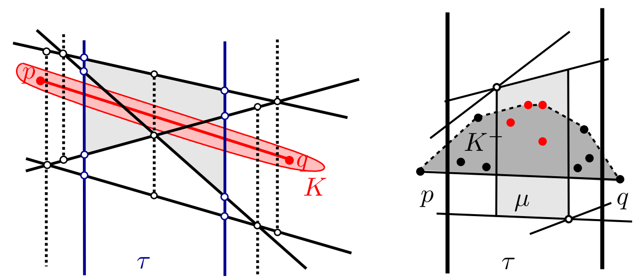

Figure 6. Left: The set K is narrow in Σ_τ because $K \cap \tau$ is contained in the zone of every segment $pq \subset K$ that crosses τ transversally. (The cells of the zone of $K \cap \tau$ within Σ_τ are shaded.) Right: The set $K^+ = \mathsf{conv}(P_K^+ \cup \{p, q\})$, and the point set $P_K(\mu)$.

To simplify the exposition, we add the lines of $\mathcal{Y}(r_0)$ to $\mathcal{R}_1$, so the arrangement $\mathcal{A}(\mathcal{R}_1)$ is a refinement of $\Lambda(r_0)$. We then construct the trapezoidal decomposition $\Sigma(\mathcal{R}_1)$ of $\mathcal{A}(\mathcal{R})$ which was described in Section II-C; see Figure 6 (left). We further subdivide each cell $\hat{\mu} \in \Sigma(\mathcal{R})$ (where necessary) into sub-trapezoids μ so that $|P \cap \mu| \leq n/r_1^2$; this can be achieved using at most $\lfloor r_1^2 |P \cap \hat{\mu}|/n \rfloor$ additional vertical walls.

A standard calculation 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 $\Sigma(\mathcal{R}_1)$, 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 $\mathcal{L} = \mathcal{L}(\Pi)$.

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 $\Sigma(r_1)$. More precisely, we have the following property.

For $t \geq 1$ let $\mathcal{L}(t)$ be the subset of all the lines in $\mathcal{L}$ that cross more than $t r_1 \log r_1$ cells of $\Sigma(r_1)$.[12] A simple double counting argument yields the following bound.

Proposition III.2. *We have* $|\mathcal{L}(t)| = O\left(\frac{m}{t}\right)$.

Let Π_t be the set of edges that span the lines of $\mathcal{L}(t)$. We fix a sufficiently large parameter t, so that $r_1 \gg t \gg r_0$. Consider the recursive instance $\mathcal{K}(t) := \mathcal{K}(P, \Pi_t, \epsilon, \sigma/2)$.

Using the bound of Proposition III.2 on $|\Pi_t| = |\mathcal{L}(t)|$, we can pierce the sets of $\mathcal{K}(t)$ by an auxiliary net Q_2 of size

$$|Q_1| = f\left(\epsilon, \frac{|\Pi_t|}{\binom{n}{2}}, \frac{\sigma}{2}\right) \leq f\left(\epsilon, \frac{m}{t\binom{n}{2}}, \frac{\sigma}{2}\right) = f\left(\epsilon, \frac{\lambda}{t}, \frac{\sigma}{2}\right).$$

[9] For $x, y \geq 1$, the notation $x \ll y$ means that $x = O\left(y^\eta\right)$. Here $\eta > 0$ is an arbitrary small but constant parameter to be fixed in the sequel, and the constants hidden by the $O(\cdot)$-notation do not depend on x and y. (For $0 < x, y \leq 1$, the notation $x \ll y$ means that $1/y \ll 1/x$.)

[10] Throughout our recurrence, σ remains bounded from below by an absolute positive constant. In the sequel, we choose $r_0 \gg 1/\sigma$ to be an arbitrary small constant positive power of $1/\epsilon$. The constants of proportionality hidden by the $O(\cdot)$-notation do not depend on σ.

[11] 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 $\mathcal{R}_1$.

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

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 t to be a very small (albeit, constant) positive power of $1/\epsilon$ guarantees that the recurrence on the density $\lambda = m/\binom{n}{2}$ 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 $\mathcal{K}(P, \Pi \setminus \Pi(t), \epsilon, \sigma/2)$. We thus remove the edges of $\Pi(t)$ from Π. 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).

For every remaining set $K \in \mathcal{K}$ that is missed by the auxiliary net Q_1, the induced edge set $\Pi_K = \Pi \cap \binom{P_K}{2}$ still contains at least $(\sigma/2)\binom{\epsilon n}{2}$ edges, each of them crossing at most $tr_1 \log r_1$ cells of the decomposition $\Sigma(r_1)$ (again, see Figure 6 (left)). In the following Section III-C we use this property to guarantee that every set $K \in \mathcal{K}$ intersects at most $tr_1 \log r_1$ 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.

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

For each slab $\tau \in \Lambda(r_0)$ we consider the subfamily $\mathcal{K}_\tau$ of all the convex sets $K \in \mathcal{K}$ so that τ is their middle slab. By Proposition III.1 (and since every remaining set $K \in \mathcal{K}$ is missed by the net Q_0 of Stage 0), we have $\mathcal{K} = \bigcup_{\tau \in \Lambda(r_0)} \mathcal{K}_\tau$. Notice that a single convex set K can belong to several such sub-families $\mathcal{K}_\tau$.

In Section III-D, we 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 $\tau \in \Lambda(r_0)$ we consider the restriction $\Sigma_\tau := \{\mu \in \Sigma(r_1) \mid \mu \subset \tau\}$.

Definition. We say that a convex set $K \in \mathcal{K}_\tau$ is *narrow in* Σ_τ if for *every* segment $pq \subset K$ that crosses τ transversally, the restriction $K \cap \tau$ is contained in the zone of pq within Σ_τ. (Notice that the cells of this zone lie in the zone of pq within the arrangement $\mathcal{A}(\mathcal{R}_1)$.) See Figure 6 (right).

Informally, the Σ_τ-narrowness of K means that its behaviour is "line-like" in Σ_τ, so the zone of K in Σ_τ can be completely "read off" from any edge of Π_K that crosses τ transversally. As a result, we obtain the following property.

Proposition III.3. *Let τ be a slab of $\Lambda(r_0)$ and K be a set of $\mathcal{K}_\tau$ that is narrow in Σ_τ. Then K intersects at most $tr_1 \log r_1$ cells of Σ_τ.*

In the full version [35] we use the structure of $\mathcal{A}(\mathcal{R}_1)$ and $\Sigma(r_1)$ to get rid of the sets $K \in \mathcal{K}_\tau$ that are not narrow in Σ_τ, by establishing the following property.

Proposition III.4. *With the previous definitions, there is a set Q_2 of cardinality $O\left(r_0^2 r_1/\epsilon\right)$ points that, for each slab $\tau \in \Lambda(r_0)$, pierce every convex set $K \in \mathcal{K}_\tau$ that is not narrow in Σ_τ.*

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 $\tau \in \Lambda(r_0)$, every remaining set of $\mathcal{K}_\tau$ is narrow in Σ_τ.

Combing the bound $|Q_2| = O\left(r_0^2 r_1/\epsilon\right)$ of Proposition III.4 with the bounds on respective cardinalities of 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

$$f\left(\epsilon, \frac{\lambda}{t}, \frac{\sigma}{2}\right) + O\left(r_0 \cdot f(\epsilon \cdot \sigma \cdot r_0) + \frac{r_0^2 r_1}{\epsilon}\right) \qquad (3)$$

points to the net Q. As previously mentioned, choosing t to be a very small (albeit, constant) positive power of $1/\epsilon$ guarantees that our recurrence (3) 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 (3) to the cardinality of Q is effectively dominated by the non-recursive term, which is roughly bounded by $1/\epsilon^{3/2}$ for $r_0 \ll r_1 = \Theta(\sqrt{1/\epsilon})$.

D. 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 $\mathcal{K}(P, \Pi, \epsilon, \sigma)$ by building a "local" net Q_τ for each family $\mathcal{K}_\tau$, which remains fixed for most of this section. To this end, we implement the paradigm of Section II-A within each slab τ of $\Lambda(r_0)$.

The setup. By definition, the slab $\tau \in \Lambda(r_0)$ is a middle slab for each convex set $K \in \mathcal{K}_\tau$. Namely, we have $|P_K \cap \tau| \geq \epsilon_0 n$ and the graph Π_K contains $\Omega\left(\frac{\sigma}{r_0}\binom{\epsilon n}{2}\right)$ edges pq that cross τ transversally. In addition, we are given a sub-partition Σ_τ of τ into trapezoidal cells so that each of these cells contains at most n/r_1^2 points of P. We also assume that each $K \in \mathcal{K}_\tau$ is narrow in Σ_τ. Hence, the zone of K in Σ_τ is composed of at most $tr_1 \log r_1$ cells; all of these cells are intersected by each of the above "witness" edges $pq \in \Pi_K$ that cross τ transversally.

Let P_τ^+ be this portion of $P_\tau = P \cap \tau$ above the line $L_{p,q}$ and put $P_K^+ = P_K \cap P_\tau^+$. 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 6 (right).

Definition. We set $\epsilon_1 := \dfrac{\epsilon_0}{80 \log 1/\epsilon}$ and $\hat{\epsilon} := \dfrac{\epsilon_0}{8tr_1 \log r_1}$.

For each cell $\mu \in \Sigma_\tau$ we denote $P_K(\mu) := P_K^+ \cap \mu$ and $g_\mu := |P_K(\mu)|$.

We say that a cell $\mu \in \Sigma_\tau$ is *full* (with respect to K^+) if $g_\mu \geq \hat{\epsilon} n$. By the Pigeonhole Principle, at least $\epsilon_0 n/5 \geq \epsilon_0 n/4 - 2r_1$ points of P_K^+ lie in (the respective interiors of) the full cells of Σ_μ, whose set we denote by Σ_K.

To implement the paradigm of Section II-A for the set $P_\tau = P \cap \tau$, the decomposition Σ_τ of τ, and the convex set K^+, we first guarantee that the points of P_K^+ are in a

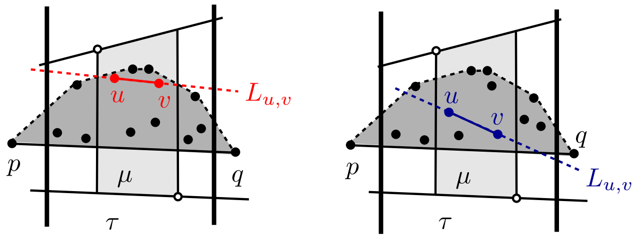

Figure 7. Left: The short edge uv is good for K because all the points of $P_K^+ \cup \{p, q\}$ that lie outside μ are to the same side of $L_{u,v}$. Right: The short edge uv is bad for K.

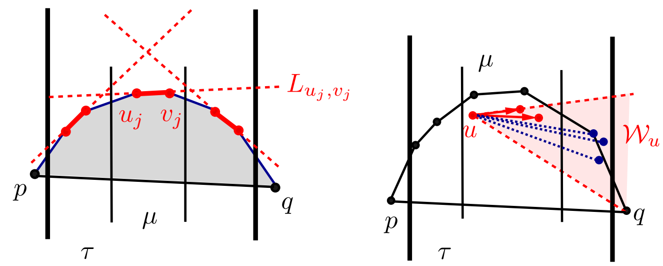

Figure 8. Left: Proposition III.5 (i) – The good edges u_jv_j with supporting lines L_{u_j,v_j} are depicted. Right: The wedge $\mathcal{W}_u$ to the right of u encompasses uq and all the outgoing good edges of u.

sufficiently convex position, and that they are sufficiently spread within Σ_τ.

(i) We construct a finer slab decomposition $\Lambda(s_0)$, where $s_0 \gg r_0$ is again an arbitrary small (albeit, fixed) constant power of $1/\epsilon$. We can assume with no loss of generality that $\mathcal{Y}(s_0) \supset \mathcal{Y}(r_0)$. Since $s_0 \ll r_1 = \Theta\left(\sqrt{1/\epsilon}\right)$, we can add the lines of $\mathcal{Y}(s_0)$ to the sample $\mathcal{R}_1$ with no affect on the asymptotic properties of $\Sigma(r_1)$. We then apply Lemma II.3 to construct a net $Q(s_0)$ that pierces every convex set that is $(\hat{C}\epsilon_1)$-crowded in $\Lambda(s_0)$. Here $\hat{C} > 0$ is a sufficiently small constant to be determined in the sequel. Notice that

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

where the last inequality uses the definition 2 of ϵ_0. Upon adding $Q(s_0)$ to Q, we can assume that each convex set K is $(C\epsilon_1)$-spread in $\Lambda(s_0)$.

(ii) We invoke Theorem II.1 to construct a strong $(\hat{C}\hat{\epsilon})$-net $Q^\triangle(\hat{\epsilon})$ 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

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

We can 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.

We now summarize the key properties of the remaining sets $K \in \mathcal{K}_\tau$, which are not pierced by the auxiliary nets $Q(s_0)$ and $Q^\triangle(\hat{\epsilon})$.

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 $\binom{g_\mu}{2} = \Omega(\hat{\epsilon}^2 n^2)$ 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 7.

Informally, the good edges span lines that are nearly tangent to K.[13] In particular, for every good edge uv the corresponding line $L_{u,v}$ must miss pq. Since uv lies above pq, the edges uv and pq are boundary edges of a convex quadrilateral. In the full version [35] we establish the following property of the remaining convex sets $K \in \mathcal{K}_\tau$, which are not pierced by $Q(s_0)$ and $Q^\triangle(\hat{\epsilon})$.

In the full version [35] we establish the following property.

Proposition III.5. *(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 8 (left).*

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

Definition. Let μ be a cell of Σ_K. We orient every good edge within μ to the right. We say that a point $u \in P_K(\mu)$ is *good* if it is adjacent to at least $g_\mu/10$ outgoing good edges.

The second part of Proposition III.5 implies the following property:

Proposition III.6. *Every full cell $\mu \in \Sigma_K$ contains at least $g_\mu/4$ good points of P_K^+, for a total of at least $\epsilon_0 n/20$ such points.*

Definition. For good point $u \in P_K^+$, let $\mathcal{W}_u$ denote the smallest planar wedge with apex u that contains uq and all the outgoing good edges uv of u (within τ) but does not contain up; see Figure 8 (right). Note that $\mathcal{W}_u$ lies entirely in the halfplane to the right of u. Let D_u denote the cardinality of $(P_\tau \cap \mathcal{W}_u) \setminus \{u\}$, that is, the number of the edges in $uw \in \binom{P_\tau}{2}$ that are adjacent to u and lie within $\mathcal{W}_u \cap \tau$.

Since $\mathcal{W}_u$ encompasses all the outgoing good edges of $u \in P_K(\mu)$, we trivially have $D_u \ge g_\mu/4 \ge \hat{\epsilon}n/10$, and D_u can be much larger than ϵn due to the additional points of $P_\tau \setminus P_K$ that potentially lie within $\mathcal{W}_u \cap \tau$.

[13] 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 witness edge $pq \in \Pi_K$ which crosses τ transversally.

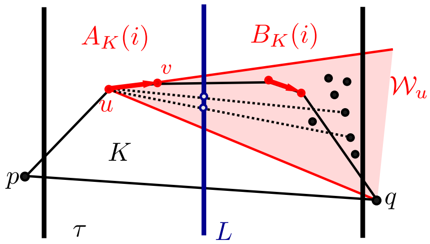

Figure 9. We use a line $L \in \mathcal{Y}(s_0)$ to split the subset $P_K(i)$ of the good points of type i into subsets $A_K(i)$ and $B_K(i)$, of cardinality $\Omega(\epsilon_1 n)$ each. For every point $p \in A_K(i)$, the wedge $\mathcal{W}_u$ contains all the points of $B_K(i)$.

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

Definition. For each i in the interval $I := [\lfloor \log(2\hat{\epsilon}/5\epsilon_1)\rfloor, \log 4/\epsilon_1]$ denote $\delta_i = 2^i\epsilon_1/4$.

We say that a good point $u \in P_K^+$ is of *type* i if $\delta_i n \le D_u < \delta_{i+1}n$. For each $i \in I$ we use $P_K(i)$ to denote the subset of all the i-type good points in P_K^+. Since $D_u \ge \hat{\epsilon}n/10$ holds for any good point, the union $\bigcup_{i\in I} P_K(i)$ covers all the good points in P_K^+.

Since $\hat{\epsilon} = \omega(\epsilon^2)$, we have $|I| \le 4\log 1/\epsilon$. Hence, the pigeonhole principle guarantees that there is $i \in I$ so that $|P_K(i)| \ge \epsilon_0 n/(80\log 1/\epsilon) = \epsilon_1 n$, in which case we say that the set $K \in \mathcal{K}_\tau$ is of *type* i. We keep the type i of our convex sets K (or, rather, their restrictions K^+) fixed during the rest of the analysis, and note that a set may belong to $O(\log 1/\epsilon)$ distinct types.

Since K is $(\hat{C}\epsilon_1)$-spread in $\Lambda(s_0)$, there must be a line $L \in \mathcal{Y}(s_0)$ so that at least $\epsilon_1 n/4$ good points in $P_K(i)$ lie to each side of L. Let $A_K(i)$ (resp., $B_K(i)$) denote the subset of the good points in $P_K(i)$ that lie to the left (resp., right) of L; see Figure 9.

Proposition III.7. *For every point $u \in A_K(i)$ the respective wedge $\mathcal{W}_u$ contains at least $g_\mu/10 \ge \hat{\epsilon}n/10$ outgoing good edges uv, and all the points of $B_K(i)$.*

Since $|B_K(i)| \ge \epsilon_1 n/4$, the proposition implies that $\delta_i \ge \epsilon_1/4$, or $i \ge 0$.

Proof of Proposition III.7: The desired number of the good edges in $\mathcal{W}_u$ follows by the construction of $\mathcal{W}_u$ (and because all the points in $A_K(i)$ are good).

To show that $\mathcal{W}_u$ contains all the points of $B_K(i)$, let $\mu \in \Sigma_K$ be the full cell that contains u. Since every point $w \in B_K(i)$ lies in a cell $\mu' \in \Sigma_K$ to the right of μ and L, the desired property follows by the first part of Proposition III.5 (as each of the points u and w is adjacent to a good edge within the respective cell). ∎

To pierce the remaining sets $K \in \mathcal{K}_\tau$ of type i, in the full version [35] we establish the following lemma.

Lemma III.8. *For each $u \in \mathcal{A}_K(i)$, its respective wedge $\mathcal{W}_u$ contains $\Omega(\delta_i n)$ edges that are adjacent to u and cross L within $K \cap L$ (for a total of $\Omega(\epsilon_1\delta_i n^2)$ such edges that cross $K \cap L$). See Figure 9.*

Lemma III.9. *There is a subset $\Pi(i) \subset \binom{P_\tau}{2}$ with the following properties:*

(i). *$\Pi(i)$ does not depend on the set $K \in \mathcal{K}_\tau$, and has cardinality $|\Pi(i)| = O\left(\delta_i n^2/\left(r_1^2\hat{\epsilon}\right)\right)$.*

(ii). *For each point $u \in A_K(i)$, the set $\Pi(i)$ contains all the edges $uw \in \binom{P_\tau}{2}$ that are adjacent to u and lie within the respective wedge $\mathcal{W}_u$.*

Proof: We partition $\mathbb{R}^2$ into $O(1/\delta_i)$ sectors $\mathcal{W}_j(p)$. 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 short edges in $\mathcal{W}_j(p)$ to $\Pi(i)$ only if this sector encompasses at least $\hat{\epsilon}n/10$ short edges.

To see the first property of $\Pi(i)$, we bound the degree of each point p by $O\left(n/(r_1^2\hat{\epsilon})\right)$. To that end, we recall that for each cell $\mu \in \Sigma_\tau$ we have that $n_\mu = |P_\mu| \le n/r_1^2$. Therefore, for each $p \in P_\mu$ there can be only $O\left(1/(r_1^2\hat{\epsilon})\right)$ sectors $\mathcal{W}_j(p)$ that satisfy $|\mathcal{W}_j(p) \cap P_\mu \setminus \{p\}| \ge \hat{\epsilon}n/10$.

For the second property, we recall that, for every good point $u \in A_K(i)$ that lies in a full cell $\mu \in \Sigma_K$, the respective wedge $\mathcal{W}_u$ contains at most $2\delta_i n$ outgoing edges uw within τ and, therefore, is contained in one of the sectors $\mathcal{W}_j(u)$. Proposition III.7 now implies that this sector $\mathcal{W}_j(u)$ contains at least $g_\mu/10 \ge \hat{\epsilon}n/10$ outgoing short edges uv. Hence, $\Pi(i)$ must include all of the edges in $\mathcal{W}_j(u)$. ∎

Notice that the density of the graph $\Pi(i)$ 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 II-A – combining Lemma III.8 and Lemma III.9 (ii) for each $u \in A_K(i)$ yields that the intercept $K \cap L$ 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 $\Pi(i)$ contains roughly n^2/r_1 edges, which gives rise to the second favourable scenario of Section II-A (e.g., via Lemma II.4, or through a direct application of Lemma III.8).

The net. Our net $Q_\tau(i)$ for the convex sets $K \in \mathcal{K}_\tau$ of type i interpolates between the above extreme cases. For every $i \in I$, and every line $L \in \mathcal{Y}(s_0)$ within τ, we select every intersection of L with an edge of $\Pi(i)$ into the set $X_L(i)$. We then select every $\lceil C'\epsilon_1\delta_i n^2\rceil$-th point of $X_L(i)$ into our net $Q_L(i)$, for a sufficiently small constant $C' > 0$.

We then define $Q_\tau(i) := \bigcup\{Q_L(i) \mid L \in \Lambda(s_0), L \subset \tau\}$. and $Q_\tau := \bigcup_{i\in I} Q_\tau(i)$. Hence, our last net Q_3 at Stage 3 is given by $$Q_3 := Q(s_0) \cup Q^{\triangle}(\hat{\epsilon}) \cup \bigcup_{\tau\in\Lambda(r_0)} Q_\tau.$$

It suffices to check, for all $\tau \in \Lambda(r_0)$ and $i \in I$, that every type-i set $K \in \mathcal{K}_\tau$ is pierced by some net $Q_L(\tau)$ whose line $L \in \Lambda(s_0)$ lies within τ. Indeed, according to Lemma III.8, every point $u \in A_K(i)$ gives rise to $\Omega(\delta_i n)$ outgoing edges that cross the intercept $K \cap L$ with some line $L \in \mathcal{Y}(s_0)$

which separates $A_K(i)$ and $B_K(i)$, for a total of $\Omega(\epsilon_1 \delta_i n^2)$ such edges. Hence, fixing a small enough constant C' (which may depend on $\hat{C}$) guarantees that K is pierced by $Q_L(i)$.

Stage 3 – the analysis. For every type $i \in I$, and every line $L \in \Lambda(s_0)$ within τ, the cardinality of $Q_L(i)$ is bounded by

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

where the first inequality uses the bound of Lemma III.9 (ii), and the second one uses the definitions of ϵ_1 and $\hat{\epsilon}$.

Recall that $\Lambda(s_0)$ is a refinement of $\Lambda(r_0)$, every slab $\tau \in \Lambda(r_0)$ contains $O(s_0/r_0)$ lines of $\mathcal{Y}(s_0)$. Using this and the definition of ϵ_0 in Section III-A, we can bound the cardinality of Q_τ by

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

Repeating this bound for each slab $\tau \in \Lambda(r_0)$ and combining it with the prior bounds on the respective cardinalities of the nets $Q(s_0)$ and $Q^{\triangle}(\hat{\epsilon})$, we conclude that the overall increase in the size of Q at Stage 3 is bounded by

$$O\left(s_0 \cdot f\left(\epsilon \cdot \frac{s_0 \cdot \sigma}{r_0 \log 1/\epsilon}\right) + \frac{t r_0 r_1}{\epsilon \sigma} \log^2 \frac{1}{\epsilon} + \frac{s_0 r_0^2 t \log^3 1/\epsilon}{\sigma^2 r_1 \epsilon^2}\right). \tag{4}$$

Note that s_0 is very small (albeit, fixed) positive power of $1/\epsilon$ that satisfies $s_0 \gg r_0 \gg 1/\sigma$, so the recursive term is again near-linear in $1/\epsilon$. Furthermore, the two non-recursive terms sum up to roughly $r_1/\epsilon + 1/(r_1\epsilon^2) = \Theta(\epsilon^{-3/2})$.

E. The final recurrence

To develop the complete recurrence for the quantity $f_2 = f(\epsilon)$. we combine bounds on the sizes of the auxiliary nets Q_0, Q_1, Q_2 and Q_3 that we constructed in Stages 0 – 3. As mentioned in Section II, we fix a suitably small constant $0 < \tilde{\epsilon} < 1$ and use the old bound $f(\epsilon) = O\left(\tilde{\epsilon}^{-2}\right) = O(1)$ of Alon *et al.* [1] whenever $\epsilon > \tilde{\epsilon}$. Assume then that $\epsilon < \tilde{\epsilon}$.

Bounding $f(\epsilon, \lambda, \sigma)$**.** To obtain the desired recursion for $f(\epsilon) = f(\epsilon, 1, 1)$, we first express $f(\epsilon, \lambda, \sigma)$ in terms of the simpler quantities $f(\epsilon')$, for $\epsilon' > \epsilon$. Assume first that $\lambda > \epsilon$. Combining the bounds in (3) and (4), and substituting $r_1 = \Theta(\sqrt{1/\epsilon})$, we obtain the bound

$$f(\epsilon, \lambda, \sigma) \leq f\left(\epsilon, \tfrac{\lambda}{t}, \tfrac{\sigma}{2}\right) + O\left(s_0 \cdot f\left(\epsilon \cdot s_0 \cdot \tfrac{\sigma}{r_0 \log 1/\epsilon}\right)\right) +$$

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

We begin with $f(\epsilon) = f(\epsilon, 1, 1)$ and recursively apply the last inequality to the "leading" term, which involves the density λ. Notice that σ is initially equal to 1 (and it will remain bounded from below by a fixed positive constant). Hence, all the parameters r_0, t and s_0 can be chosen to be arbitrary small (albeit, constant) positive powers of $1/\epsilon$ that satisfy $1/\sigma \ll r_0 \ll s_0$. (Recall that the relation $\ll$ depends on an arbitrary small constant $\eta > 0$.)

This recurrence in λ bottoms out when the value of λ falls below ϵ. Since t is a fixed positive power of $1/\epsilon$, this recurrence has depth $k = O(\log_t 1/\epsilon) = O(1)$, and σ never falls below $1/2^k = \Theta(1)$. If $\lambda < \epsilon$, we can invoke Lemma II.4 with an arbitrary small (albeit, fixed) positive power r of $1/\epsilon$ that satisfies $s_0 \ll r$. As a result, we obtain

$$f(\epsilon) = O\left(r \cdot f\left(\epsilon r^{1-\eta'}\right) + s_0 \cdot f\left(\epsilon s_0^{1-\eta'}\right) + r_0 \cdot f\left(\epsilon r_0^{1-\eta'}\right)\right) +$$

$+O\left(\frac{1}{\epsilon^{3/2+\eta'}}\right)$ for an arbitrary small constant $\eta' = \Theta(\eta) > 0$.

Bounding $f(\epsilon)$**.** The last recurrence terminates when $\epsilon \geq \tilde{\epsilon}$, in which case we have $f(\epsilon) = O(1)$. By following the standard inductive approach which applies to recurrences of this type (see, e.g., [30], and also [20], [36] and [37, Section 7.3.2]), and fixing suitably small constants η and $\tilde{\epsilon}$ (and, thereby, also $\eta' = \Theta(\eta)$), the recurrence solves to $f(\epsilon) = O\left(\epsilon^{-3/2-\gamma}\right)$, where $\gamma > 0$ is an arbitrary small constant, and the constant of proportionality depends on γ. This concludes the proof of Theorem I.1. □

IV. CONCLUDING REMARKS

1. Our analysis is largely inspired by the partition-based proof [17] of the Szemerédi-Trotter Theorem [38] 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 of P is involved in $o(1/\epsilon)$ such canonical sets (which are naturally associated with the surrounding sectors $\mathcal{W}_j(p)$). This gives rise to a sparse restriction graph Π.

2. Our proof of Theorem I.1 is fully constructive, and the resulting net includes points of the following types:

- The vertices of the decompositions $\Sigma(r_1)$ which arise in the various recursive instances.
- 1-dimensional $\hat{\epsilon}$-nets within lines $L \in \mathcal{Y}(r_0)$, 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 Π_t (in Section III-C) or $\Pi(i)$ (in Section III-D).
- 1-dimensional $\hat{\epsilon}$-nets within lines $L \in \mathcal{Y}(r_0)$, 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.
- 2-dimensional $\hat{\epsilon}$-nets of Theorem II.1 with respect to triangles in $\mathbb{R}^2$.

3. 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 $\mathbb{R}^2$.

The author conjectures that the true 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 $\mathbb{R}^d$ [22]. The connection between these notions was explored by Mustafa and Ray [32] and, more recently, by Har-Peled and Jones [21].

4. 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 that is close to $n^2/\sqrt{\epsilon}$.

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, the author is indebted to Gábor Tardos for pointing out that Proposition III.4 extends to vertical decompositions, which substantially simplified the overall exposition.

Funding acknowledgement: 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. The author was also supported by grant 1452/15 from Israel Science Foundation, by grant 2014384 from the U.S.-Israeli Binational Science Foundation, and by Ralph Selig Career Development Chair in Information Theory.

REFERENCES

[1] 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.

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

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

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

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

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

[7] 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.

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

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

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

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

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

[13] B. Chazelle, The discrepancy method: randomness and complexity, Camb. Univ. Press, New York, NY, USA, 2000.

[14] 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.

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

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

[17] 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.

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

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

[20] D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions, with applications to visibility in terrains, *Disc. Comput. Geom.* 12 (1994), 313-326.

[21] S. Har-Peled and M. Jones, How to net a convex shape, https://arxiv.org/abs/1712.02949, 2017.

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

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

[24] C. Keller, S. Smorodinsky and G. Tardos, Improved bounds on Hadwiger-Debrunner numbers, *Israel J. Math.* 225 (2) (2018), 925–945.

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

[26] L. Martínez-Sandoval, E. Roldán-Pensado, and N. Rubin, Further consequences of the colorful Helly hypothesis, in *Proc. 34th Sympos. Comput. Geom. 2018*, pp. 59:1–59:14.

[27] 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.

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

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

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

[31] 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.

[32] 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.

[33] 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.

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

[35] N. Rubin, An improved bound for weak Epsilon-nets in the plane, https://arxiv.org/abs/1808.02686, 2018.

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

[37] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cam. Univ. Press, NY, 1995.

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

[39] 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.