

Output-Sensitive Tools for Range Searching in Higher Dimensions *

Micha Sharir[†] Shai Zaban[‡]

October 16, 2013

Abstract

Let P be a set of n points in \mathbb{R}^d . A point $p \in P$ is k -shallow if it lies in a halfspace which contains at most k points of P (including p). We show that if all points of P are k -shallow, then P can be partitioned into $\Theta(n/k)$ subsets, so that any hyperplane crosses at most $O((n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k))$ subsets. Given such a partition, we can apply the standard construction of a spanning tree with small crossing number within each subset, to obtain a spanning tree for the point set P , with crossing number $O(n^{1-1/(d-1)} k^{1/d(d-1)} \log^{2/(d-1)}(n/k))$. This allows us to extend the construction of Har-Peled and Sharir [HS11] to three and higher dimensions, to obtain, for any set of n points in \mathbb{R}^d (without the shallowness assumption), a spanning tree T with *small relative crossing number*. That is, any hyperplane which contains $w \leq n/2$ points of P on one side, crosses $O(n^{1-1/(d-1)} w^{1/d(d-1)} \log^{2/(d-1)}(n/w))$ edges of T . Using a similar mechanism, we also obtain a data structure for halfspace range counting, which uses $O(n \log \log n)$ space (and somewhat higher preprocessing cost), and answers a query in time $O(n^{1-1/(d-1)} k^{1/d(d-1)} (\log(n/k))^{O(1)})$, where k is the output size.

1 Introduction

1.1 Background

One of the central themes in computational geometry is the design of efficient range searching algorithms. Typically, in such a problem one is given an input set P of n points in d dimensions, and the goal is to preprocess P into a data structure, so that for any query *range* of some type (say, a halfspace H), one can efficiently count, report, or test for emptiness the set $P \cap H$. A significant component of most of these algorithms involves space decomposition techniques, which partition the input set into subsets with some useful structure.

*Work by Micha Sharir was supported by NSF Grant CCF-08-30272, by Grant 2006/194 from the U.S.-Israel Binational Science Foundation, by Grant 338/09 from the Israel Science Fund, by the Israeli Centers for Research Excellence (I-CORE) program (Center no. 4/11), and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University. The paper is part of the second author's M.Sc. dissertation, prepared under the supervision of the first author.

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

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; shaiza@bezeqint.net.

Apart from range searching, partitions are also used in constructions of spanning trees with small crossing number, approximations, ε -nets, and in several other computational geometry problems.

In a typical approach (see, e.g., [Mat92a]), given a point set P in \mathbb{R}^d , one partitions P into subsets of approximately equal sizes, so that each of them is contained within some region of constant description complexity (e.g., simplices). We remark that the subsets are pairwise disjoint, while the containing regions might have nonempty intersections. We seek partitions with *small crossing number*, meaning that the maximum number of enclosing regions crossed by the boundary of a range (typically, a hyperplane) is small.

In this paper, we consider a variant of the partition paradigm, referred to as a partition of a *shallow point set*, which is applied to a set P of *shallow* points in \mathbb{R}^d , where a point $p \in P$ is k -shallow if there exists a hyperplane h that contains at most k points of P , including p , on one of its sides. We will construct a partition whose crossing number, namely, the maximum number of its simplices that can be crossed by a hyperplane, depends, in addition to n , also on the shallowness of P , i.e., on the maximum shallowness k of its points.

As is typical in computational geometry, we consider the dimension d as a fixed (small) constant, thus factors depending only on d will be regarded as constants.

Cuttings. One of the major ingredients of our partitioning technique is *cuttings* of arrangements of hyperplanes. A cutting is a collection of (possibly unbounded) d -dimensional closed cells with constant description complexity (e.g., simplices) with pairwise disjoint interiors, which cover the entire \mathbb{R}^d (or some specified portion thereof). Let H be a collection of n hyperplanes in \mathbb{R}^d and let Ξ be a cutting of the arrangement $\mathcal{A}(H)$. For each simplex $\Delta \in \Xi$, let H_Δ denote the collection of hyperplanes intersecting the interior of Δ . The cutting Ξ is called a $(1/r)$ -*cutting* for H if $|H_\Delta| \leq n/r$ for every simplex $\Delta \in \Xi$.

It will sometimes be convenient to work with weighted collections of hyperplanes, where such a collection is a pair (H, w) , where H is a collection of hyperplanes, and $w : H \rightarrow \mathbb{R}^+$ is a weight function on H . For each $L \subseteq H$, we write $w(L)$ for $\sum_{h \in L} w(h)$. The notions introduced for unweighted collections of hyperplanes can usually be generalized for weighted collections in an obvious way. For example, a cutting Ξ is a $(1/r)$ -cutting for (H, w) if for every simplex $\Delta \in \Xi$, the collection H_Δ has total weight at most $w(H)/r$.

Partitions: related work. Our study extends to higher dimensions the planar construction of a partition of a set of shallow points, recently developed by Har-Peled and Sharir [HS11]. This technique partitions P into $O(n/k)$ subsets, each containing $O(k)$ points and enclosed in a (possibly unbounded) triangle, so that the triangles are pairwise disjoint, and the crossing number of this partition (by any line) is only $O(\log(n/k))$.

Har-Peled and Sharir also constructed a set P of n k -shallow points in \mathbb{R}^3 (in fact, P is a set in convex position, so all its points are 1-shallow), so that the (maximum) crossing number of any partition of P into subsets of size $\Theta(k)$ is $\Omega(\sqrt{n/k})$. See below for a further discussion of this phenomenon, which was the starting point and the motivation for the study presented in this paper.

We next review the more standard partitioning techniques of Matoušek [Mat92a, Mat92b]. There P is a set of n points in \mathbb{R}^d , $d \geq 2$, which do not have to be shallow. Again, we par-

partition P into $\Theta(n/k)$ pairwise disjoint subsets, each having between k and $2k$ points (here $1 < k < n$ is an arbitrary parameter), enclosed by some simplex, where these simplices might have a nonempty intersection¹. The partition, referred to as a *simplicial partition*, has the property, that any hyperplane h crosses at most $O((n/k)^{1-1/d})$ simplices.

Let P , n and k be as above. Matoušek [Mat92b] has also developed a variant of the partitioning scheme described above, yielding a partition of $\Theta(n/k)$ subsets of size between k and $2k$, such that any k -shallow hyperplane h (that is, a hyperplane which contains at most k points of P on one side) crosses at most $O((n/k)^{1-1/\lfloor d/2 \rfloor})$ simplices of the partition.

(In contrast, the construction in [HS11], and the one presented in this paper, deal with sets whose points are shallow, but we seek a small crossing number with respect to every hyperplane.)

Applications: related work. Let us now review some applications based on the various kinds of partitions mentioned above.

Halfspace range searching. Matoušek [Mat92a] uses the standard partition recursively to construct an efficient *partition tree*, whose root stores the entire input set P and a simplicial partition thereof. Each node v of the tree stores a subset P_v of P , and a simplicial partition of P_v , and each subset of the partition corresponds to a distinct child of v . Overall, the resulting partition tree has linear size, can be constructed in $O(n \log n)$ deterministic time, such that, given a query halfspace γ , one can count the number of points in $P \cap \gamma$ in $O(n^{1-1/d} \log^{O(1)} n)$ time.

In a variant of this approach, Matoušek [Mat92b] exploits the partition machinery with respect to shallow hyperplanes to efficiently solve the halfspace range reporting (or emptiness) problem. More specifically, one can construct a data structure of size $O(n \log \log n)$, in $O(n \log n)$ time, such that, given a query halfspace γ , one can report the points in $P \cap \gamma$ in $O(n^{1-1/\lfloor d/2 \rfloor} \log^{O(1)} n + k)$ time, where $k = |P \cap \gamma|$.

Another useful application of partitions of sets of shallow points, as described above, is the construction of spanning trees with small *relative* crossing number. We recall the standard result, due to Chazelle and Welzl [CW89] (see also [Wel92]), on the existence of spanning trees with small crossing number. That is, given a set P of n points in \mathbb{R}^d , there exists a straight-edge spanning tree \mathcal{T} on P , such any hyperplane crosses at most $O(n^{1-1/d})$ edges of \mathcal{T} . Har-Peled and Sharir have refined this construction in the plane, to obtain a spanning tree \mathcal{T} that has the following property. Define the *weight* $w_l \leq n/2$ of a line l to be the smaller of the two numbers of points of P on each side of l . Then l crosses only $O(\sqrt{w_l} \log(n/w_l))$ edges of \mathcal{T} ; see [HS11] for more details. In this paper, we follow the same machinery to extend the above construction to higher dimensions.

Relative (p, ε) -approximations. The existence of a spanning tree with small relative crossing number, for a point set P in the plane, as just reviewed, facilitates the construction of *relative (p, ε) -approximations* for P , with respect to halfplane ranges, whose size is smaller than the one guaranteed by the general theory of Li et al. [LLS01] (see [HS11]). For given $0 < p, \varepsilon < 1$, a relative (p, ε) -approximation for P is a subset $A \subseteq P$ with the property that,

¹A recent refined construction by Chan[XX] constructs simplices with pairwise disjoint simplices.

for each halfspace γ ,

$$(1 - \varepsilon) \frac{|A \cap \gamma|}{|A|} \leq \frac{|P \cap \gamma|}{|P|} \leq (1 + \varepsilon) \frac{|A \cap \gamma|}{|A|},$$

provided that $|P \cap \gamma|/|P| \geq p$. Given such a spanning tree, we convert it into a spanning path with the same asymptotic crossing number, and generate a matching out of this path by selecting every other edge along it, and apply a “halving technique” on the resulting matching, in which one point of each pair of the matching is drawn independently at random, thus getting rid of half of the points. By the standard theory of discrepancy (see [Cha01]), the small crossing number of the matching implies a correspondingly low discrepancy of the halving. We repeat this halving procedure until the combined discrepancy first exceeds a given prescribed parameter, and then return the remaining points as the desired approximation. Har-Peled and Sharir [HS11] apply this construction to a planar point set P , and obtain a relative (p, ε) -approximation for P , of size

$$O\left(\frac{1}{\varepsilon^{4/3} p} \log^{4/3} \frac{1}{\varepsilon p}\right).$$

For most values of ε and p , this is better than the general bound $O(\frac{1}{\varepsilon^{2p}} \log \frac{1}{p})$ given in [LLS01]. See Section 1.2 below for more details. In this paper, we extend this construction to higher dimensions, using the extension of spanning trees with small relative crossing number to higher dimensions, as mentioned above.

Our results. We introduce a new variant of the partitioning machinery, based on the general approach that was introduced in [Mat92a] and [Mat92b], and obtain a partition of a set P of n k -shallow points in \mathbb{R}^d into $\Theta(n/k)$ subsets, each of size between k and $2k$, so that each subset is contained within a d -dimensional simplex and so that the crossing number of the partition (the maximum number of simplices crossed by a hyperplane) is at most $O((n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k))$ for $d \geq 3$ and $O(\alpha(n/k) \log^2(n/k))$ for $d = 2$, where $\alpha(\cdot)$ is the near-constant inverse Ackermann function. Note that the size of the sets in the partition is roughly the same as the shallowness parameter k , and that the exponent in the bound on the crossing number is smaller than the one provided for the general setup in [Mat92a].

We use this partition, as in [HS11], to construct a spanning tree with small relative, output-sensitive, crossing number. Specifically, any hyperplane h of weight w_h (the number of input points on its “lighter” side) crosses at most $O(n^{1-1/(d-1)} w_h^{1/d(d-1)} \log^{2/(d-1)}(n/w_h))$ edges of the resulting spanning tree for $d \geq 3$, and $O(\sqrt{w_h} \alpha(n/w_h) \log^2(n/w_h) + \alpha(n) \log^4 n)$ edges for $d = 2$.

We then extend the planar construction of relative (p, ε) -approximations for points and halfplanes introduced in [HS11] to higher dimensions. We again follow a similar machinery as the one in [HS11], based on the classical halving technique, to convert a spanning tree with small relative crossing number to a relative (p, ε) -approximation. The properties of the relative (p, ε) -approximation and the exact result are detailed in Section 1.2 below.

Another application of the shallow-points partition is an exact output-sensitive range counting algorithm for point sets and halfspace ranges in \mathbb{R}^d . We combine the range reporting algorithm of Matoušek [Mat92b] with our shallow-points partition and the standard

simplex range counting algorithm of Matoušek [Mat92a], to obtain an improved, output-sensitive range counting algorithm, whose running time is $O(n^{1-1/(d-1)}k^{1/d(d-1)}(\log n)^{O(1)})$, for $d \geq 3$, where k is the output size. This is an improvement over the previous bound $O(n^{1-1/d}(\log n)^{O(1)})$ of [Mat92a] when $k \ll n$ (although when k is very small, the range reporting of [Mat92b] will be faster). See Section 5.2 for a more precise statement of this result. (One weak aspect of our structure is that, for $d \geq 4$ we do not know how to construct it in near-linear time. The current bound on the preprocessing cost is roughly n^{c_d} for some exponent c_d that depends on d and satisfies $1 < c_d < 3/2$. See Theorem 5.7 for the precise statement.)

1.2 Known notions of approximations

We recall the result of Li et al. [LLS01], and two useful extensions of ε -approximations and ε -nets derived from it, which we will exploit later in this paper.

Let (X, \mathcal{R}) be a *range space*, where X is a set of n objects and \mathcal{R} is a collection of subsets of X , called *ranges*. The *measure* of a range $\tau \in \mathcal{R}$ in X is the quantity

$$\bar{X}(\tau) = \frac{|X \cap \tau|}{|X|}.$$

Assume that (X, \mathcal{R}) has *finite VC-dimension* δ , which is a constant independent of n ; see [Mat02] for more details.

Let $0 < \alpha, \nu < 1$ be two given parameters. Consider the distance function

$$d_\nu(r, s) = \frac{|r - s|}{r + s + \nu}, \quad \text{for } r, s \geq 0.$$

A subset $N \subseteq X$ is called a (ν, α) -*sample* for (X, \mathcal{R}) , if for each $\tau \in \mathcal{R}$ we have

$$d_\nu(\bar{X}(\tau), \bar{N}(\tau)) < \alpha.$$

Theorem 1.1 (Li et al. [LLS01]) *A random sample of X of size*

$$O\left(\frac{1}{\alpha^2\nu} \left(\delta \log \frac{1}{\nu} + \log \frac{1}{q}\right)\right)$$

is a (ν, α) -sample for (X, \mathcal{R}) with probability at least $1 - q$, with an appropriate choice of an absolute constant of proportionality.

Har-Peled and Sharir [HS11] show that, by appropriately choosing α and ν , various standard constructs, such as ε -nets and ε -approximations, are special cases of (ν, α) -samples; thus bounds on the sample sizes that guarantee that the sample will be one of these constructs with high probability are immediately obtained from Theorem 1.1. Two other, less standard special cases of (ν, α) -samples are relative (p, ε) -approximations [HS11] and *shallow ε -nets* [ShSh11]; we review them next.

Relative (p, ε) -approximations.

Definition 1.2 (*Relative (p, ε) -Approximation*). Let (X, \mathcal{R}) be a range space of finite VC-dimension δ . For given parameters $0 < p, \varepsilon < 1$, a subset $N \subseteq X$ is a *relative (p, ε) -approximation* for (X, \mathcal{R}) if, for each $\mathbf{r} \in \mathcal{R}$, we have

- (i) $(1 - \varepsilon)\overline{X}(\mathbf{r}) \leq \overline{N}(\mathbf{r}) \leq (1 + \varepsilon)\overline{X}(\mathbf{r})$, if $\overline{X}(\mathbf{r}) \geq p$.
- (ii) $\overline{X}(\mathbf{r}) - \varepsilon p \leq \overline{N}(\mathbf{r}) \leq \overline{X}(\mathbf{r}) + \varepsilon p$, if $\overline{X}(\mathbf{r}) \leq p$.

The following result shows that a relative (p, ε) -approximation is a special case of a (ν, α) -sample.

Theorem 1.3 (Har-Peled and Sharir [HS11]) *A random sample N of*

$$\frac{c_1}{\varepsilon^2 p} \left(\delta \log \frac{1}{p} + \log \frac{1}{q} \right)$$

elements of X , for an appropriate absolute constant c_1 , is a relative (p, ε) -approximation for (X, \mathcal{R}) with probability at least $1 - q$.

Shallow ε -nets. Let (X, \mathcal{R}) be a range space of finite VC-dimension δ , and let $0 < \varepsilon < 1$ be a given parameter. A subset $N \subseteq X$ is a *shallow ε -net* if it satisfies the following two properties, for some constant c that depends on δ .

- (i) For each $\mathbf{r} \in \mathcal{R}$ and for any parameter $t \geq 0$, if $|N \cap \mathbf{r}| \leq t \log \frac{1}{\varepsilon}$ then $|X \cap \mathbf{r}| \leq c(t+1)\varepsilon|X|$.
- (ii) For each $\mathbf{r} \in \mathcal{R}$ and for any parameter $t \geq 0$, if $|X \cap \mathbf{r}| \leq t\varepsilon|X|$ then $|N \cap \mathbf{r}| \leq c(t+1) \log \frac{1}{\varepsilon}$.

Note the difference between shallow and standard ε -nets: Property (i) (with $t = 0$) implies that a shallow ε -net is also a standard ε -net (possibly with a re-calibration of ε). Property (ii) has no parallel in the case of standard ε -nets — there is no guarantee how a standard ε -net interacts with small ranges.

Theorem 1.4 (Sharir and Shaul [ShSh11]) *A random sample N of*

$$\frac{c'_1}{\varepsilon} \left(\delta \log \frac{1}{\varepsilon} + \log \frac{1}{q} \right)$$

elements of X , for an appropriate absolute constant c'_1 , is a shallow ε -net with probability at least $1 - q$.

A brief overview of the analysis. The paper is structured as follows. In Section 2 we derive a technical result that shows that, informally, a random sample Z of $\Theta(n/k)$ points from a k -shallow n -point set in \mathbb{R}^d has, with some constant positive probability, the property that at least some fixed fraction of its points are vertices of its convex hull. (Curiously, this fraction cannot exceed $1/e \approx 0.368$ in the worst case.) We plug this result into the analysis in Section 3, which presents a construction of a $(1/r)$ -cutting of the zone of the boundary of a convex set in an arrangement of hyperplanes in \mathbb{R}^d . This construction is an easy adaption of standard constructions of cuttings, but seems not to have been

presented so far in the literature. Armed with this construct, we present in Section 4 the main technical contribution of the paper, which is a construction of a simplicial partition of a set of shallow points in \mathbb{R}^d that has a small crossing number (with respect to any hyperplane). The construction is not general, in the sense that the size of the subsets in the partition has to be equal to (or proportional to) the shallowness parameter k . Still, this provides the key tool for the applications in Section 5, which, as mentioned above, are (i) a construction of a spanning tree with small relative crossing number, (ii) a construction of relative (p, ε) -approximations for halfspace ranges in \mathbb{R}^d , whose size (in many cases) is smaller than the halfspace general bound in [LLS01, HS11], and (iii) an output-sensitive halfspace range searching data structure.

2 Sampling from a k -shallow set

Let P be a set of n points in \mathbb{R}^d , for $d \geq 2$, so that all its points are k -shallow, for some fixed parameter $1 \leq k \leq n$.

Let Z be a random sample of P where each point of P is chosen in Z independently with probability $p = 1/k$. Let $Y \subseteq Z$ be the set of 1-shallow points of Z (with respect to Z), i.e., Y is the set of vertices of the convex hull $\text{Conv}(Z)$. Let us assume from now on that $k \geq 2$, unless explicitly noted otherwise.

In what follows, we assume that P is in general position. In particular, no $d+1$ points of P lie in a common hyperplane, and no d lie in a common vertical hyperplane (i.e., parallel to the x_d -axis). Consequently, considering the dual set of hyperplanes $H = \mathcal{D}(P)$ (using the duality that preserves the above/below relationship; see, e.g., [Ed87]), we get that H is in general position too, which means that every d hyperplanes of H intersect in a single point, and no $d+1$ have a common point.

The following proposition, although its proof is technical and somewhat involved, asserts a fairly intuitive property of k -shallow sets. That is, if we take a random sample of about n/k points of such a set, at least some fixed fraction of the points of the sample are vertices of its convex hull. As a partial explanation of why the proof is complicated, we observe, after the proof, that the fraction of convex hull vertices of the sample cannot in general be made arbitrarily close to 1.

Proposition 2.1 *Let P , Z , and Y be as above, and assume that $k \leq \frac{n}{256}$. Then, with probability > 0.22 ,*

$$|Y| \geq \frac{n}{8k}, \quad \text{and} \quad \frac{n}{2k} \leq |Z| \leq \frac{2n}{k}.$$

Before getting into the proof, let us recall Chernoff's bound (see, e.g., [AS92], Appendix A, Theorems A.12 and A.13) for a binomial random variable X with parameters n, p , which we denote by $X \sim B(n, p)$. That is, $X = \sum_{i=1}^n X_i$, where X_1, X_2, \dots, X_n are independent random indicator variables, and $\Pr\{X_i = 1\} = p$ for each i . Put $\mathbf{E}[X] = \mu = np$. Then, for any $\delta > 0$, we have

$$\Pr\{X > (1 + \delta)\mu\} < \left(\frac{e^\delta}{(1 + \delta)^{(1 + \delta)}} \right)^\mu \equiv \varepsilon(\mu, 1 + \delta). \quad (1)$$

Alternatively, we can write it as

$$\Pr\{X > t\} < \left(\frac{e^{1-\frac{\mu}{t}}\mu}{t}\right)^t = \varepsilon(\mu, t/\mu), \quad (2)$$

for any $t > \mu$. Another variant of Chernoff's bound, for large deviations below the mean, is

$$\Pr\{X < (1-\delta)\mu\} < e^{-\frac{\mu\delta^2}{2}}, \quad (3)$$

for any $\delta > 0$. Using (1) and (2), we get the following proposition:

Proposition 2.2 *Let $X \sim B(n, \frac{1}{k})$, and let $1 < \beta \leq k$ be a parameter. For $y \in \mathbb{R}$, define*

$$S_{>y} = \sum_{t>y \mid t \in \mathbb{N}} t \cdot \Pr\{X = t\}, \quad S_{\geq y} = \sum_{t \geq y \mid t \in \mathbb{N}} t \cdot \Pr\{X = t\}.$$

Then

$$S_{>\frac{\beta n}{k}} \leq \varepsilon(n/k, \beta) \cdot \left(\frac{\beta n}{k} + 1 + \frac{\beta}{\beta - e^{1-1/\beta}}\right).$$

Proof. We have, for $t_0 \in \mathbb{N}$,

$$\begin{aligned} S_{\geq t_0} &= \sum_{t \geq t_0} t \cdot \Pr\{X = t\} = \sum_{t \geq t_0} t(\Pr\{X \geq t\} - \Pr\{X \geq t+1\}) = \\ &t_0 \cdot \Pr\{X \geq t_0\} + \sum_{t=t_0+1}^n \Pr\{X \geq t\} = t_0 \cdot \Pr\{X \geq t_0\} + \sum_{t=t_0}^{n-1} \Pr\{X > t\}. \end{aligned}$$

We note, that for any $y \in \mathbb{R}$, $S_{>y} = S_{\geq \lfloor y+1 \rfloor}$. Hence, substituting $t_0 = \lfloor \frac{\beta n}{k} + 1 \rfloor$ into the equality above, and using the bound (1), with $\mu = \frac{n}{k}$ and $1 + \delta = \beta$, we get

$$\begin{aligned} S_{>\frac{\beta n}{k}} &= S_{\geq \lfloor \frac{\beta n}{k} + 1 \rfloor} = \left\lfloor \frac{\beta n}{k} + 1 \right\rfloor \cdot \Pr\left\{X \geq \left\lfloor \frac{\beta n}{k} + 1 \right\rfloor\right\} + \sum_{t=\lfloor \frac{\beta n}{k} + 1 \rfloor}^{n-1} \Pr\{X > t\} \leq \\ &\left(\frac{\beta n}{k} + 1\right) \Pr\left\{X > \frac{\beta n}{k}\right\} + \sum_{t=\lfloor \frac{\beta n}{k} + 1 \rfloor}^{n-1} \Pr\{X > t\} \leq \\ &\left(\frac{\beta n}{k} + 1\right) \cdot \varepsilon(n/k, \beta) + \sum_{t=\lfloor \frac{\beta n}{k} + 1 \rfloor}^{n-1} \Pr\{X > t\}. \end{aligned}$$

It remains to estimate the second expression. We note that $(e^{1-a/x})/x$ is monotonically decreasing in x for $0 < a \leq x$. Hence, using (2), we obtain, for any $y > \mu$,

$$\begin{aligned} \sum_{t \geq y} \Pr\{X > t\} &\leq \sum_{t \geq y} \left(\frac{e^{1-\frac{\mu}{t}}\mu}{t}\right)^t \leq \sum_{t \geq y} \left(\frac{e^{1-\frac{\mu}{y}}\mu}{y}\right)^t < \\ &\left(\frac{e^{1-\frac{\mu}{y}}\mu}{y}\right)^y \left(\frac{y}{y - e^{1-\frac{\mu}{y}}\mu}\right) = \varepsilon(\mu, y/\mu) \cdot \left(\frac{y}{y - e^{1-\frac{\mu}{y}}\mu}\right). \end{aligned}$$

Putting $y = \frac{\beta n}{k}$, and noting that $y > \mu$ (since $\beta > 1$ and $\mu = \frac{n}{k}$), we obtain

$$\sum_{t=\lfloor \frac{\beta n}{k} + 1 \rfloor}^{n-1} \Pr \{X > t\} \leq \sum_{t \geq \frac{\beta n}{k}} \Pr \{X > t\} \leq \varepsilon(n/k, \beta) \cdot \left(\frac{\beta}{\beta - e^{1-1/\beta}} \right),$$

which completes the proof. \square

Proof of Proposition 2.1. We separately bound from below the probabilities that (i) $|Y| \geq \frac{\alpha n}{k}$, and (ii) $\frac{n}{\nu k} \leq |Z| \leq \frac{\gamma n}{k}$, in terms of α, ν, γ, n and k ; the concrete values of α, ν and γ , as given in the proposition, will be substituted later. The probability for both conditions to hold is then estimated using the union bound (on the complementary events).

(i) Bounding from below the probability that $|Y| \geq \frac{\alpha n}{k}$. Put $q = \Pr \{|Y| < \frac{\alpha n}{k}\}$, and let $1 < \beta \leq k$ be a parameter. We need to bound from below the value of $1 - q$.

Since $Y \subseteq Z$ and $|Z| \sim B(n, \frac{1}{k})$, (1) implies that

$$\Pr \left\{ |Y| > \frac{\beta n}{k} \right\} \leq \Pr \left\{ |Z| > \frac{\beta n}{k} \right\} < \varepsilon(n/k, \beta). \quad (4)$$

Recalling that Y is the subset of 1-shallow points within Z , and that all the points in the ground set P are k -shallow, we have, for each point $a \in P$,

$$\Pr \{a \in Y\} \geq \frac{1}{k} \left(1 - \frac{1}{k} \right)^{k-1} \geq \frac{1}{ke}.$$

Indeed, as a is a k -shallow point, there exists a hyperplane through a which bounds a halfspace h containing at most $k - 1$ points of P , excluding a . The lower bound above is the probability that a is chosen in Z and none of the other points in h is chosen. Using linearity of expectation, we thus have

$$\mathbf{E}\{|Y|\} \geq \frac{|P|}{ke} = \frac{n}{ke}. \quad (5)$$

We now proceed to estimate q . Note that we cannot apply Chernoff's bound directly, because the events $a \in Y$, for $a \in P$, are not independent. Instead we proceed as follows. Put $\varepsilon = \varepsilon(n/k, \beta)$, and $s = S_{> \frac{\beta n}{k}}$. We then have (assuming that $\beta > \alpha$, which will hold for specific values that we will later use)

$$\begin{aligned} \frac{n}{ke} &\leq \mathbf{E}\{|Y|\} = \sum_{t=0}^n t \cdot \Pr \{|Y| = t\} = \\ &\quad \sum_{t < \frac{\alpha n}{k}} t \cdot \Pr \{|Y| = t\} + \sum_{\frac{\alpha n}{k} \leq t \leq \frac{\beta n}{k}} t \cdot \Pr \{|Y| = t\} + \\ &\quad \sum_{t > \frac{\beta n}{k}} t \cdot \Pr \{|Y| = t\} \leq q \frac{\alpha n}{k} + (1 - q - \varepsilon) \frac{\beta n}{k} + s. \end{aligned}$$

Hence, by applying Proposition 2.2, we have

$$q \leq \frac{\beta(1 - \varepsilon) - e^{-1} + s \frac{k}{n}}{\beta - \alpha} \leq \frac{\beta - e^{-1} + \varepsilon \frac{k}{n} \left(1 + \frac{\beta}{\beta - e^{1-1/\beta}} \right)}{\beta - \alpha},$$

or

$$1 - q \geq 1 - \frac{\beta - e^{-1} + \varepsilon \frac{k}{n} \left(1 + \frac{\beta}{\beta - e^{1-1/\beta}}\right)}{\beta - \alpha}. \quad (6)$$

By choosing (somewhat arbitrarily) $\alpha = \frac{1}{8}$, $\beta = 1.2$, recalling that we have required that $k \leq \frac{n}{256}$, and using Chernoff's bound (1) to bound $\varepsilon = \varepsilon(n/k, \beta)$, we get that $\Pr\{|Y| \geq \frac{n}{8k}\} > 0.221$.

(ii) **Bounding from below the probability that $\frac{n}{\nu k} \leq |Z| \leq \frac{\gamma n}{k}$.** Since $|Z| \sim B(n, \frac{1}{k})$, estimating this probability can be done by a direct application of Chernoff's bounds. As in (4),

$$\Pr\left\{|Z| > \frac{\gamma n}{k}\right\} < \varepsilon(n/k, \gamma),$$

and by applying (3), we have

$$\Pr\left\{|Z| < \frac{n}{\nu k}\right\} < e^{-\frac{n}{k} \frac{(1-1/\nu)^2}{2}}.$$

Choosing $\gamma = 2$, $\nu = 2$, and using the assumption that $k \leq \frac{n}{256}$, we get that

$$\Pr\left\{\frac{n}{2k} \leq |Z| \leq \frac{2n}{k}\right\} > 0.999.$$

From (i) and (ii) we obtain that

$$\Pr\left\{\left(|Y| \geq \frac{n}{8k}\right) \wedge \left(\frac{n}{2k} \leq |Z| \leq \frac{2n}{k}\right)\right\} > 0.22,$$

as asserted. \square

Upper bounds on the probability of a sample with many hull vertices. We can increase the lower bound given in (6), on the probability $1 - q$ of getting a sample with at least $\frac{\alpha n}{k}$ hull vertices, by choosing β close enough to 1 and α close to 0. We then have

$$1 - q \geq 1 - \frac{\beta - e^{-1} + \varepsilon \frac{k}{n} \left(1 + \frac{\beta}{\beta - e^{1-1/\beta}}\right)}{\beta - \alpha} \approx 1 - \left(1 - e^{-1} + \varepsilon \frac{k}{n} \left(1 + \frac{\beta}{\beta - e^{1-1/\beta}}\right)\right).$$

By requiring that n/k is at least some sufficiently large constant (which depends on the choice of β ; we need it to "neutralize" the denominator $\beta - e^{1-1/\beta}$, which tends to 0 as $\beta \rightarrow 1$), we can make this probability arbitrarily close to

$$1 - (1 - e^{-1}) = e^{-1} \approx 0.368.$$

The following construction shows that the bound $1/e$ is fairly close to being worst-case tight. This construction was provided by Adam Sheffer, and we are grateful to him for this observation.

The construction is depicted in Figure 1. It consists of a set A of n points on the unit circle C , sufficiently closely clustered together near the top point of C , and of $2c + 1$ sets T, L_1, \dots, L_c , and R_1, \dots, R_c , so that (i) each of these sets consists of points lying on a ray

emanating from the center of C , outside C and sufficiently close to it, (ii) each of these sets consists of k points, except for T which consists of $k - 1$ points, and (iii) the ray containing T points upwards, the rays containing L_1, \dots, L_c point into the third quadrant, and the rays containing R_1, \dots, R_c point into the fourth quadrant. Here c is an additional parameter that we will fix later. (The points are not in convex position but a sufficiently small perturbation of them will place them in general position without affecting the analysis.)

Altogether, the resulting set P has $n' = n + (2c + 1)k - 1$ points. It is easily verified that the points can be arranged in such a way that each of them is k -shallow. Moreover, with an appropriate layout, the convex hull of a sample Z will have at most $2c + 1$ vertices if we choose at least one point of T , at least one point of $L_1 \cup \dots \cup L_c$, and at least one point of $R_1 \cup \dots \cup R_c$. The probability of this latter event is

$$q' = \left(1 - \left(1 - \frac{1}{k}\right)^{k-1}\right) \left(1 - \left(1 - \frac{1}{k}\right)^{ck}\right)^2 \approx (1 - e^{-1})(1 - e^{-c})^2.$$

Assuming that $n \gg (2c + 1)k$, so that $2c + 1 \ll n'/k$, the probability of getting a sample with $\Theta(n'/k)$ hull vertices, with an appropriate constant of proportionality ensuring that this bound is larger than $2c + 1$, is

$$1 - q' = 1 - (1 - e^{-1})(1 - e^{-c})^2,$$

which we can make arbitrarily close to

$$1 - (1 - e^{-1}) = e^{-1},$$

by choosing c sufficiently large. This shows that e^{-1} is an upper bound on the probability that we can guarantee for getting a sample with $\Theta(n/k)$ hull vertices from a set of n k -shallow points (when n is sufficiently large). As noted above, we can get arbitrarily close to this bound by requiring k/n to be sufficiently small.

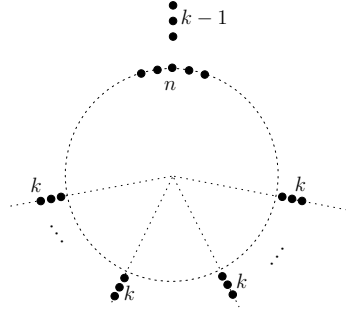


Figure 1: Construction of a k -shallow point set with large probability for a sample with a constant number of hull vertices.

Z is also a relative approximation set. Consider next the range space (P, \mathcal{R}) , where \mathcal{R} is the collection of all simplices in \mathbb{R}^d . As is well known, (P, \mathcal{R}) has finite VC-dimension, which we denote by δ . As a matter of fact, we have $\delta = O(d^2 \log d)$; see, e.g., [SA95??].

Proposition 2.1 establishes a property of a random sample of size roughly n/k from a set of n k -shallow points in \mathbb{R}^d . Of course, such a sample is also a (ν, α) -sample with high

probability, for appropriate choices of ν and α . The following proposition combines both properties.

Remark: The standard sampling method, under which Theorem 1.3 was originally proved, is to choose each subset of the prescribed size s with equal probability. However, the theorem also holds if the subset is obtained by choosing each point $x \in X$ independently with probability $s/|X|$, similar to the way Z was chosen. In this case s is only the expected size of the sample.

Proposition 2.3 *Let P, Z, n, k be as above, and let c_1 be the constant in Theorem 1.3. Then, with probability at least $q = 0.999$, Z is a relative $(\frac{c_2 k}{n} \log \frac{n}{c_2 k}, 1/2)$ -approximation for (P, \mathcal{R}) , for $c_2 \geq 8c_1(\delta + \log(1/q))$, and for $n \geq 2c_2 k$.*

Proof. As already noted, the sample Z satisfies $|Z| \geq \frac{n}{2k}$, it suffices to show that

$$\frac{c_1}{\varepsilon^2 p} \left(\delta \log \frac{1}{p} + \log \frac{1}{q} \right) \leq \frac{n}{2k},$$

for the specific values $p = \frac{c_2 k}{n} \log \frac{n}{c_2 k}$, $\varepsilon = 1/2$, $q = 0.999$, and c_2 as set above. Indeed, we have

$$\begin{aligned} \frac{c_1}{\varepsilon^2 p} \left(\delta \log \frac{1}{p} + \log \frac{1}{q} \right) &\leq \frac{4c_1 \delta}{c_2} \left(1 - \frac{\log \log \frac{n}{c_2 k}}{\log \frac{n}{c_2 k}} \right) \cdot \frac{n}{k} + \frac{4c_1}{c_2} \cdot \frac{\log \frac{1}{q}}{\log \frac{n}{c_2 k}} \cdot \frac{n}{k} \leq \\ &\frac{4c_1(\delta + \log \frac{1}{q})}{c_2} \cdot \frac{n}{k} \leq \frac{n}{2k}, \end{aligned}$$

as asserted, provided that $n \geq 2c_2 k$ (which implies that $\log \frac{n}{c_2 k} \geq 1$). \square

3 Cutting lemma for the zone of a convex set

Let H be a collection of n hyperplanes in \mathbb{R}^d . We denote the arrangement of H by $\mathcal{A}(H)$. For a simplex Δ , let H_Δ denote the collection of hyperplanes of H that *cross* Δ (i.e., intersect its interior). A $(1/r)$ -*cutting* for $\mathcal{A}(H)$ is a collection Ξ of (possibly unbounded) closed d -dimensional simplices with pairwise disjoint interiors, which cover \mathbb{R}^d , such that $|H_\Delta| \leq n/r$, for every $\Delta \in \Xi$. The *size* of a $(1/r)$ -cutting is the number of its simplices. The so-called Cutting Lemma (see, e.g., [Mat02]) asserts that for every H and $r \leq n$ there exists a $(1/r)$ -cutting of size $O(r^d)$ for H (which is asymptotically the best possible size).

Here we will need a modified version of the Cutting Lemma, where the cutting is not required to cover the entire \mathbb{R}^d , but only the zone in $\mathcal{A}(H)$ of the boundary ∂C of some fixed convex region C , where the zone of ∂C is the collection of all cells of $\mathcal{A}(H)$ whose interiors are crossed by ∂C . The simplices of our cutting might also contain points outside this zone, but they are required to cover all the cells of the zone. (In a sense, this can be regarded as a variant of the shallow cutting lemma of Matoušek [Mat92b].)

Theorem 3.1 (Cutting Lemma for the zone of a convex set). *Let H be a collection of n hyperplanes in \mathbb{R}^d , let $r \leq n$ be a parameter, and let σ be the boundary of a convex set in \mathbb{R}^d . Then there exists a $(1/r)$ -cutting Ξ for the zone of σ in the arrangement $\mathcal{A}(H)$, consisting of $O(r^{d-1} \log r)$ simplices, for $d \geq 3$, and of $O(r\alpha(r))$ simplices, for $d = 2$.*

We will need a variant of this result for weighted collections of hyperplanes, where such a collection is a pair (H, w) , where H is a collection of hyperplanes, and $w : H \rightarrow \mathbb{R}^+$ is a nonnegative weight function on H . For a subset $X \subseteq H$, we write $w(X)$ for $\sum_{h \in X} w(h)$. The notions introduced for unweighted collections of hyperplanes can usually be generalized for weighted collections in an obvious way. In particular, given the boundary σ of a convex set in \mathbb{R}^d , we say that Ξ is a $(1/r)$ -cutting for the zone of σ in $\mathcal{A}(H, w)$, if the simplices of Ξ cover all points in the zone, and for every simplex Δ of Ξ , the collection H_Δ has total weight $w(H_\Delta) \leq w(H)/r$. A simple reduction from cuttings for weighted collections of hyperplanes to unweighted cuttings is discussed in [Mat91a], and can be applied here for a $(1/r)$ -cutting for the zone of σ in $\mathcal{A}(H)$ without any change. Hence, Theorem 3.1 implies the following.

Corollary 3.2 *Let (H, w) be a weighted collection of n hyperplanes in \mathbb{R}^d , let $r \leq n$ be a parameter, and let σ be the boundary of a convex set. There exists a $(1/r)$ -cutting Ξ for the zone of σ in $\mathcal{A}(H, w)$ of size $O(r^{d-1} \log r)$, for $d \geq 3$, and $O(r\alpha(r))$, for $d = 2$.*

Before proving Theorem 3.1, we need the following technical lemma.

Lemma 3.3 *Let $X \sim B(n, p)$ and put $\mu = \mathbf{E}\{X\} = np$. Then*

- (i) $\mathbf{E}\{X^k \log(X + 1)\} = O(\mu^k \log \mu)$, for any $k \geq 1$.
- (ii) $\mathbf{E}\{X\alpha(X + 1)\} = O(\mu\alpha(\mu))$.

Proof. This follows from the fact that binomial variables are concentrated near their means, and we omit the routine and technical details. \square

The proof of Theorem 3.1 uses a triangulation of the cells of an arrangement of hyperplanes, called *canonical triangulation* (or *bottom-vertex triangulation*). The definition and some properties of this triangulation can be found in [Mat02]. For a subcollection $L \subseteq H$ of hyperplanes, let $CT(L)$ denote the set of simplices in the canonical triangulation of $\mathcal{A}(L)$. Let σ be the boundary of a convex set in \mathbb{R}^d . We denote by $CT_\sigma(L) \subseteq CT(L)$ the minimal set of simplices of $CT(L)$ which cover the zone of σ in $\mathcal{A}(L)$. We recall the following results about the canonical triangulation:

Lemma 3.4 *Let H be a collection of n hyperplanes in \mathbb{R}^d .*

- (i) [CF90] *For every simplex $\Delta \in CT(H)$, there exists a unique inclusion-minimal collection $S_\Delta \subseteq H$, such that $\Delta \in CT(S_\Delta)$. This collection S_Δ , called the defining set of Δ , consists of at most $D = d(d + 3)/2$ hyperplanes.*
- (ii) [CF90] *If $L \subseteq H$ and $\Delta \in CT(L)$ then $\Delta \in CT(H)$ if and only if its interior is intersected by no hyperplane of H .*
- (iii) [Mat02] *Assuming H is in general position, each cell in $\mathcal{A}(H)$ is a simple polytope, in which the number of simplices in its canonical triangulation is at most proportional to the number of its vertices (with the constant of proportionality depending on d).*

Let H , n , r and σ be as in Theorem 3.1. For a simplex Δ , define the *excess* of Δ to be $\max\{1, |H_\Delta| \frac{r}{n}\}$. Let R be a random sample of hyperplanes of H , where each hyperplane is drawn independently with probability $p = r/n$. Let $n_\sigma(p, t)$ denote the expected number of simplices with excess at least t in $CT_\sigma(R)$.

Lemma 3.5 (Exponential Decay Lemma). *For $t > 1$ and $r \leq n/2$,*

$$n_\sigma(p, t) = O(2^{-t} n_\sigma(p/t, 1)),$$

where $n_\sigma(p/t, 1)$ is the expected size of $CT_\sigma(R')$, for another random sample R' , in which each hyperplane is chosen independently with probability p/t .

We omit the proof of the lemma, which is essentially the same as in [Mat92a] and [Mat92b], and it can be viewed as a special instance of the general setup considered in [AMS98].

As a final step before the proof of Theorem 3.1, we recall the bound established in [APS93] on the complexity of the zone of the boundary of a convex set in a hyperplane arrangement.

Theorem 3.6 (Extended Zone Theorem). *The complexity of the zone of the boundary of an arbitrary convex set in an arrangement of n hyperplanes in \mathbb{R}^d is $O(n^{d-1} \log n)$ for $d \geq 3$, where the constant of proportionality depends on d , and $O(n\alpha(n))$ for $d = 2$.*

Proof of Theorem 3.1. First, we may assume that $r \leq n/2$, since otherwise, the bottom-vertex triangulation $CT_\sigma(H)$ can serve as the desired cutting. By Lemma 3.4 (iii), the number of simplices in $CT_\sigma(R)$ is at most proportional to the number of vertices in the zone of σ in $\mathcal{A}(R)$. Hence, by Theorem 3.6, we have

$$|CT_\sigma(R)| = \begin{cases} O(|R|^{d-1} \log(|R| + 1)), & \text{for } d \geq 3 \\ O(|R|\alpha(|R| + 1)), & \text{for } d = 2. \end{cases}$$

Since $|R| \sim B(n, p)$, Lemma 3.3 (i) implies that the expected number $n_\sigma(p, 1)$ of simplices in $CT_\sigma(R)$ is, for $d \geq 3$,

$$\begin{aligned} n_\sigma(p, 1) &= \mathbf{E}\{|CT_\sigma(R)|\} = \mathbf{E}\left\{O\left(|R|^{d-1} \log(|R| + 1)\right)\right\} = \\ &O((\mathbf{E}\{|R|\})^{d-1} \log(\mathbf{E}\{|R|\})) = O((np)^{d-1} \log(np)). \end{aligned} \tag{7}$$

For $d = 2$, a similar argument shows, by Lemma 3.3 (ii), that $n_\sigma(p, 1) = O(np\alpha(np))$. Note that these bounds hold for any $0 < p \leq 1$.

To obtain the desired cutting, we start with $\Xi_0 = CT_\sigma(R)$. For each simplex $\Delta \in \Xi_0$, let $t = t(\Delta)$ be the excess of Δ . Each simplex Δ with $t(\Delta) = 1$ is left as is. For those simplices Δ with $t(\Delta) > 1$, let Ξ_Δ be a $(1/t)$ -cutting for $\mathcal{A}(H_\Delta)$ of size $O(t^d)$. Such a cutting exists by the Cutting Lemma [Mat02]. We take the intersection of every simplex $\Delta' \in \Xi_\Delta$ with Δ , and triangulate it if necessary into $O(1)$ sub-simplices. The collection of simplices appearing in these triangulations, over all simplices $\Delta \in \Xi_0$, with $t(\Delta) > 1$, plus the simplices of Ξ_0 which were not triangulated further, form our cutting Ξ , which, as is

easily verified, is indeed a $(1/r)$ -cutting for the zone of σ in $\mathcal{A}(H)$. We bound the size of Ξ by bounding the expected value S of the sum $\sum_{\Delta \in \Xi_0} t(\Delta)^d$. We have

$$S \leq \sum_{t=1}^{\infty} t^d n_{\sigma}(p, t).$$

Using Lemma 3.5 and the bound (7), and recalling that $p = r/n$, we have, for $d \geq 3$, and for appropriate constants c and c' , depending on d ,

$$\begin{aligned} S &\leq \sum_{t=1}^{\infty} t^d \cdot c 2^{-t} n_{\sigma}(p/t, 1) \leq c' \sum_{t=1}^{\infty} t^d 2^{-t} \left(\frac{r}{t}\right)^{d-1} \log\left(\frac{r}{t}\right) \leq \\ &c' \left(\sum_{t=1}^{\infty} t 2^{-t}\right) r^{d-1} \log r = O(r^{d-1} \log r). \end{aligned}$$

For $d = 2$, arguing similarly, we have $S = O(r \alpha(r))$. This completes the proof of Theorem 3.1. \square

4 Partition theorem for shallow points in \mathbb{R}^d

Let P be a set of n points in \mathbb{R}^d . A *simplicial partition* for P is a collection

$$\Pi = \left\{ (P_1, \Delta_1), \dots, (P_m, \Delta_m) \right\},$$

where the P_i 's are pairwise disjoint subsets (called the *classes* of Π) forming a partition of P , and each Δ_i is a relatively open simplex containing P_i . We also require that the P_i 's be roughly of the same size, that is, $k \leq |P_i| \leq 2k$ for each i , for some parameter $k < n$ (so $m = \Theta(n/k)$). Assuming that the points of P are in general position, and $k \geq d + 1$, the simplices Δ_i are full-dimensional. We also note that the simplices Δ_i are not required to be pairwise disjoint, so a simplex Δ_i may also contain other points of P , in addition to those of P_i . (See however the recent work of Chan [Ch12] where the simplices can be made pairwise disjoint.)

For a hyperplane h , and a simplex Δ , we say that h *crosses* Δ if $h \cap \Delta \neq \emptyset$ and $\Delta \not\subset h$ (thus a hyperplane does not cross a lower-dimensional simplex contained in it). We define the *crossing number* of a hyperplane h (with respect to Π) as the number of simplices Δ_i crossed by h , and define the *crossing number* of Π as the maximum crossing number of any hyperplane with respect to Π .

We recall two previous versions of the partition theorems, both developed by Matoušek. The first one, used for obtaining an efficient simplex range counting algorithm, is the standard, general partition method. The second one, used for improved range reporting algorithms, is a partition method for shallow hyperplanes only, that is, the improved crossing number (relative to the first result), is guaranteed only for hyperplanes that contain up to a prescribed number of points on one side.

Theorem 4.1 (Partition Theorem [Mat92a]). *Let P be an n -point set in \mathbb{R}^d and let $1 < k < n$ be a parameter. There exists a simplicial partition Π for P of size $O(n/k)$, whose classes satisfy $k \leq |P_i| \leq 2k$, with crossing number $O((n/k)^{1-1/d})$.*

Theorem 4.2 (Partition Theorem for shallow hyperplanes [Mat92b]). *Let P be an n -point set in \mathbb{R}^d and let $1 < k < n$ be a parameter. There exists a simplicial partition Π for P of size $O(n/k)$, whose classes satisfy $k \leq |P_i| \leq 2k$, and such that the crossing number of any k -shallow hyperplane with respect to Π is $O((n/k)^{1-1/\lfloor d/2 \rfloor})$ for $d \geq 4$, and $O(\log(n/k))$ for $d = 2, 3$.*

Here is our version of the Partition Theorem, which caters to sets of shallow points and for any hyperplane.

Theorem 4.3 (Partition Theorem for shallow points). *Let P be a set of n k -shallow points in \mathbb{R}^d , where $d \geq 2$ and $2 \leq k \leq n/2$. There exists a simplicial partition Π for P , whose classes P_i satisfy $k \leq |P_i| \leq 2k$ (so their number m is $\Theta(n/k)$), and such that the crossing number of any hyperplane with respect to Π is $O\left((n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k)\right)$ for $d \geq 3$, and $O(\alpha(n/k) \log^2(n/k))$ for $d = 2$.*

Such a simplicial partition can be constructed in time $O(n^{1+\delta})$, for any fixed $\delta > 0$ (where the constants in this bound and in the bound on the crossing number depend on δ). The construction time can be improved to $O(n \log(n/k))$, provided that $k \geq n^\gamma$ for any fixed $\gamma > 0$ (again with the constants in this bound and in the bound on the crossing number depending on γ).

Remarks: (i) Note that, in contrast with Theorem 4.1 and 4.2, which hold for any $1 < k < n$, Theorem 4.3 only holds for the shallowness parameter k of the set P . (ii) We may assume that n/k is larger than some suitable constant, since otherwise we can partition P into a constant number ($\leq \lfloor n/k \rfloor$) of arbitrary classes, and the theorem then holds trivially.

The proof of Theorem 4.3 is based on the following lemma, adapted from similar lemmas in [Mat92a, Mat92b].

Lemma 4.4 *Let P , n and k be as above, and let Q be a given set of hyperplanes. Then there exists a simplicial partition Π for P , each of whose classes P_i satisfies $k \leq |P_i| \leq 2k$, such that the crossing number of every hyperplane $h \in Q$ with respect to Π is*

$$\begin{aligned} &O\left((n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k) + \log |Q|\right), \quad \text{for } d \geq 3, \text{ and} \\ &O\left(\alpha(n/k) \log^2(n/k) + \log |Q|\right), \quad \text{for } d = 2. \end{aligned}$$

Proof. We first present the proof for the case $d \geq 3$, and then consider the case $d = 2$, whose simpler proof uses the same machinery. We will inductively construct pairwise disjoint subsets P_1, P_2, \dots of P , and simplices $\Delta_1, \Delta_2, \dots$, so that $P_j \subseteq \Delta_j$ for each j . Suppose that P_1, \dots, P_i have already been constructed, and set $P'_i = P \setminus (P_1 \cup \dots \cup P_i)$ and $n_i = |P'_i|$, where we start with $P'_0 = P$. We iterate the construction as long as $n_i > Bk$, for an appropriate constant B that will be set below. Let Z_i be a random sample of P'_i , where each point is chosen independently with probability $1/k$, and let Y_i be the set of vertices of $\text{Conv}(Z_i)$, as in Section 2. By Proposition 2.1, with constant probability, Y_i and Z_i satisfy $|Y_i| \geq \frac{n_i}{8k}$ and $\frac{n_i}{2k} \leq |Z_i| \leq \frac{2n_i}{k}$ (here we use the assumption that n_i/k is sufficiently large; see below). Assume that the sample Z_i does satisfy these properties. Then, by Proposition 2.3, with probability at least $1-q$, Z_i is also a relative $(\frac{c_2 k}{n_i} \log \frac{n_i}{c_2 k}, \frac{1}{2})$ -approximation for (P'_i, \mathcal{R}) , where

\mathcal{R} is, as above, the set of all simplices in \mathbb{R}^d , $c_2 \geq 8c_1(\delta+1/q)$, c_1 is taken from Theorem 1.3, and $\delta = O(d^2)$ is the VC-dimension of (P, \mathcal{R}) . Let us set $c_2 = \max\{(d+1)/k, 8c_1(\delta+1/q)\}$, and assume that $c_2 \geq 2$.

We distinguish between two cases. First suppose that $n_i > Bk$. For a hyperplane $h \in Q$, let $\kappa_i(h)$ denote the number of simplices among $\Delta_1, \dots, \Delta_i$ crossed by h . We define a weighted collection (Q, w_i) by setting $w_i(h) = 2^{\kappa_i(h)}$ for each $h \in Q$.

We use a variant of the Cutting Lemma for weighted collections of hyperplanes. We fix a parameter t_i (again, its value will be set shortly), and construct a $(1/t_i)$ -cutting Ξ_i for the zone of the boundary σ of the convex hull of Y_i in the arrangement $\mathcal{A}(Q, w_i)$. By Corollary 3.2, Ξ_i consists of at most $c_3 t_i^{d-1} \log t_i$ simplices, for a suitable constant c_3 . We set

$$t_i = c_4 \left(\frac{n_i/k}{\log^2(n_i/k)} \right)^{1/(d-1)},$$

for an appropriate constant parameter $c_4 \leq 1$, whose concrete value will be set later. Let F be the solution of $F/\log^2 F = (1/c_4)^{d-1}$. We put $B = \max\{F, 2c_2, 256\}$ and note that (i) $t_i > 1$, (ii) $n_i \geq 256k$ (as required by Proposition 2.1), and (iii) $n_i \geq 2c_2k$ (as required by Proposition 2.3). We estimate the size of Ξ_i by

$$|\Xi_i| \leq c_3 t_i^{d-1} \log t_i \leq c_3 c_4^{d-1} \cdot \frac{n_i}{k} \cdot \frac{(d-1) \log c_4 + \log \frac{n_i}{k} - 2 \log \log \frac{n_i}{k}}{(d-1) \log^2 \frac{n_i}{k}} \leq \frac{c_3 c_4^{d-1}}{d-1} \cdot \frac{n_i}{k} \cdot \frac{1}{\log \frac{n_i}{k}}.$$

(The final inequality holds when $\log c_4 \leq 0$ and $\log \log(n_i/k) > 0$, which is indeed the case since $c_4 \leq 1$ and $n_i > 4k$.) Using this bound, we argue that there exists a (relatively open) simplex $\tau \in \Xi_i$ containing at least k points of P'_i . To see this, we first observe that each point of Y_i is contained in some simplex of Ξ_i . Hence there exists a relatively open simplex $\tau \in \Xi_i$, such that $|Y_i \cap \tau| \geq |Y_i|/|\Xi_i|$. Since $|Y_i| \geq n_i/8k$, the measure of τ in Z_i satisfies

$$\overline{Z}_i(\tau) = \frac{|Z_i \cap \tau|}{|Z_i|} \geq \frac{|Y_i \cap \tau|}{|Z_i|} \geq \frac{|Y_i|}{|\Xi_i|} \cdot \frac{1}{|Z_i|} \geq \frac{n_i}{8k} \cdot \frac{k}{2n_i} \cdot \frac{1}{|\Xi_i|} \geq \frac{k}{n_i} \cdot \frac{(d-1)}{16c_3 c_4^{d-1}} \cdot \log \frac{n_i}{k} = A \frac{k}{n_i},$$

where

$$A = \frac{d-1}{16c_3 c_4^{d-1}} \log \frac{n_i}{k}. \quad (8)$$

We can now bound from below the number of points of P'_i in τ . Put $D = \max\{2, (d+1)/k\}$. Since Z_i has the relative (p, ε) -approximation property with respect to P'_i , with the values of p and $\varepsilon = 1/2$ as specified in Proposition 2.3, we have, by Definition 1.2,

(i) If $\overline{P}'_i(\tau) \geq p = \frac{c_2 k}{n_i} \log \frac{n_i}{c_2 k}$, then, by the choice of c_2 ,

$$|P'_i \cap \tau| \geq p n_i = c_2 k \log \frac{n_i}{c_2 k} \geq c_2 k \geq \max\{d+1, 2k\}.$$

(ii) If $\overline{P}'_i(\tau) \leq p$, then

$$\overline{P}'_i(\tau) \geq \overline{Z}_i(\tau) - \varepsilon p \geq \frac{k}{n_i} \left(A - \frac{c_2}{2} \log \frac{n_i}{c_2 k} \right),$$

and thus we have

$$|P'_i \cap \tau| \geq k \left(A - \frac{c_2}{2} \log \frac{n_i}{c_2 k} \right).$$

By requiring that $A - \frac{c_2}{2} \log \frac{n_i}{c_2 k} \geq D$, we get that $|P'_i \cap \tau| \geq \max\{d+1, 2k\}$. Substituting the value of A in (8), this requirement becomes

$$\frac{d-1}{16c_3c_4^{d-1}} \log \frac{n_i}{k} \geq D + \frac{c_2}{2} \log \frac{n_i}{c_2 k},$$

or

$$\left(\frac{d-1}{16c_3c_4^{d-1}} - \frac{c_2}{2} \right) \log \frac{n_i}{k} \geq D - \frac{c_2}{2} \log c_2.$$

Since we assume that $n_i \geq 256k$, this inequality will be satisfied if

$$\frac{d-1}{2c_3c_4^{d-1}} \geq 4c_2 + D - \frac{c_2}{2} \log c_2. \quad (9)$$

This inequality in turn can be enforced by choosing c_4 to be a sufficiently small constant which depends only in d (note that the dependence of D on k does not prevent us from choosing c_4 to be a constant).

To recap, with this choice of c_4 (and thus of t_i), we can ensure, with probability at least $1 - q$, that $|P'_i \cap \tau| \geq \max\{d+1, 2k\}$, that is, τ is a full-dimensional simplex and contains at least $2k$ points of P'_i . We now set P_{i+1} to be an arbitrary k -point subset of $P'_i \cap \tau$, and put $\Delta_{i+1} = \tau$, so $P_{i+1} \subseteq \Delta_{i+1}$. See Figure 2 for an illustration.

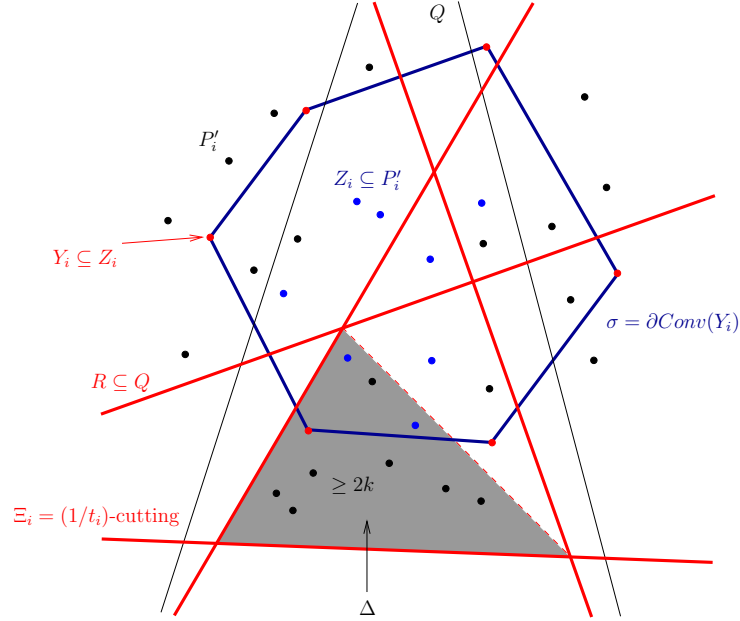


Figure 2: Δ is a simplex of the $(1/t_i)$ -cutting Ξ_i of (Q, w_i) which contains $\geq 2k$ points of P'_i .

Proceeding with this construction, we reach an index $i = q$, for which $n_q \leq Bk$. We then partition the set P'_q of the remaining points into at most B arbitrary subsets

$P_{q+1}, P_{q+2}, \dots, P_m$ of size k each, except for P_m , whose size is between k and $2k - 1$. We define the simplices $\Delta_{q+1} = \Delta_{q+2} = \dots = \Delta_m = \mathbb{R}^d$, and set $\Pi = \{(P_1, \Delta_1), \dots, (P_m, \Delta_m)\}$, thereby completing the construction. Since the last phase adds at most a constant number of (i.e., at most $\lfloor B \rfloor$) classes and simplices, it only increases the crossing number of any hyperplane by an additive constant. Hence it suffices to bound the crossing number of any hyperplane relative to $\Pi' = \{(P_1, \Delta_1), \dots, (P_q, \Delta_q)\}$.

To do so, we estimate the final total weight $w_q(Q)$ of the hyperplanes of Q in two different ways.

On one hand, the weight $w_q(h)$ of a hyperplane $h \in Q$ with crossing number κ is equal to 2^κ . Therefore, $\kappa \leq \log w_q(h) \leq \log w_q(Q)$.

On the other hand, let us consider how $w_{i+1}(Q)$ increases compared to $w_i(Q)$. Note that, since Ξ_i is a $(1/t_i)$ -cutting of the weighted collection (Q, w_i) , it follows that Δ_{i+1} is crossed by a collection of hyperplanes, denoted by Q^* , whose total weight at step i is at most $w_i(Q)/t_i$. When passing to step $i+1$, the weight of every hyperplane of Q^* is doubled, whereas the weight of any other hyperplane remains unchanged. Thus, the total weight of Q increases by at most $w_i(Q)/t_i$. That is,

$$\begin{aligned} w_{i+1}(Q) &\leq w_i(Q) + \frac{w_i(Q)}{t_i} = w_i(Q) \left(1 + \frac{1}{t_i}\right) \\ &\leq w_i(Q) \left(1 + \frac{1}{c_4} \left(\frac{\log^2(n_i/k)}{n_i/k}\right)^{1/(d-1)}\right). \end{aligned}$$

Let us put $r = n/k$, and recall that $w_0(Q) = |Q|$, $n_i = n - ik$, and $q \leq r$. Hence

$$w_q(Q) \leq |Q| \prod_{i=0}^{q-1} \left(1 + \frac{1}{t_i}\right) \leq |Q| \prod_{i=0}^{q-1} \left(1 + \frac{1}{c_4} \left(\frac{\log^2(r-i)}{r-i}\right)^{1/(d-1)}\right).$$

Taking logarithms and using the inequality $\ln(1+x) \leq x$, we get

$$\begin{aligned} \log w_q(Q) &\leq \log |Q| + \frac{1}{c_4 \ln 2} \sum_{i=0}^{q-1} \left(\frac{\log^2(r-i)}{r-i}\right)^{1/(d-1)} \\ &\leq \log |Q| + \frac{1}{c_4'} \sum_{j=1}^r \left(\frac{\log^2 j}{j}\right)^{1/(d-1)}, \end{aligned} \tag{10}$$

where $c_4' = c_4 \ln 2$. That is, we have

$$\begin{aligned} \kappa &\leq \log w_q(Q) \leq \log |Q| + \frac{1}{c_4'} \log^{2/(d-1)} r \sum_{j=1}^r \left(\frac{1}{j}\right)^{1/(d-1)} \\ &= O\left(\log |Q| + r^{1-1/(d-1)} \log^{2/(d-1)} r\right) \\ &= O\left(\log |Q| + (n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k)\right), \end{aligned}$$

for $d \geq 3$. This completes the proof for $d \geq 3$.

Consider next the case $d = 2$, and follow the same reasoning as in higher dimensions. At step i of the construction, we put

$$t_i = c_4 \cdot \frac{n_i/k}{\alpha(n_i/k) \log(n_i/k)},$$

and construct a $(1/t_i)$ -cutting Ξ_i of the zone of $\text{Conv}(Y_i)$ in $\mathcal{A}(Q, w_i)$, where w_i is defined as above. By Corollary 3.2 we have

$$|\Xi_i| \leq c_3 t_i \alpha(t_i) \leq c_3 c_4 \cdot \frac{n_i}{k} \cdot \frac{1}{\log \frac{n_i}{k}},$$

for an appropriate constant c_3 . Hence, (8) and (9), with $d = 2$, are valid for this case as well. Let F be the solution of $F/(\alpha(F) \log(F)) = 1/c_4$, and put $B = \max\{F, 2c_2, 256\}$ as in higher dimensions. Using the relative (p, ε) -approximation property of Z_i (which we may assume to hold), we obtain that there exists an open cell τ of Ξ_i , for which $|P'_i \cap \tau| \geq 2k$, provided that c_4 is chosen to be a sufficiently small constant, satisfying inequality (9). We take P_{i+1} to be an arbitrary subset of k points of $P'_i \cap \tau$, and put $\Delta_{i+1} = \tau$. We repeat this step until $n_i \leq Bk$, and then complete the construction as in the higher-dimensional case.

The analysis of the crossing number proceeds as above, except for the different value of the parameters t_i . Plugging this value into the appropriate variant of inequality (10), yields the following bound

$$\kappa \leq \log w_q(Q) \leq \log |Q| + \frac{1}{c'_4} \sum_{j=1}^r \frac{\alpha(j) \log j}{j} \leq \log |Q| + \frac{1}{c'_4} \alpha(r) \log r \sum_{j=1}^r \frac{1}{j} =$$

$$O(\log |Q| + \alpha(r) \log^2 r) = O(\log |Q| + \alpha(n/k) \log^2(n/k)),$$

where $c'_4 = c_4 \ln 2$, as above. This concludes the proof of the lemma. \square

The next step in the proof of Theorem 4.3 is to choose a small ‘test-set’ Q of hyperplanes, with the property that the crossing number of any hyperplane is at most proportional to the maximum crossing number of a hyperplane in Q .

Lemma 4.5 (Test-set lemma). *Let P , n , and k be as above. Then there exists a set Q of $O((n/k)^{d/(d-1)})$ hyperplanes, such that, for any simplicial partition $\Pi = \{(P_1, \Delta_1), \dots, (P_m, \Delta_m)\}$ satisfying $|P_i| \geq k$ for every i , the following holds: If κ_0 is the maximum crossing number of a hyperplanes of Q with respect to Π , then the crossing number of any hyperplane with respect to Π is at most*

$$(d+1)\kappa_0 + O\left((n/k)^{1-1/(d-1)}\right),$$

for $d \geq 3$, and $(d+1)\kappa_0 + 1$ for $d = 2$.

Proof. Let $H = \mathcal{D}(P)$ be the collection of hyperplanes dual to the points of P . Put $t = \Theta((n/k)^{1/(d-1)})$. Construct a $(1/t)$ -cutting Ξ of size $O(t^d)$ for $\mathcal{A}(H)$, and let Q be the set of all hyperplanes dual to the vertices of Ξ . Clearly $|Q| = O((n/k)^{d/(d-1)})$.

Fix a simplicial partition $\Pi = \{(P_1, \Delta_1), \dots, (P_m, \Delta_m)\}$ as above. Let h be any hyperplane in the primal space. The point $\mathcal{D}(h)$ dual to h is contained in a simplex $\sigma \in \Xi$. Let

G be the set of hyperplanes in the primal space dual to the vertices of σ . Clearly $G \subseteq Q$, so each of the $d + 1$ hyperplanes of G crosses at most κ_0 simplices of Π .

It remains to bound the number of simplices of Π which are crossed by the hyperplane h but by no hyperplane of G . Such simplices Δ_i must be completely contained in the zone of h in the arrangement $\mathcal{A}(G)$, and hence this zone must also contain the points of their corresponding classes P_i in its interior. It is elementary to verify (also see [Mat92a] and [Mat92b]), that any point of P lying in the interior of the zone of h in $\mathcal{A}(G)$ dualizes to a hyperplane of H intersecting the interior of the simplex σ , and there are at most n/t such hyperplanes in H . Hence, the zone of h may contain at most these many points of P , implying that there are at most $n/(tk) = O((n/k)^{1-1/(d-1)})$ simplices of Π completely contained in the zone of h in $\mathcal{A}(G)$, since $|P_i| \geq k$ for every i . In the plane, there is at most $n/(tk) = 1$ such simplex. \square

Proof of Theorem 4.3. The proof of the Partition Theorem now follows easily. Consider first the case $d \geq 3$. Given the set P of n k -shallow points, we first construct a test-set Q of $O((n/k)^{d/(d-1)})$ hyperplanes, as in Lemma 4.5. Then we apply Lemma 4.4, obtaining a simplicial partition Π such that any hyperplane $h \in Q$ has crossing number

$$O\left(\log |Q| + (n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k)\right) = O\left((n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k)\right)$$

relative to Π . (Here we use the fact that, when n/k is at least some sufficiently large constant, $\log |Q|$ is dominated by the other term in the bound.) By the ‘test-set’ property of Q , the crossing number of any hyperplane h relative to Π is at most

$$\begin{aligned} (d+1) \cdot O\left((n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k)\right) + O\left((n/k)^{1-1/(d-1)}\right) \\ = O\left((n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k)\right), \end{aligned}$$

as claimed. The case $d = 2$ is argued similarly and the resulting crossing numbers, of the lines of Q , and thus also of any line, are $O(\alpha(n/k) \log^2(n/k))$. This completes the proof of the existence and main properties of the partition.

It remains to analyze the preprocessing time of this partition. We actually carry out the standard algorithm of [Mat92a], which constructs a $(1/t_i)$ -cutting of the entire arrangement of $\mathcal{A}(Q, w_i)$, and use the preceding analysis to argue that one of the simplices in the cutting will contain (with high probability) $2k$ points of P'_i . There is no need to explicitly construct Z , Y , and the zone of the convex hull of Z , they are needed only to guarantee the existence of such a simplex. With this approach, we can thus conclude that the preprocessing time of our simplicial partition for a set of k -shallow points is at most $O(n^{1+\delta})$ for any value of k and for any fixed $\delta > 0$ (where the constants in this bound and in the bound on the crossing number both depend on δ). If $k \geq n^\alpha$, for any fixed $\alpha > 0$, the preprocessing time can be reduced to $O(n \log(n/k))$, as in [Mat92a] (where the constants in both bounds now depend on α). We note that, in the output-sensitive range counting algorithm below, we use the latter bound, since we deal there only with large values of k .

This concludes the proof of the Partition Theorem. \square

Let us note the following special case of Theorem 4.3, where the points of P are 1-shallow.

Corollary 4.6 *Let P be a set of n 1-shallow points in \mathbb{R}^d , where $d \geq 2$. There exists a simplicial partition Π for P , whose classes P_i satisfy $|P_i| = 2$ (except, possibly, for one class of size 3), such that the crossing number of any hyperplane with respect to Π is $O(n^{1-1/(d-1)} \log^{2/(d-1)} n)$ for $d \geq 3$, and $O(\alpha(n) \log^2 n)$ for $d = 2$.*

5 Applications

Overview. In this section we derive two major applications of the Partition Theorem for shallow sets (Theorem 4.3) presented in the previous section. The first application is the construction of a spanning tree with small relative crossing number, which depends on the shallowness of the crossing hyperplane. We then turn this construction into a construction of a relative (p, ε) -approximation, using the same machinery of Har-Peled and Sharir [HS11], thus extending their planar construction to higher dimensions. The second application is an output-sensitive halfspace range counting data structure, where the query time is better than that of the standard algorithm of Matoušek [Mat92a], when the hyperplane bounding the query halfspace is shallow.

5.1 Spanning trees with small relative crossing number in \mathbb{R}^d

In our first application, we extend the “weight-sensitive” version of *spanning trees with small relative crossing number* in the plane, studied by Har-Peled and Sharir [HS11], to higher dimensions. Both our extension and the construct in [HS11] refine the following classical construct of spanning trees with small crossing number, as obtained by Chazelle and Welzl [CW89], with a simplified construction given later in [Wel92].

Theorem 5.1 [CW89, Wel92] *Let P be a set of n points in \mathbb{R}^d , $d \geq 2$. Then there exists a straight-edge spanning tree \mathcal{T} of P such that each hyperplane in \mathbb{R}^d crosses at most $O(n^{1-1/d})$ edges of \mathcal{T} .*

We begin by refining this result for a set of k -shallow points.

Lemma 5.2 *Let P be a set of n k -shallow points in \mathbb{R}^d . One can construct a spanning tree \mathcal{T} for P , such that any hyperplane crosses at most $O\left(n^{1-1/(d-1)} k^{1/d(d-1)} \log^{2/(d-1)}(n/k)\right)$ edges of \mathcal{T} , for $d \geq 3$, and $O\left(\alpha(n/k) \log^2(n/k) \left(\sqrt{k} + \log(n/k)\right)\right)$ edges, for $d = 2$.*

Proof. We first consider the case $d \geq 3$. Given a set P of n k -shallow points, we construct a simplicial partition $\Pi = \{(P_1, \Delta_1), \dots, (P_m, \Delta_m)\}$ of P , by applying the Partition Theorem (Theorem 4.3), for sets of $2k$ -shallow points (clearly, the points of P are also $(2k)$ -shallow). (We note that replacing k by $2k$ does not affect asymptotically the crossing number of Π .) Each class P_i of the partition now satisfies $2k \leq P_i \leq 4k$.

Given Π , we first ignore its “tail”, namely the pairs $(P_{q+1}, \Delta_{q+1}), \dots, (P_m, \Delta_m)$, where q , as in the proof of Lemma 4.4, is the first index for which $n_q < B \cdot (2k) = B'k$. Let $\Pi' \subseteq \Pi$ denote the collection $\{(P_1, \Delta_1), \dots, (P_q, \Delta_q)\}$. For each $i = 1, \dots, q$, we construct a spanning tree \mathcal{T}_i for P_i with crossing number $O(k^{1-1/d})$ using Theorem 5.1. Then we connect those trees by segments, called *bridges*, into a single tree, whose construction will

be detailed below. Finally, we construct another spanning tree, denoted by \mathcal{T}_{q+1} , of the at most $B'k$ points in $P_{q+1} \cup \dots \cup P_m$, with crossing number $O((B'k)^{1-1/d}) = O(k^{1-1/d})$. We connect \mathcal{T}_{q+1} by a single edge to an arbitrary point in the other tree. The union of the trees $\mathcal{T}_1, \dots, \mathcal{T}_{q+1}$ together with the connecting bridges forms our spanning tree \mathcal{T} .

The connecting bridges are constructed as follows. We pick a point p_i from P_i , uniformly at random, for each $i = 1, \dots, q$. Let R denote the resulting set of $m = \Theta(n/k)$ points. The probability of any single point to belong to R is at most $1/(2k)$. (It is here that we use the larger size of the classes; see below.) Suppose that some point $a \in P$ has been chosen in R . By assumption, a is k -shallow, so there is a halfspace H which contains a and at most $k-1$ other points of P . Since each of these points appears in R with probability $\leq 1/(2k)$, the expected number of points in $R \cap H$ (conditioned on a being chosen in R) is $< 1/2$. Hence, using Markov's inequality, with (conditional) probability at least $1/2$, there is no other point apart from a in $R \cap H$, that is, a is a vertex of the convex hull of R . Hence, the expected number of vertices of $\text{Conv}(R)$ is at least $|R|/2$. We may assume that R satisfies this property, and denote by R_0 the subset of hull vertices of $\text{Conv}(R)$; Thus, $|R_0| \geq m/2$.

Apply Corollary 4.6 to the set R_0 , to obtain a partition of it into $\frac{1}{2}|R_0| \geq \frac{m}{4} - 1$ disjoint pairs. Let E denote the collection of the segments connecting the points in each pair. By construction, each hyperplane crosses at most $O(|R_0|^{1-1/(d-1)} \log^{2/(d-1)} |R_0|)$ segments of E .

The segments of E merge together some subsets of trees \mathcal{T}_i into larger trees. Since we created at least $\frac{m}{4} - 1$ bridges, the number of disconnected trees is now $m' \leq m - ((m/4) - 1) = 3m/4 + 1$. We now repeat the above construction to the new trees, or, rather, to the subsets of P that they span. That is, we choose a random point from each subset, take the subset of hull vertices of the resulting sample, and apply Corollary 4.6 to it, to create new bridges. The situation has actually improved, because these vertex sets of the current collection of trees are now larger (while the points are still k -shallow), so the probability of choosing any specific point is smaller; hence the probability of a sampled point to become a hull vertex can only grow. In any case, we create at least $\frac{m'}{4} - 1$ new bridges, and keep iterating in this manner until all trees have been merged into a single tree.

As is easily checked, the number of bridges crossed by a hyperplane h is at most

$$O\left(\sum_{j \geq 0} \left(\left(\frac{3}{4}\right)^j \frac{n}{k}\right)^{1-1/(d-1)} \log^{2/(d-1)} \left(\left(\frac{3}{4}\right)^j \frac{n}{k}\right)\right) = O\left((n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k)\right).$$

In addition, h crosses $O\left((n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k)\right)$ of the simplices Δ_i , and, within each such simplex Δ_i , it crosses $O(k^{1-1/d})$ edges of the corresponding tree \mathcal{T}_i . Also, it crosses $O(k^{1-1/d})$ edges of the tree \mathcal{T}_{q+1} . Since the bound within the simplices dominates the other two bounds, we conclude that h crosses at most

$$O\left((n/k)^{1-1/(d-1)} \log^{2/(d-1)}(n/k) \cdot k^{1-1/d}\right) = O\left(n^{1-1/(d-1)} k^{1/d(d-1)} \log^{2/(d-1)}(n/k)\right)$$

edges of \mathcal{T} . This completes the proof for $d \geq 3$.

For $d = 2$, follow the same construction as described above, and re-estimate the crossing number of a hyperplane h relative to our tree \mathcal{T} . We observe that h crosses edges from $O(\alpha(n/k) \log^2(n/k))$ trees of \mathcal{T}_i , and within each of them it crosses $O(\sqrt{k})$ edges. In

addition, h crosses $O(\sqrt{k})$ edges of \mathcal{T}_{q+1} . Following the same mechanism of the bridge construction and its analysis, the number of bridges crossed by h is at most

$$\begin{aligned} O\left(\sum_{j \geq 0} \alpha\left(\left(\frac{3}{4}\right)^j \frac{n}{k}\right) \log^2\left(\left(\frac{3}{4}\right)^j \frac{n}{k}\right)\right) &= O\left(\alpha(n/k) \sum_{j \geq 0} (\log(n/k) - j)^2\right) \\ &= O(\alpha(n/k) \log^3(n/k)). \end{aligned}$$

Thus, altogether, the crossing number of \mathcal{T} is bounded by

$$O\left(\alpha(n/k) \log^2(n/k)(\sqrt{k} + \log(n/k))\right),$$

as asserted. \square

Let now P be a set of n points in \mathbb{R}^d (without the shallowness assumption). For a (non vertical) hyperplane h , let w_h^+ (resp., w_h^-) be the number of points of P lying above (resp., below or on) h , and define the *weight* of h , denoted by w_h , to be $\min\{w_h^+, w_h^-\}$.

Theorem 5.3 *Let P be a set of n points in \mathbb{R}^d . Then one can construct a spanning tree \mathcal{T} for P , such that any hyperplane h crosses at most $O\left(n^{1-1/(d-1)} w_h^{1/d(d-1)} \log^{2/(d-1)}(n/w_h)\right)$ edges of \mathcal{T} , for $d \geq 3$, or $O\left(\sqrt{w_h} \alpha(n/w_h) \log^2(n/w_h) + \alpha(n) \log^4 n\right)$ edges, for $d = 2$.*

Proof. We construct a sequence of subsets of P , as follows. Put $P'_0 = P$. At the i^{th} step, $i \geq 1$, let P_i denote the set of (at most) 2^i -shallow points of P'_{i-1} , and put $P'_i = P \setminus (P_1 \cup \dots \cup P_i)$. We stop when P'_i becomes empty. By construction, the i^{th} step removes at least 2^i points from P'_{i-1} , because any (exactly) 2^i -shallow halfspace H contains 2^i points of P'_i , all of which are (at most) 2^i -shallow. Hence, $|P'_i| \leq |P'_{i-1}| - 2^i$, and so the process terminates after $O(\log n)$ steps. At the i^{th} step, we construct a spanning tree \mathcal{T}_i for P_i , using Lemma 5.2, with $k = 2^i$. Connect the resulting trees by $O(\log n)$ additional straight segments (in an arbitrary manner) into a single spanning tree \mathcal{T} of P .

We claim that \mathcal{T} is the desired spanning tree. Indeed, consider an arbitrary hyperplane h of weight $w_h = k$. We observe that h cannot cross any of the trees \mathcal{T}_i , for $i > U = \lceil \log_2 k \rceil$. To see this, assume to the contrary that h crosses \mathcal{T}_j for some $j > U$. That is, there exist two points $p_1, p_2 \in P_j$ which are separated by h . Without loss of generality, assume that p_1 lies in the halfspace bounded by h which contains k points. In particular, p_1 must be k -shallow in P (and thus also in any subset P'_i containing it), so it must have been removed at some step $i \leq U$ and thus cannot belong to P_j .

Thus, h crosses only the first U trees of our construction. Hence, for $d \geq 3$, the number of edges of \mathcal{T} that h crosses, excluding the $O(\log n)$ connecting edges, is at most

$$\begin{aligned} &\sum_{i=1}^U O\left(n^{1-1/(d-1)} (2^i)^{1/d(d-1)} \log^{2/(d-1)}(n/2^i)\right) \\ &= O\left(n^{1-1/(d-1)} \sum_{i=1}^U \left(2^{1/d(d-1)}\right)^i (\log n - i)^{2/(d-1)}\right) \\ &= O\left(n^{1-1/d} \sum_{j=\log n - U}^{\log n - 1} \frac{j^{2/(d-1)}}{(2^{1/d(d-1)})^j}\right). \end{aligned}$$

We have $a = \log n - U = \log \frac{n}{k} - \delta$, for some $0 \leq \delta \leq 1$, and we can estimate the sum by the integral

$$\int_a^\infty \frac{x^{2/(d-1)}}{(2^{1/d(d-1)})^x} dx.$$

By integrating in parts, it easily follows that, for $a \geq 1$, $0 < u \leq 1$, and $w > 1$,

$$\int_a^\infty \frac{x^u}{w^x} dx \leq \frac{1}{\ln w} \cdot \frac{a^u}{w^a} + \frac{u}{\ln^2 w} \cdot \frac{1}{w^a} = O\left(\frac{a^u}{w^a}\right), \quad (11)$$

where the $O(\cdot)$ notation is with respect to the growth of a . This allows us to bound the sum above by

$$O\left(n^{1-1/d}(k/n)^{1/d(d-1)} \log^{2/(d-1)}(n/k)\right) = O\left(n^{1-1/(d-1)} k^{1/d(d-1)} \log^{2/(d-1)}(n/k)\right),$$

which is the asserted bound on the crossing number of h , for $d \geq 3$.

Let us consider the case $d = 2$. The crossing number of a line l relative to \mathcal{T} is bounded by

$$\sum_{i=1}^U O\left(\alpha(n/2^i) \log^2(n/2^i) \left(\sqrt{2^i} + \log(n/2^i)\right)\right),$$

which we separate into the two sums:

$$I_1 = \sum_{i=1}^U O\left(\sqrt{2^i} \alpha(n/2^i) \log^2(n/2^i)\right), \quad \text{and} \quad I_2 = \sum_{i=1}^U O\left(\alpha(n/2^i) \log^3(n/2^i)\right).$$

I_1 is bounded by

$$\sum_{i=1}^U O\left(\sqrt{2^i} (\log n - i)^2 \alpha(n/2^i)\right) = O\left(\sqrt{n} \sum_{j=\log n-U}^{\log n-1} \frac{j^2 \alpha(2^j)}{\sqrt{2^j}}\right). \quad (12)$$

$\alpha(\cdot)$ is very slowly increasing; in particular, it satisfies $\alpha(2x) \leq \alpha(x) + 1$ for any x , and so, for $\alpha(x) \geq 3$, we have

$$\frac{\alpha(2x)}{\alpha(x)} \leq 1 + \frac{1}{\alpha(x)} \leq \frac{4}{3}.$$

We can therefore bound the sum in (12), denoted by S , by

$$S \leq \alpha(n/k) \sum_{j=\log n-U}^{\log n-1} \frac{j^2 \left(\frac{4}{3}\right)^{j-\log n+U}}{\sqrt{2^j}} \leq \alpha(n/k) \left(\frac{3}{4}\right)^{\log n-U} \sum_{j=\log n-U}^{\log n-1} j^2 \left(\frac{4}{3\sqrt{2}}\right)^j.$$

Using again the integral bound in (11), we obtain, for an appropriate constant c_1 ,

$$S \leq c_1 \cdot \alpha(n/k) \left(\frac{3}{4}\right)^{\log n-U} \log^2(n/k) \left(\frac{4}{3\sqrt{2}}\right)^{\log n-U} = O\left(\sqrt{\frac{k}{n}} \alpha(n/k) \log^2(n/k)\right),$$

from which we obtain

$$I_1 = O\left(\sqrt{k} \alpha(n/k) \log^2(n/k)\right).$$

For I_2 , we use the trivial bound

$$I_2 = \sum_{i=1}^U O(\alpha(n/2^i) \log^3(n/2^i)) = O(\alpha(n) \log^4 n).$$

Altogether, the crossing number of l is at most

$$O\left(\sqrt{k}\alpha(n/k) \log^2(n/k) + \alpha(n) \log^4 n\right).$$

This establishes the bound on the crossing number for $d = 2$. \square

Remark: Note that the bound in Theorem 5.3 for the planar case is slightly worse than the bound derived in [HS11]. Note also that we have worked harder on the estimation of I_1 to ensure that when $k = \Theta(n)$ the bound coincides with the standard bound $O(\sqrt{n})$.

5.1.1 Relative (p, ε) -approximations for halfspaces

We next turn the above construction of a spanning tree with small relative crossing number into a construction of a relative (p, ε) -approximation for a set of points in \mathbb{R}^d and for halfspace ranges. We base our construction on the machinery of Har-Peled and Sharir [HS11], thus extending their planar construction to higher dimensions.

Let P be a set of n points in \mathbb{R}^d , $d \geq 3$, and let \mathcal{T} be a spanning tree of P as provided in Theorem 5.3. We replace \mathcal{T} by a perfect matching M of P , with the same asymptotic relative crossing number, i.e., the number of pairs of M that are separated by a hyperplane of weight k is at most $O\left(n^{1-1/(d-1)} k^{1/d(d-1)} \log^{2/(d-1)}(n/k)\right)$. This is done in a standard manner — we first convert \mathcal{T} to a spanning path whose relative crossing number by any hyperplane is at most twice the crossing number of the same hyperplane with \mathcal{T} , and then pick every other edge of the path. To simplify the presentation, and to ensure that the resulting collection of edges is indeed a perfect matching, we assume that n is even.

We now construct a coloring of P with low discrepancy, by randomly coloring the points in each pair of M . Specifically, each pair is randomly and independently colored either as $-1, +1$ or as $+1, -1$, with equal probabilities. The standard theory of discrepancy (see [Cha01] and [Mat99]) yields the following variant.

Lemma 5.4 *Given a set P of n points in \mathbb{R}^d , $d \geq 3$, one can construct a coloring $\chi : P \mapsto \{-1, 1\}$, such that, for any halfspace H ,*

$$\chi(P \cap H) = \sum_{p \in P \cap H} \chi(p) = O\left(n^{(d-2)/2(d-1)} |P \cap H|^{1/2d(d-1)} \log^{(d+1)/2(d-1)} n\right).$$

The coloring is balanced — each color class consists of exactly $n/2$ points of P .

Proof. Indeed (see [Cha01, Mat99]), if the maximum crossing number is X then the above discrepancy can be bounded by $O(\sqrt{X \log n})$, and the asserted bound is then immediate by Theorem 5.3. \square

We also need the following technical lemma, taken from [HS11].

Lemma 5.5 *For any $x \geq 0$, $y > 0$, and $0 < p < 1$, we have $x^p < (x + y)/y^{1-p}$.*

As we next show, the improved discrepancy bound of Lemma 5.4 leads to an improved bound of the size of (ν, α) -samples for our range space, and, consequently, for the size of relative (p, ε) -approximations (with some constraints on the relationship between ε and p , as noted below). Let us introduce the following parameters:

$$\gamma = \frac{2d(d-1)-1}{(d-1)(d+1)}, \quad \mu = \frac{2d}{d+1}, \quad \text{and} \quad \eta = \frac{d}{d-1}.$$

Theorem 5.6 *Given a set P of n points in \mathbb{R}^d , $d \geq 3$, and parameters $0 < p < 1$ and $0 < \varepsilon < 1$, one can construct a relative (p, ε) -approximation $Z \subseteq P$ of size*

$$O\left(\frac{1}{\varepsilon^\mu p^\gamma} \log^\eta \frac{1}{\varepsilon p}\right).$$

We observe that, ignoring the power $\eta > 1$ of the logarithmic factor, this bound is an improvement of the bound in Theorem 1.3, as long as ε and p satisfy

$$\frac{1}{\varepsilon^\mu p^\gamma} < \frac{1}{\varepsilon^2 p}, \quad \text{or} \quad \varepsilon < p^{\frac{\gamma-1}{2-\mu}}.$$

Substituting the values of γ and μ , we have

$$\varepsilon < p^{\frac{d(d-2)}{2(d-1)}},$$

which is the required constraint on the relationship between ε and p , for which Theorem 5.6 does indeed yield an improvement (modulo the small "penalty" in the logarithmic factor). We also note that for $d = 2$ the bound in the theorem coincides with the bound in [HS11] (except for the power of the logarithmic factor).

Proof. Following one of the classical constructions of ε -approximations (see [Cha01]), we repeatedly halve P , until we obtain a subset of size as asserted in the theorem, and then argue that the resulting set has the desired approximation property. Formally, we set $P_0 = P$, and, at the i^{th} step, partition P_{i-1} into two equal halves, using Lemma 5.4; let P_i and P'_i denote the two halves (consisting of the points that are colored $+1$, -1 , respectively). We keep P_i , remove P'_i , and continue with the halving process. Let $n_i = |P|/2^i$ denote the size of P_i . For any halfspace H , we have

$$\begin{aligned} \left| |P_i \cap H| - |P'_i \cap H| \right| &\leq c_1 \cdot n_{i-1}^{(d-2)/2(d-1)} \cdot |P_{i-1} \cap H|^{1/2d(d-1)} \cdot \log^{\frac{d+1}{2(d-1)}} n_{i-1} \\ &\leq c_2 \cdot n_i^{(d-2)/2(d-1)} \cdot |P_{i-1} \cap H|^{1/2d(d-1)} \cdot \log^{\frac{d+1}{2(d-1)}} n_i, \end{aligned}$$

for appropriate constants c_1 and c_2 . Recalling our notation

$$\overline{P}_i(H) = \frac{|P_i \cap H|}{|P_i|} \quad \text{and} \quad \overline{P}'_i(H) = \frac{|P'_i \cap H|}{|P'_i|},$$

this can be rewritten as

$$\left| \overline{P}_i(H) - \overline{P}'_i(H) \right| \leq c_2 \frac{|P_{i-1} \cap H|^{1/2d(d-1)}}{n_i^{1-(d-2)/2(d-1)}} \log^{\frac{d+1}{2(d-1)}} n_i = c_3 \frac{\overline{P}_{i-1}(H)^{1/2d(d-1)}}{n_i^{(d+1)/2d}} \log^{\frac{d+1}{2(d-1)}} n_i,$$

for an appropriate constant c_3 . Since $|P_i| = |P'_i|$ and $P_{i-1} = P_i \cup P'_i$, we have

$$\overline{P_{i-1}}(H) = \frac{|P_{i-1} \cap H|}{|P_{i-1}|} = \frac{|P_i \cap H|}{2|P_i|} + \frac{|P'_i \cap H|}{2|P'_i|} = \frac{1}{2} \left(\overline{P_i}(H) + \overline{P'_i}(H) \right).$$

Combining together the last two relations, we have

$$|\overline{P_{i-1}}(H) - \overline{P_i}(H)| = \frac{|\overline{P_i}(H) - \overline{P'_i}(H)|}{2} \leq c_4 \frac{\overline{P_{i-1}}(H)^{1/2d(d-1)}}{n_i^{(d+1)/2d}} \log^{2(d-1)} n_i,$$

where $c_4 = c_3/2$. Applying Lemma 5.5, with $p = 1/2d(d-1)$, $x = \overline{P_{i-1}}(h)$, and $y = p$, the last expression is bounded by

$$\frac{c_4 \cdot \log^{2(d-1)} n_i}{n_i^{(d+1)/2d}} \cdot \frac{\overline{P_{i-1}}(H) + p}{p^{1-1/2d(d-1)}} \leq \frac{c_4 \cdot \log^{2(d-1)} n_i}{n_i^{(d+1)/2d} p^{1-1/2d(d-1)}} (\overline{P_{i-1}}(H) + \overline{P_i}(H) + p).$$

This implies that, in the notation of Section 1.2,

$$d_p(\overline{P_{i-1}}(H), \overline{P_i}(H)) \leq \frac{c_4 \cdot \log^{2(d-1)} n_i}{n_i^{(d+1)/2d} p^{1-1/2d(d-1)}}.$$

Since d_p is a metric, the triangle inequality implies that

$$\begin{aligned} d_p(\overline{P}(H), \overline{P_i}(H)) &\leq \sum_{k=1}^i d_p(\overline{P_{k-1}}(H), \overline{P_k}(H)) \leq \frac{c_4}{p^{1-1/2d(d-1)}} \sum_{k=1}^i \frac{\log^{2(d-1)} n_k}{n_k^{(d+1)/2d}} \\ &= O \left(\frac{\log^{2(d-1)} n_i}{p^{1-1/2d(d-1)} n_i^{(d+1)/2d}} \right). \end{aligned}$$

We can therefore ensure that $d_p(\overline{P}(H), \overline{P_i}(H)) < \varepsilon$, for every halfspace H , provided that $n_i = \Omega \left(\frac{1}{\varepsilon^\mu p^\gamma} \log^\eta \frac{1}{\varepsilon p} \right)$. Taking Z to be the smallest P_i which still satisfies this size constraint, i.e., of size $\Theta \left(\frac{1}{\varepsilon^\mu p^\gamma} \log^\eta \frac{1}{\varepsilon p} \right)$, we have thus shown that Z is a (p, ε) -sample. As follows from the equivalence between (p, ε) -samples and relative (p, ε) -approximations (see [HS11] and Theorem 1.3), Z is also a relative (p, ε) -approximation, which completes the proof. \square

Remark. Very recent work by Ezra [Ez14] obtains improved bounds for the size of relative approximations in a more general setting, which also includes the case of halfspace ranges in higher dimensions. For this latter setup, the bound obtained in [Ez14] is

$$\max \left\{ O \left(\log^{\frac{3d+1}{d+1}} n \right), O \left(\frac{\log^{\frac{3d+1}{d+1}} \frac{1}{\varepsilon p}}{p^{\frac{d+1}{d+1}} \varepsilon^{\frac{2d}{d+1}}} \right) \right\},$$

where the power of $1/p$ is smaller than the one that we have obtained.

5.2 Output-sensitive halfspace range counting in \mathbb{R}^d

Our second main application is an improved, output-sensitive algorithm for halfspace range counting in \mathbb{R}^d , for any $d \geq 3$. That is, we are given a set P of n points in \mathbb{R}^d , and want

to preprocess it into a data structure that can answer efficiently queries, in which we are given a halfspace H , and we wish to count $|P \cap H|$ (exactly). We will only consider the case where we seek a data structure with near-linear storage. Our goal is an algorithm whose query performance is sensitive to the output size $w_H = |P \cap H|$.

Before describing our solution, we recall that when the output size w_H is reasonably small, we can trivially turn the halfspace range reporting algorithm of Matoušek [Mat92b] into a range counting algorithm. This algorithm uses a data structure with $O(n \log \log n)$ storage, which can be constructed in $O(n \log n)$ time, and a halfspace range reporting (counting) query with a halfspace H can be answered in time $O(n^{1-1/\lfloor d/2 \rfloor} \log^{O(1)} n + w_H)$.² However, when w_H is large, this becomes quite inefficient. Our solution will use Matoušek’s algorithm as one of its components, but its novel contribution is for the case where w_H is large; see below for details.

Here is an informal overview of the algorithm. We partition P into a logarithmic number of subsets P_1, \dots, P_l , so that, for each i , P_i consists of points which are $O(k_i)$ -shallow in $P \setminus (P_1 \cup \dots \cup P_{i-1})$, where k_1 is specified below, and $k_{i+1} = 2k_i$, for $i \geq 1$. This partition, as we show later, has the property that a halfspace H of weight w_H will miss the convex hulls of all the sets P_i , for which $k_i \geq 2w_H$. Handling sets P_i with $k_i < 2w_H$ is done using the partition data structure for shallow sets, as provided in Theorem 4.3. This will result in a query time bound which depends on w_H in roughly the same manner as the bounds in Theorem 5.3. Specifically, we show:

Theorem 5.7 *Given a set P of n points in \mathbb{R}^d , $d \geq 3$, one can construct a data structure for the halfspace range counting problem for P , so that, for any query halfspace H , the number $w_H = |P \cap H|$ of points of P in H can be counted, for $d \geq 4$, in time*

$$Q(n) = \begin{cases} O\left(n^{1-\lfloor d/2 \rfloor} \log^{O(1)} n + w_H\right), & \text{for } w_H \leq c_1 n^{1-(d-1)/(d(d-1)-1)}, \\ O\left(n^{1-1/(d-1)} w_H^{1/d(d-1)} \log^{O(1)} n\right), & \text{otherwise,} \end{cases} \quad (13)$$

for an appropriate constant c_1 . For $d = 3$, the first bound for the query time is $O(\log n + w_H)$, and the second bound remains the same.

The data structure uses $O(n \log \log n)$ storage, and preprocessing time

$$T(n) = \begin{cases} O\left(n^{1+\frac{(d-2)(d-1)}{2(d(d-1)-1)}} \log^{O(1)} n\right), & \text{for } d \text{ even,} \\ O\left(n^{1+\frac{(d-3)(d-1)}{2(d(d-1)-1)}} \log^{O(1)} n\right), & \text{for } d \text{ odd,} \end{cases}$$

Remarks: (i) To get some feeling for these bounds, we note that for $d = 3$ the query time is $O(\log n + w_H)$ for $w_H \leq c_1 n^{3/5}$ and $O(n^{1/2} w_H^{1/6} \log^{O(1)} n)$ otherwise; ignoring the logarithmic factors, both bounds balance out for $w_H = \Theta(n^{3/5})$. For $d = 4$, the query time is $O(n^{1/2} \log^{O(1)} n + w_H)$ for $w_H \leq c_1 n^{8/11}$ and $O(n^{2/3} w_H^{1/12} \log^{O(1)} n)$ otherwise. These bounds should be compared with the respective standard bounds $O(n^{2/3} \log^{O(1)} n)$ and $O(n^{3/4} \log^{O(1)} n)$ for $d = 3$ and $d = 4$ of [Mat92a]; our bounds coincide with these “insensitive” bounds when $w_H = \Theta(n)$. A graphical illustration that shows this improvement for

²Note that, if we allow subtraction, we can achieve the same query cost when the query halfspace is the complementary one, which contains $n - w_H$ points.

$d = 4$ is given in Figure 3.

(ii) If $w_H > n/2$, denoting \overline{H} as the complementary halfspace of H , we obtain the count by counting $n - |P \cap \overline{H}|$ (see also a preceding footnote). In this case the upper bound on the query time is obtained by replacing w_H with $n - w_H$ in the above bounds.

(iii) Clearly, a weak feature of the theorem is that the bound on the preprocessing cost is not near-linear. It approaches $O(n^{3/2})$ when d is very large. For small values of d , it is near linear for $d = 3$ and is roughly $O(n^{14/11})$ for $d = 4$.

Proof. We first describe the data structure, then the procedure for answering range counting queries and its analysis, then we analyze the storage used by the structure, and finally present and analyze the (somewhat involved) construction of the data structure.

Data structure. We first construct an auxiliary data structure, based on Matoušek's range reporting mechanism [Mat92b], which allows us to report the points of P in a query halfspace H in time $O(n^{1-1/\lfloor d/2 \rfloor} \log^{O(1)} n + w_H)$, where $w_H = |P \cap H|$. This structure uses $O(n \log \log n)$ storage and can be constructed in $O(n \log n)$ time.

We next partition P into a logarithmic number of subsets P_1, \dots, P_l . For each i , P_i consists of points which are $(c'_i k_i)$ -shallow in P'_{i-1} , where $P'_0 = P$, $P'_i = P \setminus (P_1 \cup \dots \cup P_i)$, for $1 \leq i \leq l-1$, and the values of the constant parameters c'_1, \dots, c'_l will be set later. We set $k_{i+1} = 2k_i$, for $i \geq 1$, and set

$$k_1 = c_1 n^{1 - \frac{d-1}{d(d-1)-1}}, \quad (14)$$

for an appropriate constant c_1 . With each P_i we associate an auxiliary halfspace emptiness data structure, due to Matoušek [Mat92b]. Putting $m_i = |P_i|$, the structure uses linear storage, requires $O(m_i^{1+\delta})$ preprocessing time, where $\delta > 0$ is arbitrarily small but fixed, and can test whether a query halfspace contains any point of P_i , in time $O(m_i^{1-1/\lfloor d/2 \rfloor} 2^{O(\log^* m_i)})$.

Next, as in the spanning tree construction (Theorem 5.3), for each set P_i , we construct a simplicial partition Π_i , using Theorem 4.3, with $k = c'_i k_i$, in overall time $O(n \log n)$. With each class of each Π_i , consisting of a simplex Δ and a subset P_Δ of the corresponding P_i , we associate another auxiliary data structure, based on the standard simplex range counting technique of Matoušek [Mat92a]. Putting $m_\Delta = |P_\Delta|$, this structure uses $O(m_\Delta)$ storage and $O(m_\Delta \log m_\Delta)$ preprocessing time, and can count the number of points of P_i in a query halfspace in time $O(m_\Delta^{1-1/d} \log^{O(1)} m_\Delta)$. This completes the description of our data structure.

Answering a counting query. Given a query halfspace H' , our range counting algorithm proceeds as follows. We first run an emptiness test on each of the sets P_i , in ascending order of their indices, with respect to both H' and its complementary halfspace $\overline{H'}$, using the auxiliary emptiness data structures associated with those sets. Assume that the test first returns **true** at iteration m , for some $m \geq 1$. We then set H to be the halfspace for which the emptiness test returned **true**, and we refer to it as the query halfspace from now on. (If the original halfspace was the complement of H , we return the desired count by subtracting ³ $|P \cap H|$ from n , as above.) Having a logarithmic number of sets, this step takes at most $O(n^{1-1/\lfloor d/2 \rfloor} \log^{O(1)} n)$ time.

³If for some reason subtraction is not allowed in the model of computation, we apply the querying procedure only to the input halfspace H' . We will then miss the opportunity to answer the query more efficiently when $n - w_H$ is small.

Our analysis is based on Lemma 5.9, given below, from which it follows that H is disjoint from the convex hulls of all the sets P_j , for $j > U = \lceil \log(w_H/k_1) \rceil + 1$, and so it is guaranteed that $m \leq U + 1$. We distinguish between two cases. First suppose that $m \leq 2$. If $m = 1$, we have $P \cap H = \emptyset$, and thus we are done. If $m = 2$, we have $|P \cap H| \leq c'_1 k_1$, and we use the reporting data structure associated with the full set P , to count the points in $P \cap H$ in time

$$O\left(n^{1-1/\lfloor d/2 \rfloor} \log^{O(1)} n + |P \cap H|\right),$$

which establishes the first bound in (13). Assume then that $m > 2$. It suffices to proceed only with the sets P_1, \dots, P_{m-1} . For each such set P_i , with an associated simplicial partition Π_i , we check, for each class $(P_\Delta, \Delta) \in \Pi_i$, whether the simplex Δ is fully contained in H or is crossed by its bounding hyperplane h . For a simplex Δ that is fully contained in H , we add $|P_\Delta|$ to the output count, whereas for Δ that is crossed by h , we use the auxiliary range counting data structure of [Mat92a], associated with that node, to count $|P_\Delta \cap H|$. Repeating this procedure for P_1, \dots, P_{m-1} , and adding up the resulting counts, we obtain the desired count $|P \cap H|$.

We now examine the value of k_1 , as set in (14), and note that, for any halfspace H with $w_H = \Theta(k_1)$, the upper bound on the cost of answering the query with H using the range reporting data structure is roughly the same, up to polylogarithmic factor, as the cost of answering the query using the range counting procedure at each of the simplices crossed by h . Indeed, the former cost is $O(n^{1-1/\lfloor d/2 \rfloor} \log^{O(1)} n + k_1)$, for $d \geq 4$, and the latter cost is

$$O\left(\left(\frac{n}{k_1}\right)^{1-1/(d-1)} \cdot k_1^{1-1/d} \log^{O(1)} n\right) = O\left(n^{1-1/(d-1)} k_1^{1/d(d-1)} \log^{O(1)} n\right).$$

This follows because (i) for such values of w_h we only need to query in P_1 , (ii) the number of simplices of Π_1 crossed by the plane h bounding H is $O\left(\left(\frac{n}{k_1}\right)^{1-1/(d-1)} \log^{O(1)} n\right)$ by Theorem 4.3, and (iii) the cost of a query within each of the crossed classes is $O(k_1^{1-1/d} \log^{O(1)} k_1)$. As is easily checked, our k_1 does indeed satisfy

$$n^{1-1/\lfloor d/2 \rfloor} \log^{O(1)} n < k_1 = \Theta\left(n^{1-1/(d-1)} k_1^{1/d(d-1)}\right),$$

which makes the two bounds roughly the same, up to a polylogarithmic factor. For $d = 3$, the first bound is $O(\log n + k_1)$, and the same relationship holds. For larger values of w_H , i.e., for $w_H > k_1$, recalling that $U = \lceil \log(w_H/k_1) \rceil + 1$, we use the second part of the data structure, recalling that, as above, h crosses only $O\left(\left(\frac{n}{k_i}\right)^{1-1/(d-1)} \log^{O(1)}\left(\frac{n}{k_i}\right)\right)$ simplices of Π_i (see Theorem 4.3), to obtain an overall cost of

$$\begin{aligned} O\left(\sum_{i=1}^U \left(\frac{n}{k_i}\right)^{1-1/(d-1)} k_i^{1-1/d} \log^{O(1)} n\right) &= O\left(\sum_{i=1}^U n^{1-1/(d-1)} k_i^{1/d(d-1)} \log^{O(1)} n\right) \\ &= O\left(n^{1-1/(d-1)} w_H^{1/d(d-1)} \log^{O(1)} n\right). \end{aligned}$$

This establishes the second bound on the query time, as given in (13).

Comparing these bounds with the preceding techniques [Mat92a, Mat92b], we note that for $k \leq k_1$ we do not obtain any improvement — we use in fact the same algorithm of [Mat92b]. However, for $k > k_1$, the bound on $Q(n)$ is smaller than the “insensitive” bound $O(n^{1-1/d} \log^{O(1)} n)$ of [Mat92a]. Again, see Figure 3 for graphical illustrations of our

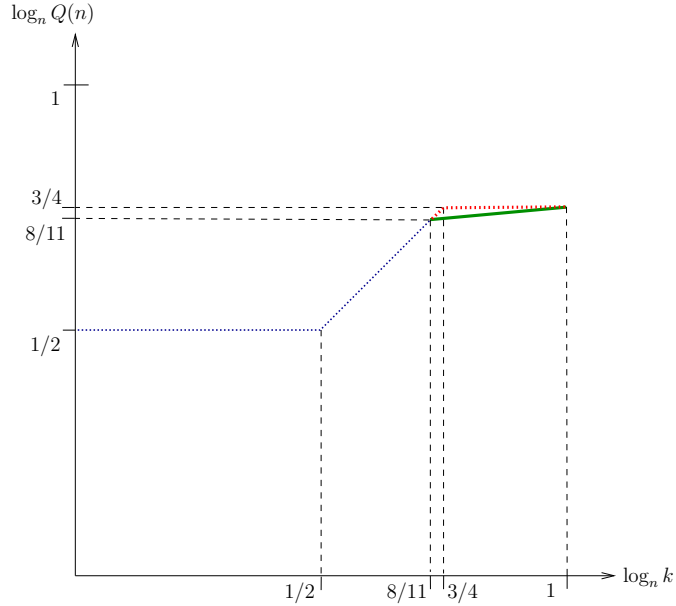


Figure 3: Graphical illustration of the output-sensitive improvement for $d = 4$.

improvement. (Note that when $k = \Theta(n)$ our algorithm has the same asymptotic bound as Matoušek’s counting algorithm [Mat92a].)

Storage. The auxiliary reporting data structure stored for the entire P uses $O(n \log \log n)$ storage; see [Mat92b]. The emptiness data structures for each of the subsets P_i are of size linear in the sizes of the respective subsets [Mat92a], and so the total space that they consume is $O(n)$. The simplicial partitions constructed for the sets P_1, \dots, P_l consume a total of $O(n)$ storage, including the storage used by Matoušek’s standard range counting structure [Mat92a] within each simplex. Hence, the overall storage is $O(n \log \log n)$, as asserted.

Preprocessing time. One of the main steps in the preprocessing (which is not required for the standard data structures of [Mat92a, Mat92b]) is the construction of the subset of k -shallow points of a set P of n points in \mathbb{R}^d , a step that we perform iteratively, to construct the sequence of sets P_1, \dots, P_l . This step is involved (and costly) in higher dimensions, but we note that we only need to handle values of k satisfying

$$k > k_1 = c_1 n^{1-(d-1)/(d(d-1)-1)},$$

because for smaller values of k we use instead the alternative range reporting machinery of [Mat92b]. For such values of k we have

$$\frac{n}{k} = O\left(n^{(d-1)/(d(d-1)-1)}\right), \quad (15)$$

a property that will be crucial in guaranteeing the (relative) efficiency of the preprocessing.

So let P , n and k be fixed. The construction proceeds as follows. Put $r = n/k$ and choose a random sample R of $c_2 r \log r$ points of P , for an appropriate sufficiently large constant c_2 . Let K denote the interior of the intersection of all halfspaces that contain at

least $|R| - c_3 \log r$ points of R , for an appropriate constant c_3 . By the centerpoint theorem (see [Mat02]), K is nonempty for $c_3 \log r < |R|/(d+1)$, which always hold for c_2 sufficiently large.

Before we describe and analyze the construction of K , we establish the following properties.

Lemma 5.8 *Let P , R , K and k be as above, and put $t = c_3 \log(n/k) = c_3 \log r$. Then an appropriate choice of the constants c_2 and c_3 guarantees that, with high probability, the following properties hold.*

- (i) *Every k -shallow hyperplane in P is t -shallow in R .*
- (ii) *Every hyperplane which is $(t+d)$ -shallow in R is $O(k)$ -shallow in P .*
- (iii) *Every point in $P \setminus K$ is $O(k)$ -shallow in P .*
- (iv) *Every point of P which is k -shallow in P lies in the complement \overline{K} of K .*

Proof. (i) For c_2 sufficiently large, R is a shallow $(1/r)$ -net with high probability (see Section 1.2 for the definition and properties of shallow nets). Since $r = 1/\varepsilon = n/k$, the claim follows from Property (ii) of shallow $(1/r)$ -nets, for an appropriate constant c_3 .

(ii) This follows from Property (i) of shallow $(1/r)$ -nets.

(iii) Let $p \in P \setminus K$, and let H be a closed halfspace containing p , disjoint from K , and parallel to a closed halfspace H_0 , which supports one of the open facets of K (H_0 might coincide with H). By the definition of K , H_0 contains at least $|R| - t$ points of R . Since at most d points lie on its bounding hyperplane, its complementary closed halfspace $\overline{H_0}$ contains at most $t + d$ points of R , and thus, by Property (ii), it is $O(k)$ -shallow in P , implying that p is $O(k)$ -shallow in P , as claimed.

(iv) We argue that no point $p \in P \cap K$ can be k -shallow in P . Indeed, let p be such a point and let H be a closed halfspace, whose bounding hyperplane h contains p . By construction, since h intersects K , it follows that $|H \cap R| > t$. By Property (i), we have $|H \cap P| > k$. Since this holds for every halfspace H containing p , p is not k -shallow in P , a contradiction which implies the claim. \square

We next derive the following lemma, promised earlier, which is crucial for the efficiency of the query mechanism, and whose proof is based on the properties established in Lemma 5.8.

We define the *depth* of a point $p \in \mathbb{R}^d$ as the minimum number of points of P contained in a closed halfspace whose bounding hyperplane passes through p .

Lemma 5.9 *Let P , its partition into P_1, \dots, P_l , and the parameters k_1, \dots, k_l be as above. Let H be a halfspace in \mathbb{R}^d of weight w_H . Then H misses the convex hulls of the sets P_2, \dots, P_l , if $w_H \leq k_1$, and of the sets P_j , for $j > U = \lceil \log(w_H/k_1) \rceil + 1$, otherwise.*

Proof. First assume that $w_H \leq k_1$. Since the points in $P \setminus P_1$ are at least $(k_1 + 1)$ -deep, H cannot contain any of them, and it therefore misses the convex hulls of the sets P_2, \dots, P_l , as claimed. We may thus assume that $w_H > k_1$, and so $U \geq 2$. Assume that H crosses the convex hull of a set P_j , for $j > U$. Thus, necessarily, there exists a point $q \in P_j$ which is

w_H -shallow. On the other hand, by Lemma 5.8 (iv), P_j consists of points which are at least $(k_{j-1} + 1)$ -deep in P'_{j-1} , and clearly so they are in P . Thus the depth of q is at least

$$k_{j-1} + 1 > 2^{j-2} k_1 \geq 2^{(\lceil \log(w_H/k_1) \rceil + 2) - 2} \cdot k_1 \geq w_H.$$

That is, q is at least $(w_H + 1)$ -deep in P , in contradiction to its being w_H -shallow. \square

Suppose that we have a procedure for constructing K . Then the construction of the sets P_i proceeds as follows. Starting with the initial set P , put $\tilde{P} = P \setminus K$. Lemma 5.8 (iii) implies that all the points of \tilde{P} are $O(k)$ -shallow in P and thus also in \tilde{P} . We now iterate this process, as follows. At step i we apply this construction to the subset P'_{i-1} of the remaining points of P (with $P'_0 = P$), and take the resulting set \tilde{P} as the next set P_i in the sequence, setting $P'_i = P'_{i-1} \setminus P_i$; that is, P'_i is the portion of P'_{i-1} inside the corresponding convex region K . We stop the process when $P'_i = \emptyset$ or when $|P'_i|/k_i$ drops below a sufficiently large constant, according to the Partition Theorem (Theorem 4.3), which happens after $O(\log(n/k_1))$ steps. If $P'_i \neq \emptyset$ upon termination, we put $P_{i+1} = P'_i$.

We finally describe the construction of K , thus concluding the description of the new aspects of our data structure construction. We pass to the dual space, where each point $p \in P$ is mapped to a hyperplane p^* , using the standard duality which also preserves the above/below relationships (see, e.g., [Ed87]). Consider the arrangement $\mathcal{A} = \mathcal{A}(R^*)$ of the hyperplanes dual to the points of the sample R . A t -shallow hyperplane q in the primal space is mapped to a point q^* which lies at level $\leq t$ in \mathcal{A} ; more precisely, its level is either among the t lowest levels of \mathcal{A} or among its t highest levels. We denote the former (resp., latter) collection of levels as $\mathcal{A}_{\leq t}^-(R^*)$ (resp., $\mathcal{A}_{\leq t}^+(R^*)$), but we use just $\mathcal{A}_{\leq t}^-$ (resp., $\mathcal{A}_{\leq t}^+$) to simplify the notation. As is well known [Mat91b] (and easy to verify), \tilde{K} is mapped in the dual space to the (open) region K^* enclosed between the upper convex hull of $\mathcal{A}_{\leq t}^-$, and the lower convex hull of $\mathcal{A}_{\leq t}^+$, in the sense that a point v is in K if and only if its dual hyperplane v^* is contained in K^* .

We note that, by Clarkson and Shor [CS89], the combinatorial complexity of $\mathcal{A}_{\leq t}^-$ and $\mathcal{A}_{\leq t}^+$ is at most

$$\begin{aligned} u &= O\left(|R|^{\lfloor d/2 \rfloor} t^{\lceil d/2 \rceil}\right) = O\left(|R|^{\lfloor d/2 \rfloor} \log^{O(1)} n\right) \\ &= O\left((r \log r)^{\lfloor d/2 \rfloor} \log^{O(1)} n\right) = O\left((n/k)^{\lfloor d/2 \rfloor} \log^{O(1)} n\right). \end{aligned}$$

Using (15) we have

$$(n/k)^{\lfloor d/2 \rfloor} = O\left(n^{\frac{(d-1)\lfloor d/2 \rfloor}{d(d-1)-1}}\right),$$

and the exponent in the right-hand side is smaller than $1/2$ when d is odd, and only slightly larger than $1/2$ when d is even, for $d \geq 3$. More precisely, u satisfies

$$u = \begin{cases} O\left(n^{\frac{1}{2}\left(1+\frac{1}{d(d-1)-1}\right)} \log^{O(1)} n\right), & \text{for } d \text{ even,} \\ O\left(n^{\frac{1}{2}\left(1-\frac{d-2}{d(d-1)-1}\right)} \log^{O(1)} n\right), & \text{for } d \text{ odd,} \end{cases} \quad (16)$$

for any $d \geq 3$.

Agarwal et al. present in [AdBMS98] a randomized algorithm for computing the $(\leq t)$ -level of an arrangement of n hyperplanes in \mathbb{R}^d within expected time $\Theta(u)$.

Let V^- (resp., V^+) denote the set of vertices of $\mathcal{A}_{\leq t}^-$ (resp., $\mathcal{A}_{\leq t}^+$). We preprocess each of V^- , V^+ for halfspace emptiness queries, using the algorithm of [Mat92b]. Then, for each $p \in P$, we test whether $p \in K$ by testing whether the upper halfspace bounded by p^* does not contain any point of V^- , and the complementary lower halfspace does not contain any point of V^+ . All the points that fail the test are placed in \tilde{P} , and those that pass it are passed to the next iteration. In this way, we report the points in \tilde{P} in expected time

$$O\left(n \cdot u^{1-1/\lfloor d/2 \rfloor} \log^{O(1)} u\right).$$

Recalling the bound (16) for u , for d even we report those points in

$$O\left(n^{1+\frac{(d-2)(d-1)}{2(d(d-1)-1)}} \log^{O(1)} n\right)$$

expected time, and for d odd in

$$O\left(n^{1+\frac{(d-3)(d-1)}{2(d(d-1)-1)}} \log^{O(1)} n\right)$$

expected time. Having a logarithmic number of steps, we construct all the sets P_i within the same asymptotic expected bound. We note that the overall time consumed by the construction of the other data structure components in each step is smaller than the bound above. We omit the description of these other steps of the construction, as they are identical to those given in [Mat92a, Mat92b]. This completes the proof of Theorem 5.7. \square

References

- [AdBMS98] P. K. Agarwal, M. de Berg, J. Matoušek, and O. Schwarzkopf, Constructing levels in arrangements and higher order Voronoi diagrams, *SIAM J. Comput.*, 27(3):657–667, 1998.
- [AM95] P. K. Agarwal and J. Matoušek, Dynamic half-space range reporting and its applications, *Algorithmica*, 13:325–345, 1995.
- [AMS98] P. K. Agarwal, J. Matoušek, and O. Schwarzkopf, Computing many faces in arrangements of lines and segments, *SIAM J. Comput.*, 27(2):491–505, 1998.
- [APS93] B. Aronov, M. Pellegrini, and M. Sharir, On the zone of a surface in a hyperplane arrangement, *Discrete Comput. Geom.*, 9(2):177–186, 1993.
- [AS92] N. Alon and J. H. Spencer, *The Probabilistic Method*, Publisher, City, 1992. Appendix A: *Bounding of Large Deviations*, pp. 233–240.
- [Ch12] T. M. Chan, Optimal partition trees, *Discrete Comput. Geom.*
- [CF90] B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica* 10 (3):229–249, 1990.
- [Cha01] B. Chazelle, *The Discrepancy Method: Randomness and Complexity*, Cambridge University Press, New York, 2001.

- [Cha93] B. Chazelle, An optimal convex hull algorithm in any fixed dimension, *Discrete Comput. Geom.*, 10:377–409, 1993.
- [CW89] B. Chazelle and E. Welzl, Quasi-optimal range searching in spaces with finite VC-dimension, *Discrete Comput. Geom.*, 4:387–421, 1989.
- [CS89] K. L. Clarkson and P. W. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.*, 4:467–490, 1989.
- [Ed87] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.
- [Ez14] E. Ezra, A size-sensitive discrepancy bound for set systems of bounded primal shatter dimension, *Proc. ACM–SIAM Annu. Sympos. Discrete Algorithms*, 2014, to appear. Also available at <http://www.cims.nyu.edu/~esther/Publications/discrepancy.pdf>
- [HS11] S. Har-Peled and M. Sharir, Relative (p, ε) -approximations in geometry, *Discrete Comput. Geom.* 45:462–496, 2011.
- [HW87] D. Haussler and E. Welzl, Epsilon-nets and simplex range queries, *Discrete Comput. Geom.*, 2:127–151, 1987.
- [LLS01] Y. Li, P. M. Long, and A. Srinivasan, Improved bounds on the sample complexity of learning, *J. Comput. Syst. Sci.*, 62:516–527, 2001.
- [Mat91a] J. Matoušek, Cutting hyperplane arrangements, *Discrete Comput. Geom.*, 6(5):385–406, 1991.
- [Mat91b] J. Matoušek, Computing the center of planar point sets, in *Computational Geometry: Papers from the DIMACS Special Year* (J. E. Goodman, R. Pollack and W. Steiger, eds.), AMS, Providence, RI, pp. 221–230, 1991.
- [Mat92a] J. Matoušek, Efficient partition trees, *Discrete Comput. Geom.*, 8(3):315–334, 1992.
- [Mat92b] J. Matoušek, Reporting points in halfspaces, *Comput. Geom. Theory Appl.*, 2(3):169–186, 1992.
- [Mat02] J. Matoušek, *Lectures on Discrete Geometry*, Springer-Verlag, Berlin, 2002; Centerpoint theorem: pp. 14–15, Cuttings: pp. 152–163, VC-Dimension: pp. 237–243.
- [Mat99] J. Matoušek, *Geometric Discrepancy*, Algorithms and Combinatorics, Vol. 18, Springer Verlag, Heidelberg, 1999.
- [ShSh11] M. Sharir and H. Shaul, Semi-algebraic range reporting and emptiness searching with applications, *Siam J. Comput.*
- [Wel92] E. Welzl, On spanning trees with low crossing numbers, in *Data Structures and Efficient Algorithms, Final Report on the DFG Special Joint Initiative*, volume 594 of *Lect. Notes in Comp. Sci.*, Springer-Verlag, Heidelberg, pp. 233–249, 1992.