Reporting points in halfspaces*

Jiří Matoušek

Department of Applied Mathematics

Charles University

Malostranské nám. 25, 118 00 Praha 1, Czechoslovakia

Abstract

We consider the *halfspace range reporting* problem: given a finite set P of points in R^d, preprocess it so that given a query halfspace γ, the points of $P \cap \gamma$ can be reported efficiently. We show that with almost linear storage, this problem can be solved substantially more efficiently that the more general simplex range searching problem. We give a data structure for halfspace range reporting in dimensions $d \geq 4$ with $O(n \log \log n)$ space, $O(n \log n)$ deterministic preprocessing time and $O(n^{1-1/\lfloor d/2 \rfloor}(\log n)^c + k)$ query time, where $c = c(d)$ is a constant and $k = |P \cap \gamma|$ (efficient solutions were known for $d = 2, 3$). For the *halfspace emptiness problem*, where we only want to know whether $P \cap \gamma = \emptyset$, we can achieve query time $O(n^{1-1/\lfloor d/2 \rfloor} 2^{c' \log^* n})$ with a linear space and $O(n^{1+\delta})$ preprocessing ($c' = c'(d)$ is a constant and $\delta > 0$ is arbitrarily small but fixed).

1 Introduction

One of central themes in computational geometry is the development of efficient *range searching algorithms*. We will consider the *halfspace range reporting*: Preprocess a set P of n points in R^d so that, given any query halfspace γ, the points in $P \cap \gamma$ can be reported efficiently. (As it is typical in computational geometry, we will consider the space dimension d as a small fixed integer, thus something depending on the dimension only will be a constant for us.) This problem is a special case of a more general *simplex range searching problem*, where we want to report all points of P contained in a query simplex σ, or (still more generally) to evaluate $\sum_{p \in P \cap \sigma} w(p)$, where $w : P \to S$ is a function assigning weights belonging to some semigroup S to the points. The general simplex range searching problem is now probably close to being solved optimally.

In a *range reporting* problem, we want to list all points contained in a query range. Chazelle [Cha86] proposed a strategy for reporting problems, called "filtering search", whose basic observation is the following: If we report k points of the answer, we must in any case spend $\Omega(k)$ time on this, and thus we can spend $O(k)$ more time on computing the answer without affecting the asymptotic efficiency. Hence for a large k, we may use a less efficient procedure

*A preliminary version appeared in Proc. 32. IEEE Symposium on Foundations of Computer Science (1991), pages 207–215.

than for a small k. With such an approach, several authors obtained solutions for the halfspace range reporting problem which are more efficient than the ones for the simplex range searching ([CGL85], [CP86], [AHT90], [Cla88]). Among these, only Clarkson's solution works for dimensions $d \geq 4$, and it gives query time $O(\log n+k)$ (k is the number of reported points), using $O(n^{\lfloor d/2 \rfloor+\delta})$ space[1]. It was not known whether one can improve over the simplex range searching results also in the case of linear or almost linear space for any dimension $d \geq 4$. In this paper we provide such an improvement, in any fixed dimension.

Before we state our results, let us briefly recall the known results on the above mentioned range searching problems.

The efficiency of a range searching algorithm depends on the amount of space (and also of preprocessing time) allowed for the data structure. For the simplex range searching problem, Chazelle [Cha89] established a lower bound for this tradeoff. This lower bound is valid only under some restrictions on the kind of algorithm used, but so far all known algorithms satisfy these restrictions and there does not seem to be much hope for circumventing the lower bound. The lower bound is as follows: In dimension d and given m units of storage, the worst-case query time is at least $\Omega((n/\log n)/m^{1/d})$. In the plane, the bound sharpens to $\Omega(n/\sqrt{m})$; this indicates that for higher dimensions, the logarithmic factor might be only a product of the proof technique.

This tradeoff has nearly been attained by Chazelle et al. [CSW90]. An improved and simpler solution was given in [Mat91b], where it is shown that that the simplex range searching problem can be solved with $O(n \log n)$ deterministic preprocessing time, using $O(n)$ space and with $O(n^{1-1/d}(\log n)^b)$ query time, $b = b(d)$ a constant. The simplex range reporting problem can be solved within the same bounds, with the number of reported points k added to the query time. Another simplex range searching algorithm, which approaches the lower bound still more tightly, was given in [Mat92]. Of the rich previous literature concerning the simplex range searching problem, let us mention e.g. [Wil82], [YY85], [HW87], [CW89].

The halfspace range reporting problem was solved optimally in dimension 2 by Chazelle et al. [CGL85]. With $O(n)$ space and $O(n \log n)$ (deterministic) preprocessing time, one can do the reporting in time $O(\log n + k)$, where k is the number of points reported. For dimension 3, a solution with nearly linear space and $O(\log n + k)$ query time was given by Chazelle and Preparata [CP86], and Aggarwal et at. [AHT90] give an algorithm with the same query time and $O(n \log n)$ space (but the preprocessing time is rather high, about cubic).

As mentioned above, Clarkson [Cla88] showed that with $O(n^{\lfloor d/2 \rfloor+\delta})$ space and the same order of magnitude of expected preprocessing time one can answer a query in $O(\log n + k)$ time. A heuristic reason which we will indicate below suggests that this amount of space might be close to optimal for this query time. We give the following counterpart for almost linear space:

Theorem 1.1 *Given a set P of n points in R^d, $d \geq 4$, one can build, in $O(n \log n)$ time, a data structure for the halfspace range reporting problem for P with space $O(n \log\log n)$ and query time $O(n^{1-1/\lfloor d/2 \rfloor}(\log n)^c + k)$, where k is the number of reported points and $c = c(d)$ is a constant (depending on the dimension).*

[1]Throughout this paper, δ stands for an arbitrary small but fixed positive real number. The constants in the asymptotic bounds may depend on the choice of δ.

Using a combination of this result with the Clarkson's data structure, we can achieve the following tradeoff:

Corollary 1.2 *Given a set of n points in R^d, $d \geq 4$, and a parameter m, $n \leq m \leq n^{\lfloor d/2 \rfloor}$, the halfspace range reporting problem can be solved with space and preprocessing time $O(m^{1+\delta})$ and query time $O(\frac{n}{m^{1/\lfloor d/2 \rfloor}} \log n + k)$.*

One might ask whether a similar improvement over the general simplex range searching result could be obtained for the simplex range reporting. At least for the pointer machine model of computation, the answer is negative: Chazelle and Rosenberg [CR92] show that the reporting in simplices (or even in slabs bounded by two parallel hyperplanes) is roughly equally difficult as the general simplex range searching problem. Also, Brönninmann and Chazelle [BC92] prove a lower bound for the general halfspace range searching, which is somewhat weaker than the bound for the simplex range searching, but substantially higher than the upper bound for the halfspace range reporting in the above theorem.

A special case of the reporting problem is an *emptiness* problem, where we only want to decide whether a query halfspace contains some point of P (interestingly, most of known applications of the halfspace range reporting use only this special case). For this problem, we can improve the query time in Theorem 1.1 a little:

Theorem 1.3 *Given a set of n points in R^d, $d \geq 4$, one can build in time $O(n^{1+\delta})$ a linear-size data structure, which can decide whether a query halfspace contains a point of P in time $O(n^{1-1/\lfloor d/2 \rfloor} 2^{c' \log^* n})$.*

A small improvement over Clarkson's data structure for the halfspace emptiness problem was recently achieved by Schwarzkopf [Sch92] (using some results of the present paper), who gave an algorithm with $O(\log n)$ query time and $O(n^{\lfloor d/2 \rfloor} / \log^{\lfloor d/2 \rfloor - \delta})$ space. The methods can also be used to get a small improvement of space for the halfspace range reporting with $O(\log n + k)$ query time, and thus improve Corollary 1.2.

It is well-known that the convex hull of an n-point set in R^d can have $\Omega(n^{\lfloor d/2 \rfloor})$ facets. It seems likely that an algorithm deciding the halfspace emptiness problem for such a point set within a polylogarithmic query time should have access to some representation of the convex hull, and thus use a space roughly proportional to its complexity. This is only a heuristic reason, and it would be interesting to prove a lower bound of such kind rigorously, perhaps extending the methods of [Cha89], [BC92], [CR92]. We even venture to conjecture that with a space m, the query time for the halfspace emptiness problem has to be $\Omega(n/m^{1/\lfloor d/2 \rfloor})$, and thus in particular that Theorem 1.1 is close to optimal. Theorem 1.3 and the above mentioned results of Schwarzkopf [Sch92] match this conjectured lower bound even more closely.

The methods used to obtain Theorem 1.1 are similar to the methods of [Mat91b], only a different version of a "cutting lemma" and of a "partition theorem" are needed; these will be developed in sections 2 and 3. Section 4 describes the halfspace range reporting algorithm. In section 5 we bring one application of our auxiliary results (concerning ε-nets with respect to halfspaces in R^3), which turns out to be very similar to an unpublished result of Clarkson [Cla89].

It turns out that the results and methods of this paper have numerous applications in other computational geometry problems; let us mention some of them. In combination with

Megiddo's parametric search technique, one obtains efficient ray shooting algorithms in various settings, algorithms for the post-office problem and k nearest neighbor queries (see [AM92]), an algorithm for linear optimization queries (where a set of constraints for a linear program in a fixed dimension is given in advance, and the queries specify the objective function) and an improved algorithm for computing extreme points of a point set ([Mat91c]).

If we sacrifice something in the query answering efficiency of our algorithm, we obtain a data structure permitting efficient insertions and deletions of points. It turns out that the above mentioned Clarkson's data structure can be dynamized as well. These data structures thus provide an efficient dynamic representation of the convex hull in a certain sense, and as such they are significant also in dimension 3. For example, one can use them for a fast computation of so-called convex layers of a point set in R^3, an output-sensitive computation of levels in hyperplane arrangements, higher order Voronoi diagrams, convex hulls in higher dimensions etc. These issues are discussed in [AM91].

2 Shallow cutting lemma

Let us begin by some definitions. We will assume that H is a collection of n hyperplanes in R^d. For a simplex Δ, let H_Δ denote the collection of hyperplanes of H intersecting its interior. A $(1/r)$-*cutting* for H is a collection Ξ of (possibly unbounded) closed d-dimensional simplices, which cover R^d and 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 concept of a $(1/r)$-cutting originated in applications of random sampling to computational geometry problems (see [HW87], [Cla88], [CF90]), and explicitly it was introduced in [Mat91a]. It turns out to be an useful tool in many computational geometry problems.

Chazelle and Friedman [CF90] proved 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). We will call this result *Cutting lemma.*

For the halfspace range reporting results in this paper, we will need a modification of Cutting lemma, where the simplices are not required to cover the entire R^d, but only a single cell in the arrangement of H. We will prove a slightly more general result. Let o be a fixed point in R^d (*origin* or *reference point*); we define the *level* of a point $x \in \mathrm{R}^d$ (relative to H) as the number of hyperplanes of H crossed by the open segment ox.

Clarkson [Cla88] proved that the number of vertices of the arrangement of H of level at most ℓ is $O(n^{\lfloor d/2 \rfloor}\ell^{\lceil d/2 \rceil})$; there exist arrangements for which this bound is tight.

Let $\ell \leq n$. We say that a collection Ξ of simplices is a $(1/r)$-*cutting for the* $(\leq \ell)$-*level of* H, provided that the simplices of Ξ cover all points of level at most ℓ and $|H_\Delta| \leq n/r$ for every $\Delta \in \Xi$. Let us remark that the simplices of Ξ typically contain also points of level greater than ℓ; obviously, if one wants to cover a polytope with many faces by few simplices, also something more has to be covered.

Theorem 2.1 (Shallow cutting lemma) *Let H be a collection of n hyperplanes in R^d and $\ell, r \leq n$ parameters, and set $q = \ell(r/n) + 1$. There exists a $(1/r)$-cutting Ξ for the $(\leq \ell)$-level of H, consisting of $O(r^{\lfloor d/2 \rfloor}q^{\lceil d/2 \rceil})$ simplices. Moreover, the complement of the union of Ξ is covered by a union of at most $4r$ halfspaces, which are defined by hyperplanes of H and do not contain the origin.*

Such a $(1/r)$-cutting can be computed in polynomial time. If $r \leq n^\alpha$, where $\alpha > 0$ is a certain constant dependent on the dimension, the computation can be performed in $O(n \log r)$ time.

In particular, a single cell of the arrangement of H can be covered by $O(r^{\lfloor d/2 \rfloor})$ simplices, such that the interior of each simplex is intersected by at most n/r hyperplanes; this is the most important special case of the above lemma.

We will need a variant of this result for weighted collections of hyperplanes. A weighted collection of hyperplanes is a pair (H, w), where H is a collection of hyperplanes, and $w : H \to \mathrm{R}^+$ is a nonnegative weight function on H. If $X \subseteq H$, we write just $w(X)$ 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, we say that Ξ is a $(1/r)$-cutting for the (≤ 0)-level of a weighted collection (H, w), if the simplices of Ξ cover all points of level 0 with respect to H and for every simplex Δ of Ξ, the collection H_Δ has total weight $w(H_\Delta) \leq w(H)/r$. The passage from unweighted cuttings to cuttings for weighted collections of hyperplanes is discussed in [Mat91a] (one first normalizes the weights so that $w(H) = n$, replaces every $h \in H$ by $\lfloor w(h) \rfloor + 1$ its copies, and computes a $(1/2r)$-cutting for the resulting multiset of hyperplanes). This method can be applied here for a $(1/r)$-cutting for the (≤ 0)-level of H without any change. We obtain

Corollary 2.2 *Let (H, w) be a weighted collection of n hyperplanes in R^d, $r \leq n$ a parameter. There exists a a $(1/r)$-cutting Ξ of size $O(r^{\lfloor d/2 \rfloor})$ for the (≤ 0)-level of (H, w). Moreover, the complement of the union of Ξ is a union of at most $4r$ halfspaces defined by hyperplanes of H. Such a $(1/r)$-cutting can be computed in polynomial time.* □

The proof of Shallow cutting lemma is similar to the proof of Cutting lemma due to Chazelle and Friedman [CF90]. Alternatively, we could also use a slightly different approach of [Cla88].

In the proof one uses a special kind of triangulation of cells of an arrangement of hyperplanes, called *canonical triangulation* (some authors use the term *bottom-vertex triangulation*). The definition and some properties of the canonical triangulation can be found in [CF90]; here we will recall only the properties directly needed for the proof. For a subcollection $R \subseteq H$ of hyperplanes, let $CT(R)$ denote the set of simplices in the canonical triangulation of the arrangement of R and let $CT_{\leq \ell}(R)$ denote the set of those simplices of $CT(R)$ all whose points have level at most ℓ (relative to H!). We need the following lemma about the canonical triangulation:

Lemma 2.3 [CF90] *Let H be a collection of hyperplanes in R^d.*

(i) *For every simplex Δ of $CT(H)$, there exists a unique inclusion-minimal collection $S(\Delta) \subseteq H$, such that $\Delta \in CT(S(\Delta))$. This collection $S(\Delta)$ has at most a constant number $D = d(d+3)/2$ hyperplanes.*

(ii) *If S is a subcollection of H and Δ is a simplex of $CT(S)$, then Δ belongs to the canonical triangulation of H if and only if its interior is intersected by no hyperplane of H.*
□

We will now formulate the main lemma needed for the proof of Shallow cutting lemma.

Let H and the parameter r be fixed. For a simplex Δ, let the *excess* of Δ be the number $\max(1, |H_\Delta|(r/n))$. Let R be a random sample of hyperplanes of H, where each hyperplane of H is drawn independently with probability $p = r/n$. Let $n(p, t, \ell)$ denote the expected number of simplices with excess at least t in $CT_{\leq \ell}(R)$.

Lemma 2.4 *For $t \geq 1$,*

$$n(p, t, \ell) = O(2^{-t} n(p/t, 0, \ell)) .$$

Proof: This lemma essentially follows from general results of [CF90] (with minor modifications, since we use a different probability space). For reader's convenience, we will sketch the proof (specialized to this particular case) here.

Let $\mathcal{T}$ denote the set of all simplices which belong to $CT_{\leq \ell}(R)$ for some $R \subseteq H$, and let $\mathcal{T}_t$ denote the simplices of $\mathcal{T}$ with excess at least t. Part (ii) of Lemma 2.3 gives us a criterion when a simplex $\Delta \in \mathcal{T}$ appears in $CT_{\leq \ell}(R)$: this is iff

(a) $S(\Delta) \subseteq R$ and

(b) $H_\Delta \cap R = \emptyset$.

The expectation $n(p, t, \ell)$ can be expressed as

$$\sum_{\Delta \in \mathcal{T}_t} p(\Delta),$$

where $p(\Delta)$ denotes the probability that a simplex Δ belongs to $CT_{\leq \ell}(R)$. Now using Lemma 2.3, $p(\Delta)$ is the probability that each of the hyperplanes of $S(\Delta)$ falls into R and none of the hyperplanes of H_Δ does, thus

$$n(p, t, \ell) = \sum_{\Delta \in \mathcal{T}_t} p^{|S(\Delta)|}(1 - p)^{|H_\Delta|}. \tag{1}$$

In order to estimate this sum from above, we consider another random sample $R' \subseteq H$, where each hyperplane $h \in H$ is chosen with probability p/t. We have

$$n(p/t, 0, \ell) = \sum_{\Delta \in \mathcal{T}} (p/t)^{|S(\Delta)|}(1 - p/t)^{|H_\Delta|} \geq$$

$$\sum_{\Delta \in \mathcal{T}_t} t^{-D} p^{|S(\Delta)|}(1 - p)^{|H_\Delta|} \cdot \left(\frac{1 - p/t}{1 - p}\right)^{|H_\Delta|} = \sum_{\Delta \in \mathcal{T}_t} p^{|S(\Delta)|}(1 - p)^{|H_\Delta|} f(\Delta) ,$$

where

$$f(\Delta) = t^{-D} \left(\frac{1 - p/t}{1 - p}\right)^{|H_\Delta|} .$$

Now we may assume that $p \leq 1/2$. Using the inequalities $1 - x \leq e^{-x}$ (valid for all x) and $1 - x \geq e^{-2x}$ (valid for $x \leq 1/2$), we estimate the above factor by

$$f(\Delta) \geq t^{-D} \exp(|H_\Delta|(-2p/t + p)).$$

Since we consider only simplices $\Delta \in \mathcal{T}_t$, i.e. with $|H_\Delta| \geq tn/r = t/p$, we get

$$f(\Delta) \geq t^{-D} e^{t-2},$$

and thus we have

$$n(p,t,\ell) \le t^D e^{-(t-2)} n(p/t,0,\ell) = O(2^{-t} n(p/t,0,\ell)).$$

This proves Lemma 2.4. □

Proof of Theorem 2.1: We begin by proving the existence of the desired $(1/r)$-cutting, and then we will discuss the algorithmic issues. We observe that the number of simplices of $CT_{\le\ell}(R)$ is at most proportional to the number of vertices of the arrangement of R which have level $\le \ell$ relative to H. For each vertex of the arrangement of H of level $\le \ell$, the probability that it becomes a vertex of the arrangement of R is p^d, and applying the Clarkson's result on the number of such vertices, we get that the expected number $n(p,0,\ell)$ of simplices in $CT_{\le\ell}(R)$ is bounded as follows:

$$n(p,0,\ell) = O(p^d n^{\lfloor d/2\rfloor} \ell^{\lceil d/2\rceil}). \tag{2}$$

The desired $(1/r)$-cutting for $(\le \ell)$-level of H is constructed by a refinement of an appropriate portion of $CT(R)$. We proceed as follows: Let Ξ_0 be the set of all simplices of $CT(R)$ which contain some point of level $\le \ell$ (relative to H). For every such simplex Δ, let us consider the collection H_Δ, and let Ξ_Δ be a $(1/t)$-cutting for H_Δ of size $O(t^d)$, where $t = t(\Delta)$ is the excess of Δ (such a cutting exists by Cutting lemma, and actually we would suffice with size $O(t^c)$ for any constant c, which is much easier to obtain than the tight bound). We take the intersection of every simplex $\Delta' \in \Xi_\Delta$ with Δ and we triangulate it. The collection of simplices appearing in these triangulations for all $\Delta \in \Xi_0$ will form our cutting Ξ. It is easy to see that is is a $(1/r)$-cutting for $(\le \ell)$-level of H. We want to bound its size, which amounts to bounding the expected value S of the sum $\sum_{\Delta\in\Xi_0} t(\Delta)^d$. Since a simplex $\Delta \in \Xi_0$ with excess t has all points of level at most $\ell + tn/r$, we can estimate S by

$$\sum_{t=1}^{\infty} t^d n(p,t,\ell + tn/r),$$

and using Lemma 2.4 and the bound (2) for $n(p/t,0,\ell)$, we get the estimates

$$S \le \sum_{t=1}^{\infty} t^d 2^{-t} O(n(p/t,0,\ell + \frac{tn}{r})) \le \sum_{t=1}^{\infty} t^d 2^{-t} O((p/t)^d n^{\lfloor d/2\rfloor} (\ell + \frac{tn}{r})^{\lceil d/2\rceil}) \le$$

$$\le \sum_{t=1}^{\infty} t^d 2^{-t} O(r^{\lfloor d/2\rfloor} t^{-\lfloor d/2\rfloor} (\frac{\ell r}{nt} + 1)^{\lceil d/2\rceil}) = O(r^{\lfloor d/2\rfloor} (\frac{\ell r}{n} + 1)^{\lceil d/2\rceil}),$$

which is just the claimed bound.

It remains to show the additional claim about covering of the complement of Ξ in Theorem 2.1. Obviously the complement is covered by halfspaces defined by the sample R in our construction. We know that the expected number of such halfspaces is r, hence with probability at least 3/4, the cardinality of the sample is not greater than 4r. Similarly with probability at least 3/4, the size of the cutting does not exceed the expectation more than four times, and so with probability at least 1/2, the sample has both these properties.

This proves the existence of the desired $(1/r)$-cutting, and it also gives a polynomial-time randomized algorithm for computing it. The whole construction can be made deterministic using the method of conditional probabilities due to Raghavan and Spencer. We omit the (routine) details, since applications of the method of conditional probabilities in very similar situations are discussed in detail e.g., in [CF90] or [Mat91a]. This yields a deterministic polynomial time algorithm.

Finally the $O(n \log r)$ time bound for small values of r is based on so-called ε-approximations; the method is discussed in [Mat91a] and [Mat91b] and we only sketch it here.

In our setting, we say that a subset $A \subseteq H$ is an *ε-approximation* for H (with respect to simplices; $0 < \varepsilon < 1$ is a real number), if for every simplex Δ it is

$$\left| \frac{|H_\Delta|}{|H|} - \frac{|A_\Delta|}{|A|} \right| < \varepsilon.$$

The desired $(1/r)$-cutting in Theorem 2.1 is computed in two steps. First, we compute a $(1/2r)$-approximation $\tilde{H}$ of cardinality $O(r^d)$ for H; for $r \leq n^\alpha$ with a small enough constant $\alpha > 0$ one can do this in $O(n \log r)$ time, by [Mat91b]. Second, set $\tilde{\ell} = \frac{|\tilde{H}|}{|H|}\ell + \frac{|\tilde{H}|}{r}$. Then we have $\tilde{q} = \tilde{\ell}(|\tilde{H}|/r)+1 \leq \ell(n/r)+2 = q+1$. Use the above discussed polynomial time algorithm to compute a $(1/2r)$-cutting Ξ of size $O(r^{\lfloor d/2 \rfloor}\tilde{q}^{\lceil d/2 \rceil}) = O(r^{\lfloor d/2 \rfloor}q^{\lceil d/2 \rceil})$ for the $(\leq \tilde{\ell})$-level of $\tilde{H}$, in time $r^{O(1)} = o(n)$. By the definition of an ε-approximation, no simplex of Ξ is intersected by more than n/r hyperplanes of H, and also every point of level at most ℓ with respect to H has level at most $\tilde{\ell}$ with respect to $\tilde{H}$. Hence Ξ is a $(1/r)$-cutting for the $(\leq \ell)$-level of H. $\Box$

3 Partition theorem for shallow hyperplanes

Let us repeat a definition from [Mat91b]. Let P be an n-point set in R^d. A *simplicial partition* for P is a collection

$$\Pi = \{(P_1, \Delta_1), \ldots, (P_m, \Delta_m)\},$$

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 the set P_i.

Let us remark at this point that the simplex Δ_i may contain also other points of P than those of P_i. Also, we admit not only full-dimensional simplices, but also simplices of any dimension $k \leq d$. This is sometimes necessary for point sets P in degenerate positions; for a general position, one may take only full-dimensional simplices in the simplicial partitions in the forthcoming Partition Theorem for shallow hyperplanes.

The simplicial partitions we will work with will satisfy the following condition:

$$\max\{|P_i|;\ (P_i, \Delta_i) \in \Pi\} < 2\min\{|P_i|;\ (P_i, \Delta_i) \in \Pi\},$$

so all the classes will have roughly the same size.

If h is a hyperplane and Δ a simplex, we say that h *crosses* Δ if $h \cap \Delta \neq \emptyset$ and $\Delta \not\subseteq h$ (thus a hyperplane does not cross a lower-dimensional simplex contained in it). Let h be a hyperplane; we define the *crossing number* of h (relative to Π) as the number of simplices among the Δ_i's crossed by h.

The basis of the efficient algorithms in [Mat91b] was the so-called Partition theorem, which states that for every a n-point set $P \subseteq \mathrm{R}^d$ and a parameter r, $1 < r < n$ there exists a simplicial partition Π for P, whose classes P_i satisfy $\lceil n/r \rceil \leq |P_i| < 2\lceil n/r \rceil$, and such that the crossing number of any hyperplane relative to Π is $O(r^{1-1/d})$.

In this section we modify the proof of this theorem to obtain a result which allows us to treat the halfspace range reporting. We say that a hyperplane h is *k-shallow* (relative to P), if one of the open halfspaces determined by h contains no more than k points of P.

Theorem 3.1 (Partition theorem for shallow hyperplanes) *Let P be an n-point set in R^d ($d \geq 2$) and let r be a parameter, $1 < r < n$. There exists a simplicial partition Π for P, whose classes P_i satisfy $\lceil n/r \rceil \leq |P_i| < 2\lceil n/r \rceil$, and such that the crossing number of any (n/r)-shallow hyperplane relative to Π is $O(r^{1-1/\lfloor d/2 \rfloor})$ (for $d \geq 4$), resp. $O(\log r)$ (for $d = 2, 3$).*

A simplicial partition as above can be computed in time $O(n^{1+\delta})$. For $r \leq n^\alpha$ (with a suitable constant $\alpha > 0$), it can be computed in $O(n \log r)$ time.

For $d \geq 4$, the crossing number is asymptotically optimal. For $d = 2, 3$, it would be interesting to see to what extent the crossing number can be improved. It is not clear whether a crossing number bounded by a constant can be attained (we conjecture that it cannot).

The whole proof is quite similar to the one in [Mat91b], which in turn uses ideas of Welzl (see [CW89]). We again begin with an existence proof, and then we discuss algorithmic issues. We will assume that r is sufficiently large (larger than a suitable constant), since otherwise the theorem is trivial.

The proof is based on the following lemma:

Lemma 3.2 *Let P, n, r be as above, and let Q be a set of $(2n/r)$-shallow hyperplanes. Then there exists a simplicial partition Π for P, whose classes P_i satisfy $\lceil n/r \rceil \leq |P_i| < 2\lceil n/r \rceil$ for every i, and such that the crossing number of every hyperplane of Q relative to Π is $O(r^{1-1/\lfloor d/2 \rfloor} + \log |Q|)$ (for $d \geq 4$), resp. $O(\log r + \log |Q|)$ (for $d = 2, 3$).*

Proof: Let o be a point (not belonging to P) such that any halfspace not containing o contains fewer than $n(1 - 2/r)$ points of P. Assuming that $r > 2(d+1)$, then the centerpoint of P is such a point (this is a point such that any halfspace not containing it contains at most $n(1 - \frac{1}{d+1})$ points of P, which always exists by Helly's theorem, see [Ede87]).

Let us set $s = \lceil n/r \rceil$. We will inductively construct the disjoint sets $P_1, P_2, \ldots \subseteq P$, and simplices $\Delta_1, \Delta_2, \ldots$, $P_i \subseteq \Delta_i$. Suppose that $P_1, \ldots, P_i$ have already been constructed, and set $P_i' = P \setminus (P_1 \cup \ldots \cup P_i)$, $n_i = |P_i'|$.

Let K be a large enough constant. We will distinguish two cases. First suppose that $n_i \geq Ks$. For a hyperplane $h \in Q$, let $\kappa_i(h)$ denote the number of simplices among $\Delta_1, \ldots, \Delta_i$ crossed by h. We define a weighted collection (Q, w_i) by setting $w_i(h) = 2^{\kappa_i(h)}$ for every $h \in Q$.

Let us choose a parameter t_i such that there exists a $(1/t_i)$-cutting Ξ_i for the (≤ 0)-level of (Q, w_i), whose simplices have at most $n_i/2s$ faces of all dimensions in total (the level is defined with respect to the above specified point o). By Corollary 2.2, we may set $t_i = c(n_i/s)^{1/\lfloor d/2 \rfloor}$, for some positive constant c, $0 < c < 1/16$. Since $n_i > Ks$, we may suppose that $t_i > 1$.

Further, we may assume that the complement of the union of Ξ is covered by the union of at most $4t_i$ halfspaces defined by hyperplanes of Q. By the choice of the origin o and since each hyperplane of Q is $2s$-shallow, each of these halfspaces contains at most $2s$ points of P.

All these $4t_i$ halfspaces together thus contain at most $8t_i s \leq 8c(n_i/s)^{1/\lfloor d/2 \rfloor} s \leq 8cn_i < n_i/2$ points of P. Hence the simplices of Ξ_i together contain at least $n_i/2$ points of P'_i.

Since the simplices of Ξ_i have at most $n_i/2s$ faces, there exists a relatively open face of some simplex of Ξ_i containing at least s points of P'_i. Let Δ_{i+1} be some such relatively open face. Since $s \geq 2$, the dimension of Δ_i is at least 1. Among the at least s points of P'_i contained in Δ_{i+1}, let us choose an s-point subset (arbitrarily) and call it P_{i+1}.

Proceeding with this construction, we reach an index $i = q$ such that $|n_q| < Ks$. We then partition the set P'_q of remaining points into sets P_{q+1}, $P_{q+2}, \ldots,$ P_m of size between s and $2s$ arbitrarily, and we define the simplices $\Delta_{q+1} = \Delta_{q+2} = \cdots = \Delta_m = \mathrm{R}^d$. We then set $\Pi = \{(P_1, \Delta_1), \ldots, (P_m, \Delta_m)\}$, which finishes the construction. Note that in this last phase, we have added at most a constant number of classes and simplices, which only increases the crossing number of any hyperplane by an additive constant. Hence it suffices to consider the crossing numbers relative to $\Pi' = \{(P_1, \Delta_1), \ldots, (P_q, \Delta_q)\}$

Let us establish the bound on the crossing numbers of the hyperplanes of Q relative to Π' (following the proof in [Mat91b] literally). The bound is obtained by estimating the final total weight $w_q(Q)$ of the hyperplanes of Q in two different ways.

The weight $w_q(h)$ of a hyperplane $h \in Q$ with crossing number κ is equal to 2^κ. Therefore $\kappa \leq \log_2 w_q(Q)$.

Let us consider how $w_{i+1}(Q)$ increases compared to $w_i(Q)$. Let Q_{i+1} denote the collection of the hyperplanes of Q crossing Δ_{i+1}. For the hyperplanes of Q_{i+1}, the weight increases twice, and for the others it remains unchanged. From this we get

$$w_{i+1}(Q) \leq w_i(Q) - w_i(Q_{i+1}) + 2w_i(Q_{i+1}) = w_i(Q)\left(1 + \frac{w_i(Q_{i+1})}{w_i(Q)}\right).$$

We claim that $w_i(Q_{i+1}) \leq w_i(Q)/t_i$. Indeed, this is clear from the definition of a $(1/t_i)$-cutting if Δ_i has dimension d. For a simplex of lower dimension we note that if a hyperplane crosses a face of a full-dimensional simplex then it also intersects the interior of that simplex; this is easily seen by induction on dimension. Hence

$$w_{i+1}(Q) \leq w_i(Q)\left(1 + \frac{1}{t_i}\right) \leq w_i(Q)\left(1 + \frac{s^{1/\lfloor d/2 \rfloor}}{cn_i^{1/\lfloor d/2 \rfloor}}\right)$$

and using $w_0(Q) = |Q|$, $n_i = n - is$, $r \leq n/s$, $q < r$, we get

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

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

$$\log w_q(Q) \leq \log |Q| + \frac{1}{c} \sum_{i=0}^{q-1} \frac{1}{(r-i)^{1/\lfloor d/2 \rfloor}} \leq \log |Q| + \frac{1}{c} \sum_{j=1}^{r} \frac{1}{j^{1/\lfloor d/2 \rfloor}}.$$

Bounding the last sum by integral, we finally obtain

$$\kappa \leq \log_2 w_m(Q) = O(\log |Q| + r^{1-1/\lfloor d/2 \rfloor})$$

(for $d \geq 4$), resp. $\kappa = O(\log |Q| + \log r)$ (for $d = 2, 3$). This concludes the proof of Lemma 3.2. $\square$

The next step in the proof of Theorep 3.1 is to choose a "test set" of hyperplanes Q, guaranteeing that the crossing number of *no* (n/r)-shallow hyperplane will exceed the required bound.

Lemma 3.3 (Shallow test set lemma) *Let P, n, r be as above. Then there exists a set Q of $O(r^{\lfloor d/2 \rfloor})$ $(2n/r)$-shallow hyperplanes, such that for any simplicial partition $\Pi = \{(P_1, \Delta_1), \ldots, (P_m, \Delta_m)\}$ satisfying $|P_i| \geq \lceil n/r \rceil$ for every i the following holds: If κ_0 is the maximum of crossing numbers of hyperplanes of Q relative to Π, then the crossing number of any (n/r)-shallow hyperplane relative to Π is bounded by $(d+1)\kappa_0 + 1$.*

Proof: Let $H = \mathcal{D}(P)$ be the collection of hyperplanes dual to the points of P (we use the duality with origin o, where o is as in the proof of Lemma 3.2). Let us choose a $(1/r)$-cutting Ξ of size $O(r^{\lfloor d/2 \rfloor})$ for the $(\leq n/r)$-level of H. Let Q_0 be the set of all hyperplanes dual to the vertices of the simplices of Ξ, and let Q be those hyperplanes of Q_0 which are $(2n/r)$-shallow. Clearly $|Q| = O(r^{\lfloor d/2 \rfloor})$.

Let h be any (n/r)-shallow hyperplane. The point $\mathcal{D}(h)$ dual to h has level at most n/r, hence it is contained in some simplex $\sigma \in \Xi$. Let G be the set of hyperplanes dual to the vertices of this σ. Since σ contains the point $\mathcal{D}(h)$ of level at most n/r and it is intersected by at most n/r hyperplanes of H, its vertices have level at most $2n/r$, and so $G \subseteq Q$. By the assumption of the lemma, each hyperplane of G crosses at most κ_0 simplices of the simplicial partition Π.

Consider a simplex Δ_i of Π which crosses the hyperplane h but none of the hyperplanes of G. This means that Δ_i is completely contained in the zone of h in the arrangement of G. It is elementary to verify that any point of that zone dualizes to a hyperplane intersecting the simplex σ. Hence, the zone can only contain at most n/r points of P. Since each Δ_i as above contains at least n/r points of $P_i \subseteq P$, there can be at most one such Δ_i. Together with the at most $(d+1)\kappa_0$ simplices of Π crossing some of the hyperplanes of G, we get at most $(d+1)\kappa_0 + 1$ simplices crossing h. $\square$

The proof of Theorem 3.1 now follows easily. Given the set P and the parameter r, we first construct a set Q of $O(r^{\lfloor d/2 \rfloor})$ $(2n/r)$-shallow hyperplanes as in Lemma 3.3. Then we apply Lemma 3.2, obtaining a simplicial partition Π such that all hyperplanes of Q have crossing number $O(r^{1-1/\lfloor d/2 \rfloor} + \log |Q|) = O(r^{1-1/\lfloor d/2 \rfloor})$ relative to Π (resp. crossing number $O(\log r)$ for $d = 2, 3$). By the "test set" property of Q, all (n/r)-shallow hyperplanes have crossing number $O(r^{1-1/\lfloor d/2 \rfloor})$ (resp. $O(\log r)$) relative to Π.

It remains to discuss algorithmic issues. Let us look at the situation when r is quite small, $r < n^\alpha$.

First we need to find the origin o as in the proof of Lemma 3.2. It is probably difficult to find a real centerpoint, but it is enough to find an approximate centerpoint, i.e. a point such that any halfspace not containing it contains at most $n(1 - \frac{1}{2(d+1)})$ (say) points of P. This is possible to do in linear time using ε-approximations, see [Mat91d].

In the construction of the simplicial partition in Lemma 3.2, all steps can be performed in time polynomial in r, the only exception being the counting of the number of points in the faces of simplices of the Ξ_i's, in order to find one with enough points in it (and then extracting these points). It turns out that this can be handled in total time $O(n \log r)$ using a suitable data structure for simplex range counting queries on the set P. The same problem is addressed in [Mat91b] and we omit the discussion here. As for Test set lemma, one computes a shallow cutting in the dual space, which can be done in $O(n \log r)$ time by Theorem 3.1. Then the hyperplanes of Q_0 which are not $(2n/r)$-shallow must be discarded. This can again be done in $O(n \log r)$ time using a suitable halfspace range counting data structure for P.

Finally for a greater value of r (closer to n), we use the technique of refining the simplicial partition. We start by computing a simplicial partition Π_0, with a small value of the parameter r, then we compute a simplicial partition for every class of Π_0, etc. This is again discussed in [Mat91b]. $\square$

4 Reporting points in halfspaces

We are now ready to describe the data structure for halfspace range reporting whose existence was asserted in Theorem 1.1. It will be a so-called partition tree. Each node v of the tree will correspond to some subset of $P_v \subseteq P$, the root will correspond to the whole P. The sets corresponding to children of a node will form a partition of the set of their father, and the sets for leaves will have size not exceeding a suitably chosen constant.

Let $\beta > 0$ be a suitable constant smaller than $1/d$ (we will put more requirements on it when discussing the preprocessing algorithm). We begin the construction in the root of the partition tree and we proceed recursively down the tree. For a node v, we set $n_v = |P_v|$, $r_v = n_v^\beta$, and we choose (according to Theorem 3.1) a simplicial partition Π_v of size $O(r_v)$ for P_v with classes no larger than $2n_v/r_v$, such that the crossing number of any (n_v/r_v)-shallow hyperplane will be at most $\kappa_v = O(r_v^{1-1/\lfloor d/2 \rfloor})$. The sets in Π_v will correspond to $O(r_v)$ children of v. We moreover associate with the node v a linear-size data structure which allows to report points of P_v in a query halfspace in time $O(n_v^{1-1/d}(\log n_v)^b + k)$, where k is the number of points reported; we use the simplex range searching result of [Mat91b], mentioned in the introduction.

Since each level of the tree uses $O(n)$ space and the depth of the tree is $O(\log \log n)$, the total space is $O(n \log \log n)$. The query answering algorithm is very simple. We start in the root, and being in a node v, we perform the following recursive procedure: We detect all simplices of Π_v crossed by the boundary hyperplane of the query halfspace γ. If the number of such simplices is greater than κ_v (in that case we know that the hyperplane is not (n_v/r_v)-shallow relative to P_v), we use the auxiliary data structure of that node to report the points of $P_v \cap \gamma$. Otherwise we report all points in simplices of Π_v completely contained in γ, and we proceed recursively down the tree for the simplices crossing the boundary hyperplane, eventually stopping at the leaves, where we answer the query trivially by inspecting all points.

First we estimate the work done when using the auxiliary data structures. Should such a data structure be used in a node v, it must be $k_v = |P_v \cap \gamma| \geq n_v/r_v = n_v^{1-\beta}$ with $\beta < 1/d$, and the auxiliary data structure uses time $O(n_v^{1-1/d}(\log n_v)^b + k_v) = O(k_v)$, thus the total work is proportional to $k = |P \cap \gamma|$. Let $T(m)$ denote the maximum running time of the query answering algorithm for an m-point set, not counting the time needed by the auxiliary data structures. Then $T(m)$ satisfies the recurrence

$$T(m) \leq T_0 \text{ for } m \leq m_0,$$

$$T(m) \leq Ar + Br^{1-1/\lfloor d/2 \rfloor} T(2m/r)$$

with $r = m^\beta$ and A, B, T_0, m_0 constants, whose solution is

$$T(n) = O(n^{1-1/\lfloor d/2 \rfloor} (\log n)^c)\,,$$

for a suitable constant c.

The preprocessing requires to compute the appropriate simplicial partitions and to build the auxiliary data structures. If β is chosen small enough, the simplicial partition in an m-point node is computed in $O(m \log m)$ time according to Theorem 3.1. The auxiliary data structure also requires $O(m \log m)$ time. This implies that the total preprocessing time is $O(n \log n)$ (since the *logarithm* of node size decreases geometrically with the depth in the partition tree). This finishes the proof of Theorem 1.1. □

Corollary 1.2 – proof sketch. The partition tree data structure constructed in the proof of Theorem 1.1 can be combined with a data structure with $O(\log n + k)$ query time (e.g., the Clarkson's one mentioned in the introduction) to obtain a continuous space-query time tradeoff. Namely, we build only an upper part of the partition tree (as in the linear-space algorithm), stopping the recursive construction as soon as the size of the point set in a node drops below a suitably chosen number s. For such a node v, instead of building the subtree for P_v, we store the data structure with logarithmic query time for P_v at the node v. A similar construction appears in several previous papers (e.g., [CSW90]), so we omit the details. □

Proof of Theorem 1.3. The key to the improvement for the emptiness problem is making the partition tree much more shallow. Namely, in a node v, we choose $r_v = n_v/(\log n_v)^C$, where C is a suitable constant. We also will not need the auxiliary data structures for the simplex range reporting in P_v, but we will use another auxiliary data structure instead. We consider the set V_v of all vertices of simplices of the simplicial partition Π_v associated with the node v, and we build the halfspace range reporting structure described in Theorem 1.1 for it. This auxiliary structure has size $O(r_v \log\log r_v) = O(n_v/\log n_v)$, and this implies that the whole data structure uses only $O(n)$ space.

The query answering for the emptiness problem for a query halfspace γ again starts in the root, and in a node v the algorithm is as follows: We set $t_v = f(|V_v|, (d+1)\kappa_v)$, where $f(n,k) = O(n^{1-1/\lfloor d/2 \rfloor}(\log n)^c + k)$ is an upper bound on the complexity of the query answering algorithm from Theorem 1.1. We run the query answering algorithm for the reporting of points of V_v in γ, terminating it after t_v time units if it has not finished yet. Now if it has run at least t_v time units, we know that $|V_v \cap \gamma| > (d+1)\kappa_v$, and thus γ cannot be empty (either too many simplices cross the bounding hyperplane, or some simplex is completely contained in γ). On the other hand, if all points of V_v in the query halfspace were reported, we can list all simplices intersecting the query halfspace γ. We thus check if some of them is completely contained in γ (then we may stop), and if not we perform a recursion for the simplices crossing the boundary of γ.

The value of r is chosen in such a way that the work spent in a node is $o(n_v^{1-1/\lfloor d/2 \rfloor})$, and so the running time will satisfy

$$T(m) \leq o(m^{1-1/\lfloor d/2 \rfloor}) + Br^{1-1/\lfloor d/2 \rfloor} T(2m/r),$$

with $r = m/(\log m)^C$, which gives $T(n) = O(n^{1-1/\lfloor d/2 \rfloor} 2^{c' \log^* n})$.

The above described data structure needs simplicial partitions with r quite close to the number of points ($r = n/(\log n)^C$), which we are only able to compute in $O(n^{1+\delta})$ time. □

5 Linear-size ε-nets in dimension 3

One interesting open problem in combinatorial geometry concerns ε-nets with respect to halfspaces. Let $X \subseteq \mathrm{R}^d$ be a finite point set. Let ε be a real number, $0 \leq \varepsilon < 1$. A subset $N \subseteq X$ is called an *ε-net for X with respect to halfspaces,* if every halfspace containing more than $\varepsilon|X|$ points of X also contains a point of N. The notion of an ε-net can actually be defined in a more general setting, and from general results of Haussler and Welzl [HW87] one gets that for every fixed d and for every ε and X, there exists an ε-net for X with respect to halfspaces, of size $O((1/\varepsilon)\log(1/\varepsilon))$. The problem is whether this bound can be sharpened, e.g. if ε-nets of size $O(1/\varepsilon)$ with respect to halfspaces always exist.

A partial result was given by Seidel et al. [SWM90], who proved (by a quite complicated geometric argument) that for every finite point set $X \subseteq \mathrm{R}^3$ and every $\varepsilon \in (0,1)$, there exists an ε-net for X with respect to halfspaces, of size $O(1/\varepsilon)$.

Later Clarkson gave an alternative probabilistic proof of this result in an unpublished note [Cla89]. By an easy application of the Shallow cutting lemma, we get a proof quite similar to Clarkson's.

It will be more convenient for us to consider a dual version of this problem, where we have a collection H of hyperplanes and we want to select a subset $N \subseteq H$ such that any point of level $> n/r$ is separated from the reference point o by some hyperplane of N.

Let us apply Shallow cutting lemma with $d = 3$, $\ell = 2n/r$ and $2r$ instead of r, choosing a $(1/2r)$-cutting Ξ of size $O(r)$ for the $(\leq 2n/r)$-level of H. Now for every simplex $\Delta \in \Xi$ which is separated from the reference point o by at least one hyperplane $h \in H$ (i.e. Δ and o lie in opposite halfspaces determined by h), let us choose one such separating hyperplane h_Δ. Let all such hyperplanes h_Δ form a set N with $O(r)$ elements. We claim that N has the desired property. Indeed, it suffices to show that for every point of level greater than n/r but smaller than $2n/r$ there is a hyperplane of N separating it from the reference point o. Let us consider one such point x and let Δ be the simplex of Ξ containing it. There are at least n/r hyperplanes separating Δ from o, and at most $n/2r$ of them may cross the simplex Δ; thus the remaining ones of them separate Δ from o. Therefore there is a hyperplane separating Δ (and also x) from o in N. □

Acknowledgement. I would like to thank the referees for thorough reading of the manuscript and valuable comments.

References

[AHT90] A. Aggarwal, M. Hansen, and T.Leighton. Solving query-retrieval problems by compacting Voronoi diagrams. In *Proc. 21. ACM Symposium on Theory of Computing*, pages 331–340, 1990.

[AM91] P. K. Agarwal and J. Matoušek. Dynamic half-space range reporting and its applications. Tech. Report CS-1991-43, Duke University 1991, 1991.

[AM92] P. K. Agarwal and J. Matoušek. Ray shooting and parametric search. In *Proc. 24. ACM Symposium on Theory of Computing*, 1992. To appear. Also published as Tech. Report CS-1991-22, Duke University 1991.

[BC92] H. Brönninmann and B. Chazelle. How hard is halfspace range searching. In *Proc. 8. ACM Symposium on Computational Geometry*, 1992. To appear.

[CF90] B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.

[CGL85] B. Chazelle, L. Guibas, and D. T. Lee. The power of geometric duality. *BIT*, 25(1), 1985.

[Cha86] B. Chazelle. Filtering search: A new approach to query-answering. *SIAM Journal on Computing*, 15(3):703–724, 1986.

[Cha89] B. Chazelle. Lower bounds on the complexity of polytope range searching. *J. Amer. Math. Soc*, 2(4):637–666, 1989.

[Cla88] K. L. Clarkson. Applications of random sampling in computational geometry II. In *Proc. 4. ACM Symposium on Computational Geometry*, pages 1–11, 1988.

[Cla89] K. Clarkson, unpublished proof outline, December 1989.

[CP86] B. Chazelle and F. P. Preparata. Halfspace range searching: An algorithmic application of k-sets. *Discrete & Computational Geometry*, 1:83–93, 1986.

[CR92] B. Chazelle and B. Rosenberg. Lower bounds on the complexity of simplex range reporting on a pointer machine. In *Proc. International Colloquium on Automata, Languages and Programming*, 1992. To appear.

[CSW90] B. Chazelle, M. Sharir, and E. Welzl. Quasi-optimal upper bounds for simplex range searching and new zone theorems. In *Proc. 6. ACM Symposium on Computational Geometry*, pages 23–33, 1990.

[CW89] B. Chazelle and E. Welzl. Quasi-optimal range searching in spaces of finite VC-dimension. *Discrete & Computational Geometry*, 4:467–490, 1989.

[Ede87] H. Edelsbrunner. *Algorithms in combinatorial geometry.* Springer-Verlag, 1987.

[HW87] D. Haussler and E. Welzl. ε-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987.

[Mat91a] J. Matoušek. Cutting hyperplane arrangements. *Discrete & Computational Geometry*, 6(5):385–406, 1991.

[Mat91b] J. Matoušek. Efficient partition trees. In *Proc. 7. ACM Symposium on Computational Geometry*, pages 1–9, 1991. Also to appear in *Discrete & Computational Geometry.*

[Mat91c] J. Matoušek. Linear optimization queries. *Journal of Algorithms*, 1991. To appear. The results combined with results of O. Schwarzkopf also appear in *Proc. 8. ACM Symposium on Computational Geometry*, 1992.

[Mat91d] 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 ed.), Amer. Math. Soc. 1991, pages 221–230.

[Mat92] J. Matoušek. Range searching with efficient hierarchical cuttings. In *Proc. 8. ACM Symposium on Computational Geometry*, 1992. To appear.

[Sch92] O. Schwarzkopf. Ray shooting in convex polytopes. In *Proc. 8. ACM Symposium on Computational Geometry*, 1992. To appear. Also published as Technical Report B–91–18, FB Mathematik, Freie Universität Berlin 1991.

[SWM90] R. Seidel, E. Welzl, and J. Matoušek. Netting a lot with a little: Small ε-nets for disks and halfspaces. In *Proc. 6. ACM Symposium on Computational Geometry*, pages 16–22, 1990.

[Wil82] D. E. Willard. Polygon retrieval. *SIAM Journal on Computing*, 11:149–165, 1982.

[YY85] F. F. Yao and A. C. Yao. A general approach to geometric queries. In *Proc. 17. ACM Symposium on Theory of Computing*, pages 163–168, 1985.