

Geometric Range Searching

JIŘÍ MATOUŠEK

*Department of Applied Mathematics, Charles University, Malostranské nám. 25,
118 00 Praha 1, Czech Republic*

In geometric range searching, algorithmic problems of the following type are considered. Given an n -point set P in the plane, build a data structure so that, given a query triangle R , the number of points of P lying in R can be determined quickly. Similar questions can be asked for point sets in higher dimensions, with triangles replaced by simplices or by more complicated shapes. Algorithms of this type are of crucial importance in computational geometry, as they can be used as subroutines in solutions to many seemingly unrelated problems, which are often motivated by practical applications, for instance in computer graphics (ray tracing, hidden-surface removal etc.). We present a survey of theoretical results and the main techniques in geometric range searching.

Categories and Subject Descriptors: E.1 [Data]: Data Structures; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*geometric problems and computations*; G.m [Mathematics of Computing]: Miscellaneous; H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—geometric algorithms, languages, and systems

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Computational Geometry, lower bounds in arithmetic model, partition tree, range searching

1. INTRODUCTION

This survey considers an algorithmic problem called *range searching*. We describe the problem and outline current theoretical knowledge about it, including the main ideas of several proofs and constructions.

Together with the main material, we also include various extensions, historical notes, and so on. Many of them are somewhat subjective. Several notions and results of computational geometry appearing in our discussion but not directly related to range searching are explained

in separate boxes. Table 1 summarizes the current best-known complexity bounds for simplex and halfspace range searching.

Computational Geometry — General Remarks

The problems being considered belong to the area of *computational geometry*. In the rest of this section we briefly introduce this field and mention some features and conventions which seem particular to it; readers familiar with computational geometry may want to skip the

Part of the work on this survey was done while the author was visiting the Computer Science Institute, Free University Berlin. The work has been supported by the German-Israeli Foundation of Scientific Research and Development (G.I.F.).

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1994 ACM 0360-0300/94/1200-0421 \$03.50

CONTENTS

- 1 INTRODUCTION
- 2 GEOMETRIC RANGE SEARCHING
- 3 INTUITION AND LOWER BOUNDS
- 4 ALGORITHMS WITH APPROXIMATELY LINEAR SPACE
- 5 ALGORITHMS WITH LOGARITHMIC QUERY TIME
- 6 EXTENSIONS AND EXAMPLES OF APPLICATIONS
 - 6.1 Halfspace Range Reporting and Other Special Situations
 - 6.2 Dynamization
 - 6.3 Multilevel Data Structures
 - 6.4 Searching with More General Ranges
 - 6.5 Ray Shooting and Linear Optimization
 - 6.6 More on Applications
- 7 A NOTE IN CONCLUSION

rest of Section 1. Older computational geometry monographs of a wide scope are Preparata and Shamos [1985], Edelsbrunner [1987] and Mehlhorn [1984]; a recent one is Mulmuley [1993]. Useful survey papers can be found in Pach [1993] and Goodman et al. [1991].

The subject of computational geometry is the design and analysis of efficient algorithms for computing various properties and parameters of finite configurations of geometric objects. The last sentence may sound rather cryptic, and the reader can get a better picture from several examples of simple problems studied in computational geometry:

- (P1) Given an n -point set P in the plane (or in a Euclidean space of dimension d),¹ find a pair of points of P with smallest distance.
- (P2) Given n segments in the plane, compute the number of their intersections.
- (P3) Given a convex polytope P in d -dimensional space specified as an intersection of n halfspaces, determine a point of P minimizing a given linear function.

¹This is probably the most frequent preamble to a problem in computational geometry.

- (P4) Given a polygon P in the plane, not necessarily convex, and two points a, b in its interior, find a shortest path from a to b inside P , or (another version) a path from a to b consisting of a minimum possible number of segments.

And so on. In computational geometry it is deceptively easy to formulate problems. Many problems have practical or theoretical motivations, and for almost any problem one can make up an acceptable-looking application *ex post*. In the pioneer era, say around the year 1980 and the following few years, it was not too difficult to find an unsolved elementary and interesting problem. Today, more complicated and more powerful methods are known, and finding a really remarkable new problem is becoming quite hard.

The subject of computational geometry is similar to the areas of interest in other fields, such as geometric aspects of combinatorial optimization. For instance, a careful reader might have recognized Problem P3 as a linear programming problem. The approach of computational geometry is usually distinguished by two typical features: the problems are considered in fixed dimension, and an infinite precision model of computation is used.

Dimension and Hidden Constants

The original problems in computational geometry were formulated in the plane or in three-dimensional space. While generalizations to arbitrary dimension are also studied, we almost always assume that the dimension is quite small, say at most 10, while the number of objects (points, segments, etc.) in the problem is large.

Formally the dimension is usually considered as a constant. The efficiency of typical computational geometry algorithms decreases with the dimension quite rapidly; for example, the multiplicative constants in the asymptotic bounds for the complexity of the algorithms are usually exponential in the

Table 1. Results on the Simplex and Halfspace Range-Searching Problems

Dim	Query Type	Space	Query Time	Preprocessing	Update Time (amortized)	Source
Simplex range searching—lower bounds						
2	general ¹	m	n/\sqrt{m}	—	—	[Chazelle 1989]
$d \geq 3$	general ¹	m	$n/(m^{1/d} \log n)$	—	—	[Chazelle 1989]
$d \geq 2$	reporting ²	m	$(n^{1-\varepsilon}/m^{1/d}) + k$	—	—	[Chazelle and Rosenberg 1992]
Simplex range searching—upper bounds						
2	general	n	$\sqrt{n} \log n$	polynom. ⁴	—	[Chazelle and Welzl 1989]
3	general	$n \log n$	$n^{2/3} \log^2 n$	polynom. ⁴	—	[Chazelle and Welzl 1989]
$d \geq 2$	general	n	$n^{1-1/d}$	$n^{1+\varepsilon}$	—	[Matoušek 1993c]
$d \geq 2$	general	n	$n^{1-1/d} \log^c n$	$n \log n$	$\log^2 n$	[Matoušek 1992b]
$d \geq 2$	general	n^d	$\log^{d+1} n$	$n^d (\log n)^\varepsilon$	—	[Matoušek 1993c]
$d \geq 2$	general	$n^{d+\varepsilon}$	$\log^{d+1} n$	$n^{d+\varepsilon}$	$n^{d-1+\varepsilon}$	[Chazelle et al. 1992] ³
$d \geq 2$	general	m	$(n/m^{1/d}) \log^{d+1} \frac{m}{n}$	$n^{1+\varepsilon} + m(\log n)^\varepsilon$	—	[Matoušek 1993c]
$d \geq 2$	general	$m^{1+\varepsilon}$	$(n/m^{1/d}) \log^{d+1} \frac{m}{n}$	$m^{1+\varepsilon}$	$m^{1+\varepsilon}/n$	[Chazelle et al. 1992] ³
Halfspace range searching—lower bound						
$d \geq 2$	general ¹	m	$(n/\log n)^{1-(d-1)/d(d+1)}$	—	—	[Brönnimann et al. 1993]

Continued

Explanations. The $O()$ symbol for upper bounds and the $\Omega()$ symbol for lower bounds are omitted. The letter m denotes a parameter that can be chosen at will; for simplex range searching and general halfspace range searching, its range is $[n, n^d]$, and for halfspace emptiness or halfspace reporting it is $[n, n^{d/2}]$. The letter ε denotes a positive constant that can be made arbitrarily small by adjusting the parameters of the algorithm; c is a generic letter for various specific constants (depending on d). The algorithms that achieve the preprocessing times are deterministic unless stated otherwise.

¹The bounds hold in the arithmetic model.

²The bound holds in the pointer machine model.

³Dynamization is due to Agarwal and Sharir [1993].

⁴The preprocessing can be made $O(n^{1+\varepsilon})$ using known methods (this has not been published).

Table 1. Continued

Dim	Query Type	Space	Query Time	Preprocessing	Update Time (amortized)	Source
Halfspace range searching—upper bounds ⁵						
$d \geq 2$	general	$n^d / \log^d n$	$\log n$	$n^d / \log^{d-\epsilon} n$	—	[Matoušek 1993c]
$d \geq 2$	general	m	$n/m^{1/d}$	$n^{1+\epsilon} + m(\log n)^\epsilon$	—	[Matoušek 1993c]
2	reporting	n	$\log n + k$	$n \log n$	—	[Chazelle et al. 1985]
3	reporting	$n \log n$	$\log n + k$	$n \log n$	—	[Agarwal et al. 1990]
$d \geq 4$	reporting	$n \log \log n$	$n^{1-1/(d/2)} \log^c n + k$	$n \log n$	—	[Matoušek 1992a]
$d \geq 4$	emptiness	n	$n^{1-1/(d/2)} \log^c n$	$n \log n$	—	[Matoušek 1992a]
$d \geq 4$	emptiness	n	$n^{1-1/(d/2)} 2^c \log^* n$	$n^{1+\epsilon}$	—	[Matoušek 1992a]
$d \geq 2$	reporting	n	$n^{1-1/(d/2)+\epsilon} + k$	n	$\log^2 n$	[Agarwal and Matoušek 1991]
$d \geq 4$	emptiness ⁷	$n^{(d/2)} \log^c n$	$\log n$	$n^{(d/2)} \log^c n$	—	[Matoušek and Schwarzkopf 1993]
$d \geq 4$	emptiness	$n^{(d/2)}/\log^{(d/2)-\epsilon} n$	$\log n$	$n^{(d/2)}/\log^{(d/2)-\epsilon} n$	—	[Matoušek and Schwarzkopf 1993]
$d \geq 2$	reporting	$n^{(d/2)+\epsilon}$	$\log n + k$	$n^{(d/2)+\epsilon}$	$n^{(d/2)-1+\epsilon}$	[Clarkson 1988b] ⁹
$d \geq 4$	emptiness ⁷	$m \log^c n$	$n/m^{1/d} \log n$	$m \log^c n$	—	[Matoušek and Schwarzkopf 1993]
$d \geq 2$	reporting	$m^{1+\epsilon}$	$(n/m^{1/d}) \log n + k$	$m^{1+\epsilon}$	$m^{1+\epsilon}/n$	[Agarwal and Matoušek 1991]

⁵All upper bounds for simplex range searching also apply.

⁶Within this time, a randomized preprocessing algorithm produces a data structure having the given parameters with high probability.

⁷The algorithms also work for ray shooting in a convex polytope, and can also be adjusted for halfspace reporting queries with some additional effort (the published version only deals with emptiness queries).

⁸Randomized expected time.

⁹Dynamization is due to Agarwal and Matoušek [1991].

Box 1: Point Location

The *point location* problem is the following: given a subdivision of \mathbb{R}^d into convex cells, construct a data structure that can quickly determine the cell of the subdivision containing a query point.

Let m denote the combinatorial complexity of the subdivision (the total number of cells of all dimensions). In the plane, several optimal data structures are known, with $O(m)$ space and preprocessing time and with $O(\log m)$ query time, e.g., see Seidel [1991a] for a modern randomized method, or the references to older works in Edelsbrunner [1987]. For dimension 3, a method with a polylogarithmic query time and nearly linear space was described in Preparata and Tamassia [1992], but for dimension 4 and higher, no data structure with a comparable efficiency has been found for general subdivisions.


An arrangement of n hyperplanes in \mathbb{R}^d is a very special subdivision, and here efficient point location methods are known in an arbitrary dimension. The first such algorithm, due to Clarkson [1987], requires $O(n^{d+\epsilon})$ space and preprocessing time, and it performs the location in $O(\log n)$ time. By various technical improvements of the same basic idea, Chazelle and Friedman [1994] obtained an algorithm with the same query time and with optimal $O(n^d)$ space. A conceptually much simpler method was discovered by Chazelle [1993a], with the same space and query time and also with $O(n^d)$ preprocessing. The algorithms from these three papers are based on hierarchies of progressively refined cuttings, similar to the algorithm described in Section 5.

dimension (this is also the case for range searching). This fixed-dimension assumption contrasts, for example, with the theory of linear programming algorithms, where the dimension is comparable to the number of constraints. Let us remark that some recent work (Clarkson [1994] and Clarkson et al. [1993] and others), often motivated by direct practical applications, puts much more emphasis on a good dependence of the running time on the dimension, and may indicate a change in the above-mentioned traditional computational geometry philosophy.

In theoretical research, the efficiency of algorithms is measured, almost exclusively, by the asymptotic order of growth in their complexity (as a function of the input size). One can raise various objections against this criterion, mainly from the practical point of view, but no better theoretical criterion seems to be available.

For simpler algorithms, mainly older ones, effectiveness expressed in this way agrees with the intuitive notion of “how fast the algorithm computes,” and even with the speed of actual implementations for real-life problems. For more complicated algorithms, this touch with reality is often lost. Sometimes it is even apparent that an asymptotic improvement has nothing to do with the speed of computa-

tion for any physically feasible inputs. Attempts to improve an algorithm with $O(n \log \log n)$ complexity to an $O(n \log^* n)$ algorithm might look like a peculiar kind of game.



In defense of such activities one can say that the actual goal is, or at least should be, a better understanding of the problem being studied. A discovery of a complicated and quite impractical algorithm with a good asymptotic complexity indicates the possibility of an improvement, and often a simpler and easily implementable algorithm of a similar complexity is found afterward. A good example is the point location problem in the plane (see “Point Location”). It is instructive to compare the first asymptotically optimal solution in Lipton and Tarjan [1980] to, say, Seidel [1991a]. The ideas from such research can have unexpected application in other areas also. For instance, methods developed for an

asymptotically optimal simplex range-searching algorithm in Matoušek [1993c] (which seems to be a canonical example of a practically useless algorithm in its present form) helped to answer long-standing questions in discrepancy theory (discrepancy of point sets with respect to halfspaces [Matoušek 1994] and the discrepancy of arithmetic progressions [Matoušek and Spencer 1994]).

Model of Computation: Rounding Problems

Algorithms in computational geometry are designed for an ideal computer (model of computation) called *Real RAM*. This is an analog of the usual RAM, i.e., an abstraction of an actual computer programmed in machine code. However, a Real RAM can store an arbitrary real number in one register and perform arithmetic operations with real numbers in unit time, while the usual RAM works with integers whose size is bounded by a polynomial in the input size. In the sequel we will also use the Real RAM model of computation.

The infinite-precision computation in the Real RAM model is very convenient theoretically, but it often behaves quite differently than an actual computation with limited precision. Not only that, the result of a calculation will be imprecise—a carelessly programmed algorithm may give completely wrong results with a limited precision, as the combinatorial information derived from comparisons of imprecise numbers may be erroneous or inconsistent. Difficulties with rounding errors are a very serious potential obstacle (perhaps the most serious one) to successful implementation of geometric algorithms.

Several promising attempts at developing program packages for a sufficiently precise and relatively quick arithmetic have appeared (e.g., Fortune and Van Wyk [1993]). However, we probably have to expect that if a more complicated geometric algorithm is implemented reliably (i.e., in such a way that it cannot give an erroneous result because of rounding errors), it will cause a significant slowdown

compared to a straightforward implementation that ignores these aspects.

General Position Assumption

In many computational geometry algorithms, one has to deal separately with configurations of input objects that are degenerate in some sense; for instance, when three input points lie on a common line. Such exceptions complicate both the description and implementation of algorithms. Luckily for theoreticians, it is known how to systematically avoid such degenerate cases, although at the cost of a slowdown by a constant factor (at least in versions published so far). Conceptually, one perturbs the input objects by infinitesimal amounts to bring them into a general position. Such methods are called *simulation of simplicity*. The idea of infinitesimal perturbations can be realized in various ways; see Edelsbrunner and Mücke [1990], Emiris and Canny [1991; 1992], and Seidel [1994], and the best method is still being sought. Another recent work in this area is by Burnikel et al. [1994].

Simulation of simplicity has theoretical drawbacks also. Since we are in effect replacing a given input problem by a different problem (although an infinitesimally close one), an algorithm run on this different problem may sometimes yield a different answer than the correct one for the original problem, and it may actually be quite complicated to recover a correct answer to the original problem. Also, the infinitesimal perturbation can increase the output size considerably (when computing a convex hull of n points in \mathbb{R}^d , the actual convex hull may have a linear complexity only, while a perturbation can make this complexity as high as $\Omega(n^{\lfloor d/2 \rfloor})$).

Of course, we may say that the input numbers for a “real-world” problem are inaccurate anyway, and so the answer for the perturbed problem is equally as appropriate as the one for the original problem. However, this is not suitable for many applications, where degeneracies in the input are not mere “coincidences,”

but are important for the problem structure.

Deterministic and Randomized Algorithms

The complexity of computational geometry algorithms is most often estimated in the worst case, which means that the estimate must hold for each admissible input of a given size (one may imagine that the input is given by an adversary who knows the algorithm and tries to make it as slow as possible).

Many fewer results concern the average case, where the input is considered as a random variable, and the expected complexity of the algorithm for such an input is estimated (e.g., for input points selected independently of the uniform distribution in the unit square). The most problematic point of the average-case analysis is the choice of the probability distribution on the inputs. Often various natural possibilities exist which give significantly different expected behavior for the algorithm. For geometric problems it is usually much easier to give an algorithm with a good average-case complexity than an algorithm with a comparable worst-case complexity. This is the case for geometric range searching also, which is almost trivial for points uniformly distributed in the unit square.

Older algorithms are, with few exceptions, *deterministic*, which means that their computation is uniquely determined by the input. So-called *randomized algorithms* appeared only recently in computational geometry. These algorithms randomly choose one from several possible continuations at some stage of their computation.² For a fixed input, the

complexity of such an algorithm is a random variable. One estimates its expectation, or further parameters (such as the probability of large deviations from the expectation, *tail estimates*), and then considers the worst case over all inputs (thus the input is chosen by the adversary, but he cannot influence the random decisions of the algorithm).

Sometimes so-called Monte Carlo algorithms are also studied. These may sometimes give a wrong result, but with only a very small probability (again, such a bound holds for every input).

In a seminal paper [Rabin 1976] on randomized algorithms, a typical computational geometry problem was given as one of two examples; at that time computational geometry did not yet exist as a field. The paper by Chew [1985] with a very elegant application of randomness, as well as earlier work by Clarkson also went almost unnoticed. In the last few years, however, randomized algorithms became dominant in computational geometry, promoted by the work of Clarkson [1987; 1988a; 1988b], Haussler and Welzl [1987], Sharir, Chazelle, Edelsbrunner, Guibas (e.g., Clarkson et al. [1990]), Mulmuley [1989; 1991a; 1991b], Seidel [1991a; 1991b], and others. In most cases, randomized algorithms are simpler, more effective, and easier to implement than deterministic ones. For most cases, problem-deterministic algorithms with a similar or only slightly worse asymptotic complexity were found (e.g., Chazelle and Friedman [1990], Matoušek [1990; 1991a], Chazelle [1993a; 1993b], Chazelle and Matoušek [1993]), but for practical purposes the randomized algorithms are still winning because of their simplicity and smaller hidden constants.

2. GEOMETRIC RANGE SEARCHING

Let \mathcal{R} be a system of subsets of the d -dimensional Euclidean space \mathbb{R}^d . The sets of \mathcal{R} are called *ranges*. The following are typical cases:

\mathcal{R}_{orthog} = *axis-parallel boxes*, i.e., all sets of the form $\prod_{i=1}^d [a_i, b_i]$, $a_1, b_1, \dots, a_d, b_d \in \mathbb{R}$.

²In actual implementations, we do not have true randomness at our disposal, and random decisions are simulated using pseudorandom numbers. Then the randomness in the algorithm is restricted to the initial setting of the random number generator, which represents only few random bits. Empirical evidence indicates that even this weak randomness is sufficient for most of the algorithms, and recent work has already confirmed this for some classes of algorithms; see Mulmuley [1992].

\mathcal{R}_{halfsp} = the set of all (closed) halfspaces in \mathbb{R}^d .

$\mathcal{R}_{simplex}$ = the set of all (closed) simplices in \mathbb{R}^d .

\mathcal{R}_{ball} = the set of all (closed) balls in \mathbb{R}^d .

Further, let P be a given n -point set in \mathbb{R}^d . One of the geometric range-searching problems is the following:

Design an efficient algorithm, which, for a given range $R \in \mathcal{R}$, finds the number of points of P lying in R .

If the set P and the range R were given together, and the problem finished by determining the number $|P \cap R|$, it would be simplest to go through all points of P one by one and count those lying in R . In fact, in such a situation we can hardly do anything better. In our problem, however, the point set P is given in advance, and we can prepare some auxiliary information about it and store it in a suitable data structure. Then we will repeatedly be given various ranges of \mathcal{R} as queries. Each such query is to be answered as soon as it appears (on-line) and as efficiently as possible. Assuming that the number of queries will be large, we see that it will be advantageous to invest some computing time into building the data structure (this phase is called *pre-processing*), if this makes answering the query faster.

Example 2.1. Let us look at the situation in dimension $d = 1$, with intervals as ranges. Without any preprocessing, query answering (counting the number of points in a given interval) requires time proportional to n . However, having stored our points in a linear increasingly sorted array, we can answer queries in $O(\log n)$ time, as follows. We locate the position of the endpoints of the query interval among the points of P by binary search, and then we find the required number by a subtraction of indices. Preprocessing, in our case sorting the points of P , can be performed in $O(n \log n)$ time, and $O(n)$ memory is sufficient for storing the data structure.

Counting points in a given range (a *range-counting query*) is only one of possible range-searching problems. Another natural problem is to compute a list of all points of P lying in a query range (we speak of a *range-reporting query*), or we may only want to ask if the query range contains any point of P at all (*range emptiness query*). Also, each point of P can be assigned some weight (e.g., a real number), and we may be interested in the sum of weights of points in a given range, or in the maximum weight.

All these problems are quite convincingly motivated by direct practical applications, most often in various database systems. A rather banal example: the points of P might correspond to employees of some company, with the coordinates representing age, salary, time of employment, etc. Then, in order to find an employee with age between 30 and 35, with at least 8 years of experience, and salary below 4000 dollars per month,³ one might use a query with an axis-parallel box from \mathcal{R}_{orthog} . A circular range query might, for example, serve to locate airports where a defective airplane may still succeed in landing. Perhaps more important than such direct applications are applications of geometric range searching as subroutines in the design of algorithms for more complicated geometric problems.

In more recent papers, a unifying generalization of various range-searching problems is usually investigated. We assume that every point $p \in P$ is assigned a weight $w(p) \in S$, where $(S, +)$ is some commutative semigroup (common to all the points). The objective of a query is to find the sum of weights of all points of P lying in a given range,

$$\sum_{p \in R \cap P} w(p).$$

For example, for range-counting queries, $(S, +)$ will be the natural numbers with

³The reader may supply a currency unit according to his or her own preference.

addition, and all weights will be equal to 1. For queries on maximum weight, the appropriate semigroup will be the real numbers with the operation of taking a maximum of two numbers. For emptiness queries, we may choose the boolean values {false, true} for S , with the operation of logical disjunction (OR); all the point weights will be “true.”

Concerning computational aspects, we usually assume that the weights can be stored in a single computer word, and that the semigroup operation can be executed in constant time.

The range-reporting queries have a somewhat special position. We could also include them in the semigroup model (discussed above)—the semigroup $(S, +)$ would be the set of all subsets of P ; the operation would be the union; and the weight of each point $p \in P$ would be the one-point set $\{p\}$. However, in reasonable models of computation, a subset cannot be represented in a constant amount of memory, and the union operation requires a nonconstant time. From the algorithmic point of view, a further specialty of range reporting is the following.

If the answer consists of k points, then we need time of order k only to output the answer. Thus, we may spend further $O(k)$ time for auxiliary computations, without decreasing the overall asymptotic complexity. If we know that the answer is going to be large, we can afford to compute it slowly. This was observed and cleverly applied in Chazelle [1986]; Chazelle calls this the *filtering search*. The query complexity of range-reporting algorithms is usually expressed in the form $O(f(n) + k)$, where k is the number of points in the answer, and $f(n)$ is some function of the total number of points in P , called the *overhead*.

Let us return to the data structure from Example 2.1, which was a simple sorted array. We can use it to answer interval-reporting queries in $O(\log n + k)$ time—and it is easy to generalize it for queries seeking the sum of real weights in a given interval. In this case, as well as for the interval-counting queries, we make use of the possibility to

subtract the weights (that is, the semigroup $(S, +)$ can be embedded into a group). If this is not possible (in cases such as finding the maximum weight), we need a different (and more complicated) data structure.

Here is a simple and well-known example of such a data structure:

Example 2.2. Consider the points of $P \subset \mathbb{R}^1$ in sorted order, and build a balanced binary tree T with the points of P as leaves. With every node v of T , store the total weight of points in the leaves of the subtree rooted at v (note that these points form an interval in P ; such intervals are called the *canonical intervals*). Given a query interval (a, b) , we search the position of a and b among the leaves of T , and the search paths give us a decomposition of the set of points of P lying between a and b into a disjoint union of $O(\log n)$ canonical intervals. Indeed, let π_a and π_b denote the search paths for a and b , respectively, and let x be the node where they branch. Then for every node $v \in \pi_a$ lying below x and such that the next node of π_a is the left son of v , we take the right subtree of v , and it will define one of the canonical intervals. We proceed symmetrically for the portion of π_b below x , obtaining the desired decomposition. Thus, to answer the query, it remains to sum the precomputed weights of the canonical intervals in the decomposition. Note that this does not use the weight subtraction, so that we can find, e.g., the point of maximum weight in a given interval.

An analogous situation appears for more complicated geometric range-searching problems also. Reporting queries and queries with weights, which can be subtracted, sometimes allows a somewhat simpler and/or more efficient solution than the general case, whose prototypes are the queries asking for maximum weight (e.g., see Matoušek and Welzl [1992] for an example of a simple two-dimensional triangle range-searching algorithm which uses subtraction crucially). For the case of subtraction, we can usually express the answers for more

complicated ranges using the answers for several simpler ranges. For instance, in our one-dimensional example, we have implicitly expressed a query interval as the difference of two semiinfinite intervals.

For a full specification of a range-searching problem, we need, in addition to the already mentioned objects (the set P , the semigroup $(S, +)$, the weights of points, and the set of admissible ranges \mathcal{R}), limits on the maximal permissible storage and preprocessing time for the data structure.

Such a limitation does not show up in Example 2.1, where the query time $O(\log n)$ is optimal and, at the same time, to answer queries at all, we need at least linear storage. However, for halfspace and simplex range searching, one cannot achieve both linear storage and a logarithmic query time, and so we have to choose the most efficient query-answering algorithm, depending on the amount of memory and preprocessing time we have at our disposal. Memory requirements are usually considered more important; this has roots in database applications, where preprocessing needs to be done only once, and we can spend a relatively long time at it, while the memory or disk space is allocated permanently and in each copy of the database. On the other hand, in many applications the preprocessing time is as critical as the storage.

In this survey we will consider the simplex and halfspace range-searching problems. These problems turned out to be crucial in computational geometry. They are even universal in some sense, since many other problems with more general ranges can be reduced to them (see Agarwal and Matoušek [1994], Yao and Yao [1985] and Section 6.4).

The algorithms for simplex and halfspace range searching problems have been applied to problems which look much more complicated. A nice example is in De Berg et al. [1994], which was motivated by a computer graphics problem. The range-searching problem, which the authors needed to solve as part of

their hidden-surface removal algorithm, is as follows.

Problem 2.3. Given a set s_1, \dots, s_n of segments in \mathbb{R}^3 , we imagine that there is a semiinfinite curtain $c(s_i)$ hanging from each segment s_i downward; we have

$$c(s_i) = \{(x, y, z) \in \mathbb{R}^3; \\ (x, y, z_0) \in s_i \text{ for some } z_0 \geq z\}$$

(the curtains may intersect). We want to design a data structure for answering queries of the following type. Given a point o and a direction θ , find the first curtain hit by the ray sent from o in direction θ .

Both the original solution of this problem in De Berg et al. [1994] and a more efficient method described in Agarwal and Matoušek [1993] reduce this problem to a combination of simplex range searching and halfspace range searching.

The orthogonal range-searching problems (the ones with axis-parallel boxes as ranges) are no less important than simplex range searching, perhaps even more important for direct applications. In this survey we will not consider them in detail, but for the reader's convenience we briefly summarize the main results concerning orthogonal range searching. The basic idea of data structures for this problem, the so-called *range tree*, was described by Bentley [1980]. Its straightforward use gives query time $O(\log^d n)$ with $O(n \log^{d-1} n)$ storage and preprocessing time in dimension d . Various improvements in some parameters, mostly by logarithmic factors, were achieved by Willard [1985], Willard and Lueker [1985], Gabow et al. [1984], and Chazelle [1986; 1988]. The exact complexities differ depending on the type of the problem (range-counting queries, range-reporting queries, general semigroup queries, etc.) and on the model of computation (pointer machine, RAM, etc.). Lower bounds for the computational complexity of orthogonal range searching were given by Chazelle [1990a; 1990b], and they match the known upper bound almost exactly. A

variant where the points lie on a grid of a bounded size was studied in Overmars [1988]; in this case somewhat better results were obtained.

3. INTUITION AND LOWER BOUNDS

In this section we begin to consider the simplex and halfspace range-searching problems in more detail. Results about their computational complexity can be summarized as follows.

Consider a simplex range-searching problem for an n -point set $P \subset \mathbb{R}^d$, with weights from a semigroup $(S, +)$, and with storage and preprocessing time at most m , where m lies in the range from n to approximately n^d . Then the query time is approximately

$$\frac{n}{m^{1/d}}. \quad (1)$$

The word “approximately” in the previous sentence means “up to a multiplicative factor bounded by $O(\log^c n)$, c a constant.”

In particular, for an approximately linear storage the query time is approximately $n^{1-1/d}$, and in order to achieve a polylogarithmic query time, one needs space and preprocessing approximately n^d . The complexity of halfspace range queries is believed to be very similar, except for a few special cases. The most significant such case is halfspace range reporting, which is considerably faster. For instance, one can achieve $O(\log n + k)$ query time with approximately linear storage in dimensions 2 and 3 (k denotes the number of points reported; see Section 6.1).

In this section we formulate known lower bounds, which show that under certain restrictions on the algorithm and the model of computation the query time (1) is approximately optimal. It is possible (although currently it does not seem very likely) that one could circumvent these lower bounds somewhat using algorithms that do not satisfy the corresponding assumptions. It would not be the first such case in theoretical computer science.

First we will try to give the reader some intuitive explanation as to where Eq. (1) comes from. The explanation is quite far from a proof, and in reality the lower and upper bounds work in a more complicated manner; it is just meant for basic orientation. We consider the two extreme cases, polylogarithmic query time and roughly linear space.

Logarithmic Query Time

First we consider halfspace queries. It is not difficult to see that for an n -point set P in a general position there are $\Theta(n^d)$ different subsets of the form $P \cap R$, where R is a halfspace. (This is best seen in the dual setting, where distinct subsets correspond to distinct cells in an arrangement of hyperplanes. (See boxes 2 and 3.) Storage of the order n^d thus means that we can store the answers for all essentially different halfspaces that can ever appear. Actual algorithms are indeed based on this principle. A naive attempt on extending this idea to simplex range searching results in a much larger space than n^d . A suitable method preserving storage close to n^d is more complicated and was discovered only recently [Chazelle et al. 1992].

Approximately Linear Storage

Here we assume that the set P is chosen randomly, by n independent random draws from the uniform distribution in the unit square. (We consider the planar case first.) We say $t = \lfloor \sqrt{n} \rfloor$, and we cover the unit square by a $t \times t$ square grid, each grid square having side $1/t$ (see Figure 1). With high probability, almost every grid square then contains only a small number of points of P (bounded by a constant).

Let R be a given halfplane. We note that the boundary line h of R intersects at most $2t$ squares of the grid. (If its slope is at most 1, then it intersects at most 2 squares in every column, and for slope > 1 , we apply a similar argument with rows.)

For the squares intersected by h we go through all the points of P lying in them,

Box 2: Geometric Duality

In the plane, duality is a mapping assigning points to lines and lines to points. There are various versions of duality considered in the literature; for most applications, the differences among them are unessential. For definiteness, we work with the duality \mathcal{D} defined as follows: a point $p = (a, b)$ has a dual line $\mathcal{D}(p) = \ell = \{(x, y); y = 2ax - b\}$, and conversely the image $\mathcal{D}(\ell)$ of such a line ℓ is the point $p = (a, b)$. This defines the image of every point and of every line, except for vertical ones (vertical lines correspond to points in the infinity of the projective plane).

It is easy to check that for any point q and a nonvertical line h , q lies on h iff the point $\mathcal{D}(h)$ lies on the line $\mathcal{D}(q)$. Moreover, q lies above h iff $\mathcal{D}(h)$ lies above $\mathcal{D}(q)$. These properties of geometric duality often allow a much more intuitive formulation of geometric problems by passing to dual objects.

Duality is similarly defined between the points and hyperplanes in d -dimensional space. The duality \mathcal{D} maps the point $p = (a_1, a_2, \dots, a_d)$ to the hyperplane $\mathcal{D}(p) = \{(x_1, x_2, \dots, x_d); x_d = 2a_1x_1 + 2a_2x_2 + \dots + 2a_{d-1}x_{d-1} - a_d\}$ and conversely. An analog of the above-mentioned properties for the planar case holds (duality preserves incidences and above/below relations). For example, if we want to count points lying below a hyperplane, the dual problem is counting the number of hyperplanes above a given point. Further information about geometric duality can be found in Edelsbrunner [1987].

Box 3: Hyperplane Arrangements

Let us consider a finite set H of lines in the plane. These lines divide the plane into convex sets (called *cells*; sometimes the word *faces* is also used) of various dimension; see the figure. The cells of dimensions 0, or 0-cells, for short, are the intersections of the lines of H , and we call them *vertices*. If we remove all vertices lying on a line $h \in H$, the line h is split into two open semiinfinite rays and a finite number of open segments. These segments and rays form the 1-cells or *edges*. Finally, by removing all the lines of H the plane is partitioned into open convex polygons (also into unbounded ones), which are the 2-cells. Similarly, a finite set H of hyperplanes in \mathbb{R}^d defines a decomposition of \mathbb{R}^d into cells of dimensions $0, 1, \dots, d$.

This decomposition is a cell complex in the sense of algebraic topology, and it is called the *arrangement of H* . For a fixed d , the total number of cells of all dimensions is bounded by $O(|H|^d)$. This bound is best possible for H in a general position; in this case it is not too difficult to obtain an exact formula for the number of cells of each dimension. Hyperplane arrangements are one of the basic objects of study in computational geometry, e.g., see Edelsbrunner [1987] for more information.




Figure 1. A simple halfspace range-searching method for uniformly distributed point sets

and for each such point we test its membership in R . The uniform distribution implies that the number of points processed in this phase is $O(t) = O(\sqrt{n})$. (Such points are marked as full circles in Figure 1.)

It remains to account for the weight of points in grid squares that are completely contained in R . This can be done row by row, using the fact that such squares form a contiguous interval in every row. The total weights of points in each such segment of each row are computed in advance; thus we only need a constant time per row for query answering. The total memory requirement is $O(n)$.

For a higher dimension d we can proceed quite similarly, dividing the unit cube into a grid of cubes with sides $n^{-1/d}$. The bounding hyperplane of any given halfspace R intersects only $O(n^{1-1/d})$ grid cubes. The cubes completely contained in the query halfspace can be processed by columns parallel to one (arbitrarily chosen) coordinate axis. In this way we get a data structure with $O(n)$ storage and $O(n^{1-1/d})$ query time

for uniformly distributed point sets in the unit cube, as required by Eq. (1).

This time also the generalization to simplex range searching is straightforward. We associate a suitable one-dimensional data structure for range searching in intervals with every column of the grid; we leave the details to the reader.

It is quite conceivable that this simple data structure is the most practical option for simplex range searching with linear space, at least for roughly uniformly distributed sets. (This method resembles the *quadtrees*; see Preparata and Shamos [1985].) For point sets that are not uniformly distributed, this simple approach fails, and all known methods with query time close to $n^{1-1/d}$ are considerably more complicated.

Lower Bounds

Lower bounds are due to Chazelle and his students; previous weaker results were obtained by Fredman [1981]. The basic work [Chazelle 1989] bounds from below the computational complexity of simplex range searching with weights from a suitable semigroup. In order to formulate this result, we have to define an appropriate model of computation, the so-called *arithmetic model*, originally introduced by Yao and Fredman as a means for lower-bound proofs.

Roughly speaking, we can say that the arithmetic model considers only the number of semigroup operations needed to compute the answer to a query. Certain partial sums of the point weights are precomputed and stored in memory; such partial sums are called *generators*. Their number cannot exceed the prescribed storage, m . The arithmetic model does not consider at all various auxiliary operations of the algorithm needed in query answering, such as finding out which generators are to be added to yield the result. This is a strength of lower bounds in the arithmetic model. On the other hand, this model puts various restrictions on the query-answering algorithm. One of the most significant restrictions is the impossibility of sub-

tracting the weights, even if the weights happen to belong to a group. Let us now pass to exact definitions.

A commutative semigroup $(S, +)$ is called *faithful* if the following holds. For any $n > 0$, any two distinct subsets $A, B \subseteq \{1, 2, \dots, n\}$ and natural numbers $\alpha_i > 0$ ($i \in A$), $\beta_j > 0$ ($j \in B$), there exist elements $s_1, \dots, s_n \in S$ so that

$$\sum_{i \in A} \alpha_i s_i \neq \sum_{j \in B} \beta_j s_j.$$

In other words, two linear forms over variables ranging over S are never identically equal unless they have the same set of variables. Faithfulness is a quite weak condition. An example of a semigroup that is not faithful is $\mathbb{Z}/2\mathbb{Z}$, i.e., $\{0, 1\}$ with addition modulo 2, where the identity $2x = 2y$ holds. On the other hand, real numbers, both with the operation of addition and the operation of a pairwise maximum, do form faithful semigroups.

In the sequel, let $(S, +)$ be a fixed faithful semigroup. Any linear form $g(s_1, \dots, s_n) = \sum_{i=1}^n \alpha_i s_i$, where s_i are variables ranging over S and α_i are non-negative integers, is called a *generator*.

It turns out to be advantageous to consider geometric range-searching problems in a more abstract form. To each problem with point set $P = \{p_1, p_2, \dots, p_n\} \subset \mathbb{R}^d$ and with a set \mathcal{R} of admissible ranges, we assign a set system (P, \mathcal{P}) , where \mathcal{P} is the system of all subsets of P definable by ranges from \mathcal{R} , that is,

$$\mathcal{P} = \{P \cap R; R \in \mathcal{R}\}. \quad (2)$$

Then we work with this set system only, thus restricting the geometric contents of the problem to a minimum. This more abstract approach turned out to be fruitful for geometric range searching (but not only for geometric range searching); see Haussler and Welzl [1987] and Matoušek [1991a].

A set $\Gamma = \{g_1, \dots, g_m\}$ of generators is called a *scheme* for the set system (P, \mathcal{P}) , if for every set $Q \in \mathcal{P}$ there exist nonneg-

ative integer coefficients β_1, \dots, β_m such that for *any* choice of weights $w(p_1), \dots, w(p_n) \in S$, we have

$$\sum_{p \in Q} w(p) = \sum_{i=1}^m \beta_i g_i(w(p_1), \dots, w(p_n)). \quad (3)$$

It is thus required that in using the generators of the scheme one express the weight of any set $Q \in \mathcal{P}$, and this expression must be of the same form for any choice of weights of the points of P .

A scheme Γ for (P, \mathcal{P}) is called a (t, m) -scheme if $m = |\Gamma|$ and for every $Q \in \mathcal{P}$ the coefficients in (3) can be chosen in such a way that at most t of the numbers β_1, \dots, β_m are nonzero. Then it is natural to define that the geometric range-searching problem with the point set P , set of ranges \mathcal{R} , and with weights from $(S, +)$ has query complexity at least t in the arithmetic model for storage m if there is no $(t - 1, m)$ scheme for (P, \mathcal{P}) , where \mathcal{P} is as in (2).

Simple Example. Let us pause for a moment to see what the generators might look like in an actual algorithm. Considering Example 2.2, we have one generator for every canonical interval, namely the sum of point weights in that interval. Hence each generator is just the sum of weights over some subset of P , and the answer to a query is computed by expressing the point set in the query range as a disjoint union of canonical subsets. Therefore all coefficients in (3) are either 0 or 1. Other range-searching algorithms for semigroup weights also use this more special form of generators and answer computation. In fact we are not aware of any single instance where the more general form allowed by the arithmetic model would be used.

Lower bounds in the arithmetic model hold only for algorithms computing the answer (the total weight of points in a query range) using a scheme in the sense just defined. There might thus exist, in principle, a better algorithm using the specific weights in a given range-search-

ing problem (while an algorithm covered by the arithmetic model must work “uniformly” for any choice of weights). Lower bounds do not apply to algorithms using weight subtraction either (if $(S, +)$ is a group). Proving lower bounds valid for this group case is one of the main challenges in this area; another is a proof of lower bounds for emptiness queries with halfspaces or simplices. Let us remark that an orthogonal range-searching-type problem is known where the complexity in the group model is provably better than in the semigroup model (see Chazelle and Rosenberg [1991]).

Now that we have clarified all these definitions, we can state the main result [Chazelle 1989].

THEOREM 3.1 (CHAZELLE). *Let $(S, +)$ be a faithful semigroup. Then for any fixed dimension d and parameters n, m there exists an n -point set $P \subset \mathbb{R}^d$ such that the simplex range-searching problem with point set P , weights from S , and storage at most m has query complexity at least*

$$\Omega\left(\frac{n}{\log nm^{1/d}}\right) \quad (4)$$

(for $d \geq 3$), resp., at least

$$\Omega\left(\frac{n}{\sqrt{m}}\right) \quad (5)$$

for $d = 2$ in the arithmetic model.

In fact, the proof of this theorem gives somewhat stronger results. First of all, P is not an artificially constructed pathological set. Rather it is sufficient to choose it randomly from the uniform distribution in the unit cube; we obtain a “hard” set with high probability. Also the hard query simplex need not be chosen in any special way. The proof uses, instead of simplices, so-called *slabs*, which are the parts of space bounded by two parallel hyperplanes. The proof shows that for a randomly chosen slab of a suitable width (among the slabs intersecting the unit cube) the query complexity is at least that shown by the lower bound formulas.

In this sense, (4) and (5) bind not only the worst case, but also the average case.

It is quite likely that Eq. (4) holds without the logarithmic factor in the denominator as well (as in dimension 2). Such an improvement has an interesting relation to a generalization of a famous problem of combinatorial geometry, the so-called Heilbronn problem. The Heilbronn problem itself can be formulated as follows:

Problem 3.2. For a set $P \subset \mathbb{R}^2$, let $a(P)$ denote the area of a smallest triangle with vertices in P . What is the asymptotic behavior of the function $a(n) = \sup\{a(P); P \subset [0, 1]^2, |P| = n\}$?

This is a very nice problem and has been worked on by many excellent mathematicians. Its complete solution still seems to be quite far off (see, e.g., Komlós et al. [1982] and Roth [1976] for a survey of results and references). The following is a generalization related to Chazelle's lower-bound proof method:

Problem 3.3. For a set $P \subset \mathbb{R}^d$, denote by $a_d(P, k)$ the smallest volume of the convex hull of a k -tuple of points of P . What is the asymptotic behavior of the function $a_d(n, k) = \sup\{a_d(P, k); P \subset [0, 1]^d, |P| = n\}$?

Chazelle has shown that for a suitably chosen set P in the unit cube, the volume of the convex hull of each k -tuple is at least proportional to k/n for any $k \geq c \log n$, with a sufficiently large constant c . In other words,

$$a_d(n, k) = \Omega(k/n) \quad (6)$$

for $k \geq c \log n$. This is essentially a result about uniform distribution of the set P , stating that no k -tuple is too clustered (in the sense of volume). If (6) could also be proved for smaller k , an improvement of (4) would follow immediately.

Such an improvement may not be easy. Known results for the Heilbronn problem imply that (6) is false for $d = 2$, $k = 3$ (which may contradict intuition somewhat). However, it is not known that (6) would not hold for larger but constant k . On the other hand, this connection shows

that by finding an algorithm for simplex range searching with $o(n/m^{1/d})$ query complexity in the arithmetic model, one would also get a nontrivial result for the difficult combinatorial problem mentioned above.

At this point the reader might ask "where does the better bound (5) in the plane come from when the generalized Heilbronn problem is open for dimension 2 as well as for higher dimensions?" The main reason is that, in dimension 1, it is easy to construct a point set satisfying (6) for every $k \geq 2$. One particular construction, although not the most straightforward one, is to take a random uniformly distributed set and throw out a suitably selected half of the points. In Chazelle's proof, subsets of P defined by slabs are considered, and what matters is the one-dimensional distribution of the orthogonal projection of each such subset onto the axis of the corresponding slab. To obtain (5) it suffices to show that, for most of the slabs, one can select half of the points from the projection with a good one-dimensional distribution in the sense of (6).

In spite of Theorem 3.1, one might hope for better algorithms in some particular cases, such as for halfspace range searching, or for simplex emptiness queries. A better algorithm is known for only one special case, namely for halfspace emptiness queries, or for halfspace range-reporting queries (see Section 6.1). On the other hand, it seems that both simplex emptiness queries and halfspace range counting should be approximately equally difficult as the general simplex range-searching problem. This was partially substantiated in Brönnimann et al. [1993] and Chazelle and Rosenberg [1992]. Brönnimann et al. [1993] showed that under similar assumptions, as in Theorem 3.1, the *halfspace* range-searching problem for an n -point set with storage at most m has query complexity at least

$$\Omega\left(\left(\frac{n}{\log n}\right)^{1-(d-1)/d(d+1)} \Big/ m^{1/d}\right)$$

in the arithmetic model. This bound might still be lower than actual complexity, but at least it shows that the halfspace range-searching problem with arbitrary weights is more difficult than the halfspace emptiness problem, where a query time of approximately $n/m^{1/\lfloor d/2 \rfloor}$ can be achieved.

Chazelle and Rosenberg [1992] consider the simplex range-reporting problem. Their model of computation is different than in the previous results: they use the so-called *pointer machine*. They prove that, for any data structure occupying space at most m , there exists a halfspace for which the reporting query requires time at least $\Omega(n^{1-\epsilon}/m^{1/d} + k)$, where k denotes the number of points in that halfspace.

The papers of Chazelle et al. on lower bounds may please the reader as a nice piece of mathematics, but the conclusions for the simplex range-searching problem are quite pessimistic. For instance, if we want to improve the query complexity K -times, compared to the trivial algorithm (consisting of inspection of every point of P) in dimension 10, storage costs are of the order K^{10} . Since we pass from a trivial algorithm to a more complicated one, we may expect that the hidden constants of proportionality will work against us. Nontrivial algorithms can thus be practically useful only for a really small dimension. This is, however, a frequent feature of computational geometry algorithms, as we have already mentioned in the Introduction.

4. ALGORITHMS WITH APPROXIMATELY LINEAR SPACE

We divide our discussion of simplex range-searching algorithms into two parts. First we consider algorithms using linear or nearly linear storage, then algorithms attaining a short—polylogarithmic—query time. As was explained above, the latter algorithms require a large space, of the order n^d . These two groups of algorithms have been investi-

gated separately in the literature, and more attention has been paid to linear space algorithms, which are usually more complicated and more interesting. Algorithms with memory requirements in between these two extremes can often be obtained by a more or less straightforward combination of algorithms of the two types.

Most nontrivial algorithms with linear space (nontrivial meaning with substantially sublinear query time) are based on the idea of *partition trees*, due to Willard [1982]. We will explain the idea in a simple form for halfplane range searching.

Let P be an n -point set in the plane. For simplicity we assume that n is of the form 4^k and that P is in general position. Let ℓ_0 be an arbitrarily chosen line halving P ($n/2$ points lie below ℓ_0 and $n/2$ points above ℓ_0). The classical *ham-sandwich cut theorem* (see Edelsbrunner [1987]) guarantees the existence of another line ℓ_1 such that each of the four regions A_1, A_2, A_3, A_4 into which ℓ_0, ℓ_1 partitions the plane contains exactly $n/4$ points of P ; see Figure 2(a). Consider an arbitrary halfplane R with the bounding line h . It is easy to see that h cannot intersect all four regions A_1, A_2, A_3, A_4 , and hence one of them, A_i , lies completely outside R or completely inside R .

If we precompute the total weight of points in each A_j , we can process all points of the region A_i missed by h in a single step, thus saving 25% of the work on query answering, compared to the trivial method. This is not significant for asymptotic complexity, but a similar saving can be repeated recursively in each of the remaining three regions. Continuing in a similar manner but in greater depth, the complexity decreases significantly.

A data structure based on this idea, a partition tree, has a root containing the coordinates of the partitioning lines ℓ_0, ℓ_1 as well as the total weights of points in the four regions. The root has four sons, corresponding to the four regions A_1, \dots, A_4 , and the i th son is the root of a partition tree constructed simi-




Figure 2. (a) Ham-sandwich cut theorem, (b) two levels of the recursive construction. For the halfplane R , the regions bounded by thick lines are processed as wholes.

larly for the set $P \cap A_i$. This recursive construction terminates after reaching small-enough sets, say one-point ones.

To answer a query with a halfplane R , we start in the root of the partition tree. We find the region missed by the boundary of R , and we process it immediately (ignoring it if it lies outside R , otherwise storing its weight to a global variable for accumulating the total weight). For the remaining three regions, we proceed similarly in the corresponding three sons of the root. When we reach a leaf of the partition tree with a trivially small set, we process the point stored there directly.

If $T(n)$ denotes the query time for an n -point set, we get the recurrence

$$T(4^k) \leq C + 3T(4^{k-1}),$$

with an initial condition $T(1) \leq C$, where C denotes a suitable constant. This gives a bound $T(4^k) = O(3^k)$. Both the construction and the analysis can be extended to an arbitrary value of n , and we obtain $T(n) = O(n^{\log_4 3}) \approx O(n^{0.792})$ —a significantly sublinear function indeed. It is easy to verify that the data structure described occupies $O(n)$ space only.

The partition tree defined above can also be used for answering triangular range queries. The query-answering algorithm is very similar. In a current node of the partition tree, we process the regions missed by the boundary of the query triangle σ directly (since they lie completely inside σ or completely outside σ), and for regions intersected by the boundary of σ we proceed recursively in the appropriate sons of the current node.

We cannot argue here that in every node the recursion visits at most three of its sons. But we note that σ is an intersection of three halfplanes, and it is not difficult to see that any node of the partition tree visited by the query-answering algorithm for the triangle σ will also be visited when answering the query with at least one of the three halfplanes. Thus the query time for triangles is of the same order as the query time for halfplanes.

Easy Improvements and Generalizations of Partition Trees

The idea of a partition tree is simple but by no means obvious. Improving the con-

struction described above and generalizing it to higher dimensions were not easy either (its history somewhat resembles athletic records). Willard himself found a better construction than the one we have described, with query time of $O(n^\alpha)$ with $\alpha \approx 0.774$.

Edelsbrunner and Welzl [1986] improved the method above by a better application of the ham-sandwich cut theorem. Instead of dividing each of the regions A_1, \dots, A_4 independently, they bisect both A_1 and A_4 simultaneously by a single line, and similarly for A_2 and A_3 . Proceeding recursively in this way, they obtain the recurrence $T(n) \leq C + T(n/2) + T(n/4)$, leading to the bound $\alpha \approx 0.695$. The preprocessing time for the partition tree construction is $O(n \log n)$, and it relies on a sophisticated linear time algorithm for the ham-sandwich cut theorem for linearly separated point sets due to Megiddo.

A first generalization for dimension 3, with $\alpha \approx 0.98$, was obtained by Yao [1983], who showed the possibility of partitioning a point set in \mathbb{R}^3 into 8 parts of relative size at most $1/24$ by 3 planes. This was improved by Dobkin and Edelsbrunner [1984] ($\alpha \approx 0.916$), Edelsbrunner and Huber [1984] ($\alpha \approx 0.909$), and Yao et al. [1989] ($\alpha \approx 0.8988$). The last authors rediscovered a result of Hadwiger, namely that any point set in \mathbb{R}^3 can be partitioned into 8 equal parts by 3 planes. They organize the recursive construction in a similar way, as Edelsbrunner and Welzl [1986]. Avis [1984] noted that a partitioning of a point set in \mathbb{R}^d into 2^d equal parts by d hyperplanes is not possible in general for $d \geq 5$. For $d = 4$, the existence of such a partition is still open.

Cole [1985] found a construction in dimension 4 giving $\alpha \approx 0.977$, and shortly after that Yao and Yao [1985] showed the existence of a nontrivial simplex range-searching algorithm with linear space in any fixed dimension, although their exponent $\alpha = \lceil \log_2(2^d - 1) \rceil / d$ differs from 1 by a very small amount only. Let us describe the beautiful construction of Yao and Yao's partition scheme.

First, the discrete point set P is re-

placed by a continuous, everywhere positive mass distribution μ : each point $p \in P$ is replaced by a small dense ball of radius ϵ and weight $1 - \epsilon$ plus a light nebula of total weight ϵ spreading out to infinity. It is easy to show that if $\epsilon > 0$ is small enough, a good partition scheme for μ will also be good for the original point set P . This passage to continuous setting enables us to apply topological arguments (and it is used, for example, in proofs of the ham-sandwich cut theorem).

The construction of the partition scheme proceeds by induction on the space dimension d . Given μ , one constructs a point $C(\mu)$ (the *mass center*) and a set $\Pi(\mu)$, which consists of a finite number of pieces of hyperplanes and whose removal partitions \mathbb{R}^d into the desired regions. For $d = 1$, $C(\mu)$ is the (unique) point bisecting μ and $\Pi(\mu) = \{C(\mu)\}$. Now let $d > 1$, and let S be the hyperplane perpendicular to the x_d -axis and bisecting μ . For any vector v with a positive x_d -coordinate, we define two $(d - 1)$ -dimensional mass distributions in S , denoted by μ_v^+ and μ_v^- . The mass distribution μ_v^+ arises by projecting the mass from the halfspace above S into S in direction of $-v$, and similarly μ_v^- arises by projecting the mass from the lower halfspace along $+v$. The hyperplane S can be identified with \mathbb{R}^{d-1} , and thus the points $C(\mu_v^+), C(\mu_v^-) \in S$ are defined. A key claim Yao and Yao prove by topological means is that there exists v such that $C(\mu_v^+)$ and $C(\mu_v^-)$ coincide. For such a v , we set $C(\mu) = C(\mu_v^+)$ and

$$\begin{aligned} \Pi(\mu) &= S \cup \{x + tv; t > 0, \\ &x \in \pi(\mu_v^+)\} \cup \{x - tv; t > 0, \\ &x \in \Pi(\mu_v^-)\}. \end{aligned}$$

This finishes the definition of the partition; it turns out that $C(\mu)$ and $\Pi(\mu)$ are unique. Moreover, the partition defines 2^d regions, each containing 2^{-d} fraction of the total mass, and that each hyperplane avoids at least one of the regions.

All the algorithms mentioned above are based on theorems of a topological flavor, such as the ham-sandwich cut theorem

or the Borsuk-Ulam theorem. They always use some *partition scheme*, which is a way of partitioning the space into several regions (usually bounded by hyperplanes), depending on the given point set. The crucial property of such a partition is, that for any hyperplane, the regions intersected by it contain significantly fewer points than the whole set. The quality of a partition scheme is determined by the maximum number of regions which can be intersected by a hyperplane and by the maximum number of points contained in such regions. From a given partition scheme one then builds a partition tree in a more or less standard way. These algorithms usually have a polynomial, but slow, preprocessing (especially in higher dimensions).

Many researchers tried to find more efficient partition schemes, but it seems that simple constructions could not be improved; and the analysis of more complicated constructions was too involved.

A Randomized Partition Scheme

A significant breakthrough was made by Haussler and Welzl [1987]. They were among the first (together with Clarkson and a few others) to bring to a larger extent than previously probabilistic methods into computational geometry. They introduced an abstractly formulated range-searching problem (specified as a set system in Eq. (2)), and started to build a theory of such abstract problems. The partition scheme used in their algorithm is also of a new type, with a large and adjustable number of regions. (Asymptotic query complexity improves with an increase in the number of regions.) The partitioning scheme can be described quite simply. (We do it in the plane.) For a suitable large constant r , pick a random r -tuple of points of P and draw all the $\binom{r}{2}$ lines determined by these points. These lines partition the plane into the desired regions. Haussler and Welzl proved that, with high probability, the partition tree arising in this way guarantees query time $O(n^\alpha)$ with $\alpha = 2/3 + \delta$, where δ is a positive constant tending to 0 with $r \rightarrow \infty$. In dimen-

sion d one gets $\alpha = 1 - 1/[d(d-1) + 1] + \delta$. The partition tree can be constructed in $O(n \log n)$ expected time by a randomized algorithm; later a deterministic construction of the same asymptotic complexity was also found [Matoušek 1990].

Spanning Trees with Low Crossing Numbers

Query time close to \sqrt{n} in the plane was first achieved by Welzl [1988]. He abandoned the principle of a partition tree constructed recursively using a partition scheme with a constant-bounded number of regions, and replaced it with the idea of a *spanning tree with a low crossing number*. This notion is worth explaining.

Consider a point set P in the plane and some spanning tree T on P , i.e., a graph-theoretic tree having P as its vertex set. We say that a halfplane R *crosses* an edge $\{u, v\}$ of T , if $|R \cap \{u, v\}| = 1$. The *crossing number of T with respect to the halfplane R* is the number of edges of T crossed by R , and the *crossing number of T* is the maximum of crossing numbers of T with respect to all halfplanes.

Let T be a spanning tree on P with a (possibly small) crossing number κ . For simplicity, let us assume moreover that T is a path. (From known algorithms for constructing a spanning tree with low crossing numbers, we can always obtain a path if we want one.) For any given halfplane R , there are at most $\kappa + 1$ components of the induced subgraph of T on the set $P \cap R$. Each such component is an interval along the path T . Thus, having a suitable one-dimensional interval range-searching data structure for intervals along T , we can answer the halfplane query by performing $\kappa + 1$ interval queries; provided, of course, that we know the edges of T crossed by the halfplane R . This explains the significance of the spanning trees with a low crossing number for halfplane range searching.

The notion of a crossing number of a spanning tree can immediately be generalized to an arbitrary dimension d (with halfspaces in place of halfplanes), and it even makes sense for an arbitrary set

system (P, \mathcal{P}) . (A set $Q \in \mathcal{P}$ crosses an edge $\{u, v\}$ if $|\{u, v\} \cap Q| = 1$, and one continues as in the previous definition.) Welzl proves a general existence result for spanning trees with a low crossing number in this abstract formulation. The resulting crossing number is a function of a certain parameter of the considered set system, the so-called *dual shatter function*.⁴

We cannot resist giving the definition of this important notion. Let $\mathcal{A} \subseteq \mathcal{P}$ be a subsystem of sets from \mathcal{P} . Let us call two points $x, y \in P$ \mathcal{A} -equivalent if $x \in Q \Leftrightarrow y \in Q$ for every $Q \in \mathcal{A}$. Then, for an integer m , we define the value of the dual shatter function $\pi_{\mathcal{A}}^*(m)$ as the maximum number of \mathcal{A} -equivalence classes on P over all m -element $\mathcal{A} \subseteq \mathcal{P}$. For example, for the set system defined by triangles on an n -point set in the plane, the dual shatter function satisfies $\pi_{\mathcal{A}}^*(m) = O(m^2)$. This is because drawing the contour of m triangles the plane partitions into $O(m^2)$ regions. In general, if $\pi_{\mathcal{A}}(m) \leq Cm^d$ for constants C, d , and $1 \leq m \leq n$, Welzl's theorem guarantees the existence of a spanning tree with crossing number $O(n^{1-1/d} \log n)$ for (P, \mathcal{P}) .

Returning to the case of halfspaces in \mathbb{R}^d , we understand that Welzl's theorem implies the existence of a spanning tree with crossing number $O(n^{1-1/d} \log n)$ for any n -point set in \mathbb{R}^d . Later on, Chazelle and Welzl [1989] improved the bound on the crossing number in this geometric setting to $O(n^{1-1/d})$, which is asymptotically optimal.^{5, 6} In this way they proved an almost optimal upper bound for query complexity for the simplex range-searching problem in the arithmetic model

⁴A *primal shatter function* also exists, but it is not significant for spanning trees with a low crossing number

⁵A recent result of Haussler [1991] implies that the bound without $\log n$ also holds in Welzl's abstract theorem.

⁶The problem of a spanning tree with a low crossing number in the plane belongs to the relatively few in computational geometry where someone worked on a more exact determination of the constant of proportionality. Welzl [1992] showed that the optimal crossing number lies between \sqrt{n} and $(2.31 + o(1))\sqrt{n}$ in the worst case.

(which we have discussed in connection with the lower bounds). They even obtained a certain characterization of set systems admitting a sublinear query complexity in the arithmetic model. However, their method does not automatically give an efficient algorithm in the usual sense. We have already mentioned where the difficulty lies: we know that the boundary of the query simplex or halfspace crosses only a few edges of the spanning tree, but it is not clear how to find these edges efficiently. In general, this problem seems almost as difficult as the original simplex range-searching problem. Chazelle and Welzl managed to solve this additional complication in dimensions 2 and 3, thus obtaining the first almost optimal algorithms. Before leaving spanning trees with low crossing numbers, let us mention that many other algorithmic applications were found, e.g., Agarwal [1990] and Edelsbrunner et al. [1989], and even purely combinatorial applications (see Matoušek et al. [1993] or the contribution of Pach in Goodman et al. [1991]).

Further Developments

The paper by Matoušek and Welzl [1992] stands somewhat aside from this "mainstream." It describes an essentially different method for answering halfplane range queries, based on a combinatorial lemma of Erdős and Szekeres, and yields $O(\sqrt{n} \log n)$ query time with $O(n \log n)$ space. This method can also be used for triangles, but only if the point weights can be subtracted. There does not seem to be much hope for generalizing this method into higher dimensions. On the other hand, it is an easily implementable algorithm with small constants of proportionality, which cannot be said about most of the other algorithms.

Chazelle et al. [1992] discovered a simplex range-searching algorithm for an arbitrary dimension, with $n^{1+\epsilon}$ space and $O(n^{1-1/d+\epsilon})$ query time, for an arbitrarily small positive constant ϵ . They thus approached the lower bound up to a factor of n^ϵ . (For this reason they call the algorithm quasioptimal.) Their method

retrns to recursively constructed partition trees. For a single point set, however, they construct not one but several partition schemes at the same time, in such a way that for any halfspace R at least one of these schemes is sufficiently efficient. The construction then recurses for the point set in each of the regions of each partition scheme, which leads to a larger memory requirement, $n^{1+\epsilon}$. As indicated above, if one partition scheme is good for all halfspaces, then it is also good for simplices, but if we must use several schemes it may happen that none of them is good for simplices. For this reason Chazelle et al. apply so-called multilevel data structures to handle the simplex queries (see Section 6.3).

Simplicial Partitions

The paper by Matoušek [1992b] returns to a single partition scheme whose parameters are asymptotically optimal for all halfspaces, and thus also for simplices. The partition is constructed by a suitable generalization of Welzl's construction of spanning trees with a low crossing number [Welzl 1988]. Let us describe these partition schemes. For simplicity, we only formulate the following definition for a point set in a general position.

Let P be an n -point set in \mathbb{R}^d in a general position. A *simplicial partition* for P is a collection

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

where the P_i are disjoint subsets of P (called the *classes*) forming a partition of P , and each Δ_i is a d -dimensional simplex containing the set P_i .

Let us remark that the simplices need not be disjoint, and the simplex Δ_i may also contain other points of P than those of P_i (see Figure 3). Although it is not clear why this should make the construction of a partition scheme any easier, no partition schemes of a comparable efficiency to Theorem 4.1, but with disjoint regions, are known.

If h is a hyperplane and Δ a simplex, we say that h *crosses* Δ if it intersects it. (For simplices of lower dimension than d




Figure 3. A simplicial partition (points of different classes are marked by different symbols).

which appear when we deal with sets in degenerate positions, the definition of crossing is somewhat more complicated.) Further, we define the *crossing number* of the simplicial partition Π with respect to h as the number of simplices of Π crossed by h , and the *crossing number* of Π is the maximum of the crossing numbers with respect to all hyperplanes. The main theorem of Matoušek [1992b] is as follows:

THEOREM 4.1. *Let P be an n -point set in \mathbb{R}^d ($d \geq 2$), r be a parameter, $1 < r \leq n/2$. Then there exists a simplicial partition for P satisfying $n/r \leq |P_i| \leq 2n/r$ for every class P_i (thus with $O(r)$ classes) and with crossing number $\kappa = O(r^{1-1/d})$.*

The crossing number in this result is asymptotically optimal. Matoušek [1992b] also gives an algorithm for constructing such a simplicial partition, with $O(n \log r)$ running time for sufficiently small r ($r \leq n^\beta$ for a certain small constant $\beta = \beta(d) > 0$).

Construction of Simplicial Partitions

The construction is based on an application of *cuttings*. (See box 4.)⁷ The first

⁷The use of cuttings replaces an argument from the otherwise similar Welzl [1988] construction of spanning trees with a low crossing number. The application of cuttings makes more use of geometric properties. This is not accidental, as Alon et al. [1987] showed that for certain set systems there exists no efficient partition scheme with a constant-bounded number of regions, although the dual shatter function is small and a spanning tree with a low crossing number thus exists.

A *cutting* is a finite set of closed simplices (here a simplex means an intersection of $d + 1$ halfspaces; hence unbounded “simplices” are also allowed) with disjoint interiors covering \mathbb{R}^d . Let H be a set of n hyperplanes in \mathbb{R}^d and r a parameter, $1 < r \leq n$. A cutting Ξ is called a $(1/r)$ -cutting for H provided that the interior of no simplex of Ξ is intersected by more than n/r hyperplanes of H . Sometimes it is useful to consider a weighted version also, where a nonnegative weight function $w: H \rightarrow \mathbb{R}$ is given, and the sum of weights of hyperplanes intersecting any simplex of a $(1/r)$ -cutting should not exceed $1/r$ of the total weight of all hyperplanes of H .

The notion of a cutting is a basis of many geometric algorithms of the “divide-and-conquer” type dealing with hyperplanes. A computational problem concerning a set H of hyperplanes is partitioned into subproblems defined by the simplices of a $(1/r)$ -cutting, each involving $r \times$ fewer hyperplanes than the original problem. This dividing strategy was developed in the work of Clarkson, Haussler, and Welzl, and others. Cuttings are considered explicitly in Chazelle and Friedman [1990] and Matoušek [1990]; the name and the current definition were given in Matoušek [1991c].

For applications, it is important that the number of simplices in a cutting is as small as possible. Chazelle and Friedman [1990] showed that for every H and r there exists a $(1/r)$ -cutting consisting of $O(r^d)$ simplices. This result is asymptotically optimal, and applies to the weighted version also. Efficient computation of cuttings is considered in Chazelle and Friedman [1990], Matoušek [1991a; 1991c], and Chazelle [1993a].

step is the choice of a suitable finite “test set” H of hyperplanes, such that whenever the crossing number of a simplicial partition Π with respect to every $h \in H$ is bounded by some number κ , then the crossing number of Π with respect to any hyperplane is $O(\kappa + r^{1-1/d})$. The test set H has size $O(r)$, and it is constructed using cuttings in the dual space.

With such a test set H at our disposal, the relation of cuttings to a small crossing number of a simplicial partition is as follows: if Ξ is a $(1/t)$ -cutting for H , then the total number of incidences among the simplices of Ξ and the hyperplanes of H is at most $|\Xi||H|/t$ (since any simplex of Ξ has at most $|H|/t$ incidences), and hence the average number of simplices of Ξ crossed by a hyperplane of H is $|\Xi|/t$. The total number of simplices in Ξ is $r = O(t^d)$, and hence the crossing number with respect to an average hyperplane of H is $O(t^{d-1}) = O(r^{1-1/d})$. The cutting itself thus gives a simplicial partition whose crossing number with respect to an average hyperplane of H is asymptotically optimal. There are two shortcomings, however.

First, the points of P need not be equally distributed among the simplices of Ξ , and second, although the average crossing number is optimal, there might be some hyperplanes with bad crossing numbers.

These are fixed by an incremental construction of Π using the so-called *reweighting strategy*, imitating Welzl’s construction. The simplicial partition Π is constructed step by step, one pair (P_i, Δ_i) at a time. Suppose that (P_1, Δ_1) through (P_i, Δ_i) have already been constructed, with $|P_j| = \lfloor n/r \rfloor$, $j = 1, \dots, i$. The set $\tilde{P}_i = P \setminus (P_1 \cup \dots \cup P_i)$ of points not yet covered has $n_i = n - i \lfloor n/r \rfloor$ points. For a hyperplane $h \in H$, let $\kappa_i(h)$ be the number of simplices among $\Delta_1, \dots, \Delta_i$ crossed by h , and assign a weight $w_i(h) = 2^{\kappa_i(h)}$ to every $h \in H$. We let Ξ_i be a $(1/t)_i$ -cutting for H, w_i ; that is, the sum of the w_i weights of the hyperplanes intersecting any simplex Δ of Ξ_i is at most $(1/t)$ of the total w_i weight of all hyperplanes. We choose t as large as possible, but in such a way that the total number of simplices of Ξ_i does not exceed $r(n_i/n)$, and therefore we can find a simplex of Ξ_i containing at least n/r points from \tilde{P}_i . Such a simplex then becomes Δ_{i+1} , and the set P_{i+1} is chosen as some $\lfloor n/r \rfloor$ points of \tilde{P}_i among those contained in Δ_{i+1} .

This finishes the description of the construction, and we now have to prove that the resulting simplicial partition has the required crossing number. One bounds the increment of the total weight of all hyperplanes of H caused by adding one new simplex Δ_{i+1} to the simplicial parti-

tion. This yields an estimate on the total weight of all hyperplanes after the last stage, and from this bound it follows that no hyperplane can give too big a crossing number; we refer to Matoušek [1992b] for a detailed calculation.

By a recursive application of the simplicial partition from Theorem 4.1, a partition tree is created in a standard way. For a given set P we find a simplicial partition Π (with a suitable choice of the parameter r), and we store its simplices as well as the total weights of points in the classes of Π in the root of the partition tree. Each class of the simplicial partition corresponds to one subtree of the root, where the construction is used recursively for the points in the corresponding classes. When answering a query with a halfspace R , we process the simplices lying completely inside R or completely outside R directly in the current node, and for the simplices crossing the boundary of R we descend recursively into the appropriate subtrees. For the query time, we thus obtain the recurrence

$$T(n) \leq f(r) + \kappa T(2n/r), \quad (7)$$

where $\kappa = O(r^{1-1/d})$ is the crossing number of the simplicial partitions, and $f(r)$ is the cost of the auxiliary computations in one node of the partition tree, i.e., finding the crossing simplices and computing the total weight of points in simplices lying completely inside the query halfspace R . If we perform a simple inspection of all the simplices of the simplicial partition and test their position with respect to R , we have $f(r) = O(r)$.

We must now choose a suitable value of the parameter r . If we take a large constant for r (as was done in previous partition tree constructions), Eq. (7) yields $T(n) = O(n^{1-1/d+\epsilon})$, where ϵ tends to 0 as $r \rightarrow \infty$. We can, however, also choose r as some function of n (more exactly, of the number of points in the current node of the partition tree), e.g., a small power of n . We are even interested in making r as large as possible: A larger r means more branching, and thus a

shallower partition tree, and since we lose a constant factor in query-answering efficiency at every level of the partition tree, the shallower the tree the better. On the other hand, we must not overdo it, since otherwise the term $f(r)$ in recurrence (7) becomes an obstacle. This means that if the number of simplices of the simplicial partition were, say, comparable to n , then the auxiliary computations (finding the crossing simplices, summing the weights) would take more time than we can afford. By a suitable compromise, where r is a sufficiently small power of n , we obtain query time $O(n^{1-1/d} \log^c n)$, where c is some constant (which we should not try to estimate). In this situation the preprocessing time is $O(n \log n)$.

In order to use a shallower partition tree without $f(r)$ getting too large, we can start building various auxiliary data structures for the above-mentioned computations in each node of the partition tree. This method is pursued in Matoušek [1992b], and with much more complicated data structures one gets an asymptotically better query time $O(n^{1-1/d} (\log \log n)^c)$.

Matoušek [1993c] succeeded in reducing the query complexity for simplex range searching with linear space to $O(n^{1-1/d})$. The method is technically somewhat complicated and is based on ideas similar to those mentioned above plus an application of so-called hierarchical cuttings from Chazelle [1993a]. This result is optimal in dimension 2 and most likely also in higher dimensions; see Section 3. However, the preprocessing time is $O(n^{1+\epsilon})$, while an ideal one would be $O(n \log n)$. Thus, there is still room for improvement, and there is also the challenge of finding a simpler optimal algorithm.

5. ALGORITHMS WITH LOGARITHMIC QUERY TIME

Halfspace Range Searching

A logarithmic query time is essentially achieved by precomputing all possible answers and storing them in a suitable data structure. It is more natural to con-

sider an equivalent dual version of the problem. (See box 2 for the definition of duality.)

Problem 5.1. Let H be a set of n hyperplanes in \mathbb{R}^d equipped with weights. For a query point q , find the sum of weights of the hyperplanes lying above q .⁸

The answer obviously remains constant for points q lying in the same cell of the arrangement of H . (See box 3.) We can thus construct the arrangement, compute and store the weight for each cell, and then it suffices to detect the cell containing a query point and return the precomputed answer. The preprocessing can be done in $O(n^d)$ time, e.g., see Edelsbrunner [1987] and Chazelle et al. [1992]. For dimensions 2 and 3, efficient methods are known for fast point location in an arbitrary subdivision of the space into convex cells. For higher dimensions, special methods have been developed for point location in hyperplane arrangements. (See box 4.) This gives a halfspace range-searching method with $O(n^d)$ space and preprocessing time and $O(\log n)$ query time. The space and preprocessing time can even be reduced (approximately by a factor of $\log^d n$) with the same query time, as was shown in Matoušek [1993c].

Here we explain a simple algorithm formulated directly for Problem 5.1; point location algorithms for hyperplane arrangements are similar. It is a nice application of the notion of cuttings. (See box 1.)

The data structure is, as usual, a tree, and each node corresponds to some subset of hyperplanes of H . The construction starts in the root, which corresponds to the whole H . We choose a large constant r , and we construct a $(1/r)$ -cutting $\Xi =$

$\{\Delta_1, \dots, \Delta_s\}$ for H . By theorems on the existence of cuttings, we can assume $s \leq Cr^d$, where C is an absolute constant independent of r . For every simplex $\Delta_i \in \Xi$, let H_i be the set of hyperplanes of H intersecting its interior, and let w_i be the total weight of hyperplanes lying completely above the interior of Δ_i . The cutting Ξ and the weights w_i are stored in the root, and for every $i = 1, 2, \dots, s$, we build a subtree corresponding to H_i by a recursive application of the same procedure. The recursion terminates as soon as we reach sets of hyperplanes smaller than a suitably chosen constant. Then, instead of a further recursion, we store the corresponding set of hyperplanes in a leaf of the constructed tree.

The total weight of hyperplanes above a query point q is determined as follows. For the cutting stored in the root, we find the simplex Δ_i containing q (in constant time, since r was a constant). This gives us the total weight w_i of the hyperplanes from $H \setminus H_i$ lying above q , and it remains to determine the total weight of hyperplanes of H_i above q . This is computed by a recursive application of the same method on the subtree corresponding to H_i . In this way we reach a leaf of the tree after $O(\log n)$ steps, and we process the several hyperplanes stored there directly.

Let us look at the space requirements of this data structure. The space $S(n)$ needed for n hyperplanes can be bounded using the recurrence relation

$$S(n) \leq Cr^d + Cr^d S(n/r).$$

The resulting bound is $S(n) = O(n^{d+\epsilon})$, where $\epsilon > 0$ can be made arbitrarily small by choosing r large enough.

The hierarchy described above, of progressively refining cuttings, is somewhat wasteful, however. On each level we lose one additional constant factor, which is reflected in the total space requirement. This effect can be restricted by making r larger, but then the search in the cutting in a single node becomes expensive. Chazelle [1993a] succeeded in overcoming this. His method produces a hierar-

⁸This reformulation follows from our special choice of the duality transform. In general, one can consider the problem of finding the total weight of the hyperplanes separating a query point from a fixed point o . Such problems can be converted to each other by projective transforms.

chy of progressively refining cuttings, where the finest cutting on the last level has a total size of only $O(n^d)$, instead of $n^{d+\epsilon}$. Chazelle's method is also explained in Matoušek [1993b] in a somewhat simplified form.

Simplex Range Searching

Since about n^6 subsets of an n -point set in the plane can be defined by a triangle, storing all possible answers for the triangle range-searching problem requires at least about n^6 space. Such an approach was used by Edelsbrunner et al. [1982], and their solution requires space of the order n^7 . A significant reduction of the space requirement, to $O(n^{2+\epsilon})$, was achieved by Cole and Yap [1985]; their method is based on an idea resembling multilevel data structures (see Section 6.3). They also give an $O(n^{10})$ space data structure for three-dimensional simplex range searching. Another paper, on considering the planar case, is by Paterson and Yao [1986]. A general method for simplex range searching with storage approximately n^d for arbitrary point weights and arbitrary dimension was found relatively recently by Chazelle et al. [1992] (see Section 6.3). Space and preprocessing time are $O(n^{d+\epsilon})$ and query time is $O(\log^{d+1} n)$. The space and preprocessing times were improved in Matoušek [1993c], using Chazelle's hierarchy of cuttings [Chazelle 1993a].

6. EXTENSIONS AND EXAMPLES OF APPLICATIONS

6.1 Halfspace Range Reporting and Other Special Situations

The halfspace range-searching problem with only emptiness queries or reporting queries turned out to have more efficient solutions than the general simplex range-searching problem. It seems that the actual query complexity for halfspace range reporting with space at most m might be approximately

$$\frac{n}{m^{\lfloor d/2 \rfloor}} + k, \quad (8)$$

where k denotes the number of points in the query halfspace. In particular, in dimensions 2 and 3, algorithms exist with almost linear storage and $O(\log n + k)$ query time; see Chazelle et al. [1985], Chazelle and Preparata [1986], and Aggarwal et al. [1990]. For a higher dimension, Clarkson [1988a] found an algorithm with space and preprocessing $O(n^{\lfloor d/2 \rfloor + \epsilon})$ and query time also $O(\log n + k)$. This result was complemented in Matoušek [1992a] by an algorithm with $O(n \log \log n)$ space, $O(n \log n)$ preprocessing time and $O(n^{1-1/\lfloor d/2 \rfloor} \log^c n + k)$ query time. (The method is quite similar to the one for simplex range searching in Matoušek [1992b], discussed at the end of Section 4.) By combining these two methods, the complexity given by Eq. (8) can almost be achieved in the whole spectrum of memory requirements. Further small improvements were described by Schwarzkopf [1992]. No lower bounds are known.

Very roughly speaking, we can say that the exponent d in Eq. (1) expressing the complexity of simplex range searching originates in the fact that the combinatorial complexity of an arrangement of n hyperplanes in \mathbb{R}^d is of the order n^d . Similarly, the exponent $\lfloor d/2 \rfloor$ in Eq. (8) is closely related to the worst-case complexity of a polytope defined as an intersection of n halfspaces in \mathbb{R}^d , which is of the order $n^{\lfloor d/2 \rfloor}$, e.g., see Edelsbrunner [1987]. When we must quickly decide the emptiness of a halfspace or—in the dual form—quickly test whether a given point lies above all hyperplanes of a given set H , we deal essentially with a problem of point location in a convex polytope, namely in the upper unbounded cell of the arrangement of H .

In the dual version of halfspace range reporting, we essentially restrict our attention to a single cell of the arrangement of the given hyperplanes, and as we saw, this yields a considerable improvement. Another such special situation is when our point set lies on a fixed lower-dimensional algebraic variety of bounded degree, or if all hyperplanes bounding the query ranges are tangent to such a

variety, then the range-searching results can sometimes also be improved. (The restricted point set case helps the linear-space case; the restricted ranges case helps the large space case, e.g., see Agarwal and Matoušek [1994] for a discussion.) These improvements use a combinatorial result, a zone theorem of Aronov et al. [1993]. The situation with points on a surface is by no means rare—it arises, for example, when dealing with lines in 3-dimensional space; see Chazelle et al. [1989].

6.2 Dynamization

Until now we have considered static range-searching problems, where the point set is given once and for all. In many applications we need to insert new points or delete old ones from time to time, and it is advantageous if we can just modify the data structure instead of rebuilding it from scratch after every change. This problem is called the *dynamization* of a data structure. For a dynamic data structure storing n points, we probably cannot expect a better time for one modification than $1/n$ fraction of the time needed for building the whole (static) structure anew. For the simplex and halfspace range-searching problems, dynamic structures with this efficiency (up to small factors, typically of order n^ϵ) are known; see Argrawal and Matoušek [1991], Argrawal and Sharir [1993], Matoušek [1992b], and Schipper and Overmars [1991]. For dynamization in general, see Bentley and Saxe [1980], Mehlhorn [1984], and Overmars [1983]. It would still be interesting to get rid of the n^ϵ extra factors for, say, the dynamic halfspace empty data structures.

Dynamizing the simplex range-searching data structures of Chazelle et al. [1992] or Matoušek [1992b] is quite straightforward, and it is equally easy to dynamize the linear-space halfspace reporting data structure of Matoušek [1992a]; see Agarwal and Matoušek [1993]. A dynamic version of Clarkson's halfspace reporting data structure is more complicated. The source of difficul-

ties here is the fact that the maximal asymptotic complexity of the convex hull of n points in an odd dimension d and the preceding even dimension $d - 1$ are the same. The change of the convex hull caused by adding or deleting a single point is proportional to the complexity of a certain convex hull in dimension $d - 1$. For instance, consider the convex hull of n points in \mathbb{R}^3 . This is a convex polyhedron with $O(n)$ vertices, edges, and facets, and some vertices can have degrees close to n . If we keep inserting and deleting such vertices, it seems impossible to maintain an explicit representation of the convex hull (e.g., as a planar graph) in time substantially smaller than n for one operation.

For these reasons, among others, several authors investigated dynamic data structures under the assumption that the update sequence is random in a suitably defined sense, and obtained very good update times for this case; see Mulmuley [1991c; 1991d] and Schwarzkopf [1991]. Agarwal and Matoušek [1991] found a dynamic algorithm, which is also efficient in the worst case, for an arbitrary update sequence.⁹ This algorithm does not maintain an explicit representation of the convex hull; rather it works with certain implicit representations. Nevertheless, this data structure in connection with other techniques (see Section 6.5) supports fast answering of various queries concerning the convex hull, such as deciding if a query point lies inside or outside the current convex hull, computing a tangent hyperplane to the convex hull containing a given line, and so on.

6.3 Multilevel Data Structures

First we explain the idea of multilevel data structures with an example and then introduce some abstract notions for their descriptions.

⁹In the amortized sense, i.e., an individual update time can sometimes be large, but the average over a sequence of n updates is small.

Let $S = \{s_1, \dots, s_n\}$ be a set of segments in the plane. We want to construct a data structure that quickly computes the number of segments of S intersected by a query line h . We permit roughly linear space for the data structure. Let a_i, b_i be the endpoints of the segment s_i . A denotes the set of all a_i and B the set of all b_i . We consider computing the number of s_i such that a_i is above h and b_i below h . The opposite case is solved symmetrically, and the case of vertical query lines can be treated separately.

Let us consider some partition tree for the set A ; for certainty, let it be the partition tree based on simplicial partitions from Theorem 4.1, with r being a large constant. Using such a partition tree we can determine, in roughly \sqrt{n} time, the number of points of A in the halfplane R above h . This does not solve our problem yet, but we look more closely at how the answer is obtained from the partition tree. The weight of every point from $A \cap R$ is accounted for in some of the visited nodes of the tree. In each such node we find the simplices of the corresponding simplicial partition lying completely inside R , and the weights of their respective classes are accounted for as wholes.

For each node of the partition tree, let us call the classes of the simplicial partition stored there the *canonical sets*. We see that the partition tree provides a partition of the set $A \cap R$ into roughly \sqrt{n} canonical sets of various sizes. The total number of canonical sets in the partition tree is $O(n)$, and the sum of their sizes is easily estimated to be $O(n \log n)$.

For our problem with segments, we augment the partition tree for the set A as follows. For every canonical set $M \subset A$ we create a partition tree for the set $M' = \{b_i; a_i \in M\}$, and we store it with the corresponding node of the partition tree for A . (The partition tree for A is called *primary* or *first-level*; the trees for the sets M' are called *secondary* or *second-level*.)

How does one apply the resulting data structure to the segment-counting problem? First we express the set $A \cap R$ as a

disjoint union of certain canonical subsets M_1, \dots, M_m , and then for each such M_i we use the appropriate secondary partition tree to count the points of M_i' lying below the line h . Adding these counts together over all canonical sets M_i , we obtain the desired number of segments s_i with a_i above h and b_i below h .

On first sight it seems that the described two-level data structure should be much less efficient than the original partition trees used for its construction. By a simple calculation we can convince ourselves that this is not the case, and that the required space is $O(n \log n)$ only (this is because of the total size of the canonical subsets), and the query time remains still close to \sqrt{n} . This is because there are only few large canonical sets in the decomposition of $A \cap R$, and the computation on the second level is fast for small canonical sets.

The principle used in the above example is quite universal. Usually, it is applicable whenever the query is a conjunction of several conditions (or, geometrically, an intersection of several regions), and for each condition (region) we already have a suitable efficient data structure.

The idea of multilevel data structures appears in Bentley's data structures for orthogonal range searching [Bentley 1980]. In the context of partition trees such a construction was introduced in Dobkin and Edelsbrunner [1987]. Recently it has been used quite frequently, e.g., see Overmars et al. [1990], Chazelle et al. [1992], Matoušek [1992b], and Agarwal and Sharir [1993], and many others. The descriptions of such data structures often look complicated, especially if there are more levels. Matoušek [1993c] proposed an abstract framework for description and analysis of such data structures.

Let us restate the principle of multilevel data structures in a somewhat abstract setting and give two more examples. First we introduce the notion of a decomposition scheme. Consider a range-searching problem in an abstract form, such as we mentioned in Section 3.

It is given by a set system (P, \mathcal{P}) , usually defined as in Eq. (2). From an abstract point of view, in most range-searching algorithms one proceeds as follows. Another system $\mathcal{E}(P)$ of subsets of P , the so-called *canonical subsets*, is defined, and a rule is given how to express each range $R \in \mathcal{P}$ as a disjoint union of canonical sets from $\mathcal{E}(P)$. Let us call this pair, the set system $\mathcal{E}(P)$ plus the rule how to decompose sets of \mathcal{P} , a *decomposition scheme* for (P, \mathcal{P}) .

Such a decomposition scheme can be turned into a range-searching data structure (provided that the decompositions can be found efficiently) by simply storing the total weight of points for each canonical subset; this is how most of the range-searching data structures work, although the canonical subsets are not mentioned explicitly.

For partition trees as described in Section 4, canonical subsets are the point sets lying in regions of the partition schemes at the nodes of the partition tree. For the one-dimensional data structure for range searching with intervals in Example 2.2, the canonical subsets are just canonical intervals, and each interval can be partitioned into $O(\log n)$ canonical intervals. In the data structure for halfspace range searching with logarithmic query time explained in Section 5, each node of the tree defines canonical subsets corresponding to simplices of the cutting stored in it. If the node corresponds to a subset G of hyperplanes, and Δ_i is one of the simplices in its cutting, then the corresponding canonical subset is formed by all the hyperplanes of G lying completely above Δ_i . Any range, which is the set of hyperplanes lying above a query point, can be expressed as a disjoint union of $O(\log n)$ canonical sets.

For our subsequent development it is necessary that a decomposition scheme operate not only on the system (P, \mathcal{P}) itself, but also on systems induced by subsets of P . This means that for a subset $P' \subset P$, we also have a decomposition scheme for $(P', \{P' \cap R; R \in \mathcal{P}\})$. This is usually trivially true for geometric problems, since if we can build a decomposi-

tion scheme for a set, we can also build one for its subset. Such a decomposition scheme will be called a *hereditary* one.

Now let P be a set of basic objects (not necessarily points; in the initial example of this section these would be the segments). Let $\mathcal{P}_1, \mathcal{P}_2$ be two set systems on P , and suppose that we have a decomposition scheme \mathcal{D}_1 for (P, \mathcal{P}_1) and a hereditary decomposition scheme \mathcal{D}_2 for (P, \mathcal{P}_2) . In the example with segments, \mathcal{P}_1 would be all subsets of the form $\{s_i \in P; a_i \in H\}$ for some halfplane H , and similarly $\mathcal{P}_2 = \{s_i \in P; b_i \in H\}$ for some halfplane H . The decomposition scheme for (P, \mathcal{P}_1) is given by building the partition tree on the set A of the a_i -endpoints, and the decomposition scheme for (P, \mathcal{P}_2) by a partition tree on the b_i -endpoints. (Thus formally we work with sets of segments, although the partition trees actually deal with their endpoints.)

Our goal is to derive a decomposition scheme for a more complicated system (P, \mathcal{P}) , where \mathcal{P} consists of all sets of the form $R_1 \cap R_2$ with $R_1 \in \mathcal{P}_1$ and $R_2 \in \mathcal{P}_2$. In the example with segments, the ranges we are interested in are indeed of this form, namely the sets of segments whose a_i -endpoints lie in an upper halfplane H_1 and b_i -endpoints in a lower halfplane H_2 ; in fact, H_1 and H_2 are complementary halfplanes, but we take no advantage of this. We define what is called the *composition* of the decomposition schemes \mathcal{D}_1 and \mathcal{D}_2 , denoted by $\mathcal{D} = \mathcal{D}_1 \circ \mathcal{D}_2$. To decompose a set $R = R_1 \cap R_2 \in \mathcal{P}$, we first decompose, using \mathcal{D}_1 , the set $R_1 \in \mathcal{P}_1$ into canonical subsets $C_1, C_2, \dots, C_m \in \mathcal{E}_1(P)$. Then for each C_i , we consider the decomposition scheme \mathcal{D}_2 operating on the subsystem of \mathcal{P}_2 induced by C_i , and decompose the set $C_i \cap R_2$ into canonical sets $C_{i1}, C_{i2}, \dots, C_{ik_i} \in \mathcal{E}_2(C_i)$. The union of these collections for all $i = 1, 2, \dots, m$ gives us a decomposition of $R_1 \cap R_2$. The canonical sets \mathcal{E} in the resulting decomposition scheme \mathcal{D} will thus be all canonical sets from $\mathcal{E}_2(C)$ for some $C \in \mathcal{E}_1(P)$.

What we did in the example with segments can thus be rephrased as first

composing the two decomposition schemes and then turning the resulting decomposition scheme into a range-searching data structure by precomputing the weights of all canonical sets. Knowing the parameters of both the decomposition schemes, it is a routine calculation to derive the parameters of the composed scheme. The total number of canonical subsets determines the space requirements of the data structure, and the maximal number of canonical sets appearing in a decomposition of a range is related to query time.

As one basic example, let us consider orthogonal range searching. Let P be an n -point set in \mathbb{R}^d . We let \mathcal{P}_i be the set system defined on P by intervals on the x_i -axis. For each (P, \mathcal{P}_i) we have the decomposition scheme \mathcal{D}_i with canonical sets being the canonical intervals along the x_i -axis (see Example 2.2). A range of the form $R_1 \cap R_2 \cap \dots \cap R_d$ with $R_i \in \mathcal{P}_i$ corresponds to a subset of P lying in an axis-parallel box. The d -wise composition $\mathcal{D} = \mathcal{D}_1 \circ \mathcal{D}_2 \circ \dots \circ \mathcal{D}_d$ is thus a decomposition scheme for the set system defined by axis-parallel boxes. By turning this decomposition scheme into a range-searching data structure, we recover an abstract version of the *range trees* of Bentley [1980].

Another important example is the application of multilevel data structures for the already-mentioned simplex range searching with polylogarithmic query time from Chazelle et al. [1992]. The halfspace range-searching method described in Section 5 yields a decomposition scheme \mathcal{D}_0 which partitions the point set in a query halfspace into $O(\log n)$ canonical subsets (we stated this above in the dual form). Since a simplex is an intersection of $d + 1$ halfspaces, we can obtain a simplex decomposition scheme by a $(d + 1)$ -wise composition of \mathcal{D} with itself, and this gives the desired simplex range-searching data structure. Easy calculations show that the resulting data structure has query time $O(\log^{d+1} n)$ and occupies memory $O(n^{d+\epsilon})$.

Let us conclude with few remarks. Asymptotically, one usually loses a factor of n^ϵ in space and query time with each level of a multilevel data structure (compared to the efficiency of the data structures used for the levels). In Matoušek [1993c], decomposition schemes are described that allow one to build multilevel data structures while losing a polylogarithmic factor per level only. On the other hand, it seems that a practical efficiency of a data structure will be lost quite quickly with an increasing number of levels. (Again, this is not substantiated by any implementation experience.) Thus, it makes a big difference if we use a simplex decomposition scheme directly or if we create it by composing halfspace decomposition schemes $d + 1$ times.

We should also point out that while most of the simplex range-searching algorithms provide decomposition schemes in the sense explained above, the halfspace emptiness and halfspace reporting algorithms discussed in Section 6.1 do not provide a halfspace decomposition scheme. (In some sense, they can only decompose very “shallow” halfspaces, i.e., those that contain a few points of P only.) Thus, if we want to build a data structure for testing the emptiness of complicated ranges, where one of the defining conditions is a halfspace, we may only use the halfspace emptiness data structure in the lowest (last) level of a multilevel data structure.

6.4 Searching with More General Ranges

In the previous sections, we considered range searching with ranges bounded by hyperplanes. Many applications lead naturally to searching in ranges with nonlinear, curved boundaries. How should general query ranges be considered? It seems that nontrivial results can be expected mainly for ranges determined by a small (constant) number of real parameters. The most important such case is subsets of \mathbb{R}^d defined by a conjunction of at most k polynomial inequalities of maximum degree D , where d , k , and D

are constants. We call such sets *elementary cells*, for short.¹⁰

Range-searching problems with various kinds of elementary cells include many specific problems motivated by applications. Some of the results mentioned below can be formulated for still somewhat more general ranges, but writing out all the necessary assumptions would become too lengthy.

Perhaps the simplest among nonlinear ranges are circular disks in the plane and balls in higher dimensions. The corresponding range-searching problem arises when we are interested in points lying in at most a given distance from a given point. Notice that balls do not appear in this second, “application-oriented,” formulation directly. This is quite typical, since range-searching problems with nonlinear ranges are usually obtained by a suitable reformulation of the original specification of the problem. For instance, if we are asking for points whose distance from a given segment does not exceed a given number, an equivalent formulation is range searching with ranges of the shape of a racecourse. If we want to detect points that see a segment s_1 under a larger angle than a segment s_2 , we get a range-reporting query with a rather complicated range defined by inequalities of degree 6, and so on.

Ball range searching and related problems, usually referred to as *proximity problems*, have a rich literature, e.g., see Chazelle et al. [1986], Aggarwal et al. [1990], Chazelle and Welzl 1989], Clarkson [1994], and others. Range searching with other ranges was paid less attention to. An important earlier work considering very general geometric range-searching problems is by Yao and Yao [1985]; it contains several important ideas which were further developed and found many applications later on. A recent work considering the problem of

range searching in elementary cells is by Agarwal and Matoušek [1994].

To some extent, methods developed for simplex and halfspace range searching can also be applied for searching with elementary cells. We now outline the results. We concentrate on ranges defined by a *single* polynomial inequality. Ranges described by a conjunction of several inequalities can be handled similarly, sometimes by applying multilevel data structures, where each level processes one of the inequalities.

A disk $C = C(a_1, a_2, a_3) \subset \mathbb{R}^2$ with center (a_1, a_2) and radius a_3 is described by the inequality $(x_1 - a_1)^2 + (x_2 - a_2)^2 \leq a_3^2$. More generally, we consider ranges of the form

$$R_f(a) = \{x \in \mathbb{R}^d; f(x, a) \geq 0\},$$

where f is a fixed polynomial in $d + p$ variables $x_1, \dots, x_d, a_1, \dots, a_p$. The polynomial f specifies the type of the considered ranges (disks, conics, cylinders, ...), and a is a p -dimensional parameter vector determining the particular range of the given type (a particular disk, etc.). The set of admissible ranges for such a problem is $\mathcal{R}_f = \{R_f(a); a \in \mathbb{R}^p\}$.

One possible way of solving the range-searching problems with ranges from \mathcal{R}_f is a direct transformation to halfspace range searching in space of higher dimension, the so-called *linearization* of the problem. Let us consider the example with disks, where

$$\begin{aligned} f(x_1, x_2, a_1, a_2, a_3) &= a_3^2 - (x_1 - a_1)^2 - (x_2 - a_2)^2 \\ &= a_3^2 - a_1^2 - a_2^2 + 2a_1x_1 \\ &\quad + 2a_2x_2 - x_1^2 - x_2^2. \end{aligned}$$

We define a mapping $\varphi: \mathbb{R}^2 \rightarrow \mathbb{R}^4$ by $\varphi(x_1, x_2) = (x_1, x_2, x_1^2, x_2^2)$. A key property of this mapping is the following.

For every disk $C = C(a_1, a_2, a_3)$ there exists a halfspace $H = H(a_1, a_2, a_3)$ in \mathbb{R}^4 such that $C = \varphi^{-1}(H)$ or, in other words, the points of the plane mapped by φ into the halfspace H are exactly those of the disk C . It is easy to see that a

¹⁰ In the literature, the term *Tarski cells* is used for a similar but somewhat more general notion.

suitable halfspace is $H = \{(t_1, t_2, t_3, t_4) \in \mathbb{R}^4; 2a_1t_1 + 2a_2t_2 - t_3 - t_4 + a_3^2 - a_1^2 - a_2^2 \geq 0\}$. A mapping φ with the above-described property is called a *linearization*. A linearization allows us to transform the disk range-searching problem in the plane with a point set P to the half-space range-searching problem in \mathbb{R}^4 with the point set $\varphi(P)$.

The linearization φ was obtained by introducing one new coordinate t_j for each monomial in the variables x_1, x_2 occurring in the polynomial f . A problem with an arbitrary polynomial f can be linearized in the same way—the linearization arising in this way is known in algebraic geometry as the *Veronese map*; see Harris [1992]—and thus we obtain nontrivial range-searching algorithms for ranges from \mathcal{R}_f . This important observation was made in Yao and Yao [1985].

The example with disks, which we used to demonstrate the general method, also shows that this method is sometimes wasteful in the dimension of the image space. It is well known that the disks in the plane can be linearized in the above sense in dimension 3. Geometrically, each disk can be obtained as a vertical projection of the intersection of a certain halfspace in \mathbb{R}^3 with the paraboloid $z = x^2 + y^2$ onto the plane $z = 0$. Thus a better linearization is obtained by mapping the plane onto the paraboloid by the mapping $(x, y) \mapsto (x, y, x^2 + y^2)$. This transformation, sometimes called a *lifting*, has many important consequences in computational geometry, e.g., a close relationship between Voronoi diagrams in the plane and convex polytopes in \mathbb{R}^3 ; see Edelsbrunner [1987]. A linearization of a minimum dimension for a given polynomial f can be found effectively by solving a system of linear equations; see Agarwal and Matoušek [1994].

An important example of linearization in this sense is frequently used in problems dealing with lines in \mathbb{R}^3 . A line in \mathbb{R}^3 is naturally described by four real parameters (coordinates). Given two lines ℓ, ℓ' in \mathbb{R}^3 , the predicate “does ℓ lie above ℓ' ” is nonlinear when expressed

using this four-parameter description of the lines. However, the predicate becomes linear if expressed using the so-called Plücker coordinates; these coordinates correspond to points in (projective) 5-dimensional space. Here the situation is further complicated by the necessity of dealing with oriented lines, which leads to interesting combinatorial and algorithmic problems. An initial work in this direction is by Chazelle et al. [1989]; further developments are given by Agarwal [1993] and Pellegrini et al. (e.g., Pellegrini [1992], Pellegrini and Shor [1992], and Agarwal [1993]).

For some polynomials, f , we get the best known range-searching algorithms by a direct application of linearization. In other cases faster algorithms are obtained by generalizing the methods for halfspaces into the “nonlinear” setting. For example, for disk range searching in the plane, the reduction to halfspace range searching in dimension 3 yields roughly $n^{2/3}$ query time with linear space. However, for example, one can imitate the algorithm of Matoušek [1992b], replacing halfplanes by circles and generalizing the necessary results appropriately, and arrive at a linear space algorithm with query time close to \sqrt{n} .

We will briefly discuss the complexity of algorithms constructed in this way for range-searching problems with ranges in \mathcal{R}_f . More information can be found in Agarwal and Matoušek [1994]. Roughly speaking, we can say that the exponent in the query complexity for a linear space depends on d , the dimension of the space containing the point set, while the space requirement for a logarithmic query time depends mainly on p , the number of parameters determining a range. (Here we talk about bounds for specific algorithms; the true complexity might perhaps sometimes be smaller.) For disk range searching in the plane we have $d = 2$, and the query time for linear space is approximately $n^{1-1/2} = \sqrt{n}$; the number of parameters is $p = 3$, and storage close to $O(n^3)$ guarantees a logarithmic query time. From this the reader might conclude that for a given d, p we

obtain bounds close to $n^{1-1/d}$, resp., about n^p . This could (should?) indeed be the case, but a proof is only known if $d \leq 3$, resp., $p \leq 3$. For higher dimensions there is an obstacle: the unsolved problem of an efficient decomposition of an arrangement of algebraic surfaces.

We explain the required notions on a very informal level only, since exact definitions would occupy too much space. Let $\sigma_1, \dots, \sigma_n$ be algebraic hypersurfaces in \mathbb{R}^d , i.e., subsets described by a single algebraic equation—the reader may imagine (hyper)spheres. Analogously to a hyperplane arrangement, the hypersurfaces $\sigma_1, \dots, \sigma_n$ partition \mathbb{R}^d into cells (not necessarily convex ones anymore), each cell being characterized by a particular sign pattern for the polynomials defining $\sigma_1, \dots, \sigma_n$. The number of such cells is $O(n^d)$. (The hidden constant depends on d and on the degrees of the defining polynomials.) For imitating the techniques for halfspace range searching, we still need to partition each of the cells into elementary cells (as defined in the beginning of the current section). It is not important how the particular elementary cells used for the decomposition look. It is essential that each of them is described by a constant number of parameters. For spheres, we could for instance use intersections of constantly many balls, spheres, and complements of balls as the elementary cells.

If $\sigma_1, \dots, \sigma_n$ are hyperplanes, it is not difficult to partition all cells of their arrangement into $O(n^d)$ simplices in total. It is conjectured that for general algebraic hypersurfaces of bounded degree there also exists a decomposition into approximately n^d elementary cells, but proving it seems to be hard. The best known general construction, due to Chazelle et al. [1991], gives approximately n^{2d-3} elementary cells (for $d \geq 3$; in the planar case it is easy to obtain $O(n^2)$ elementary cells). This is almost optimal in dimension 3, but for higher dimensions there remains a significant gap between the lower bound of n^d and upper bound of approximately n^{2d-3} .

The relationship between this decomposition problem to range-searching algo-

gorithms for ranges from \mathcal{R}_f is as follows. If we can decompose the arrangement of any m hypersurfaces in \mathbb{R}^d of a certain type (determined by the polynomial f) into $O(m^b)$ elementary cells, we obtain, for an n -point set, query complexity roughly $O(n^{1-1/b})$ with linear storage. Similarly, decompositions into $O(m^b)$ elementary cells for any m hypersurfaces in \mathbb{R}^p (again determined by f) imply logarithmic query time with space $O(n^{b+\epsilon})$.

Let us remark that the previously described decomposition problem has several other motivations besides geometric range searching. For instance, it is very important for decision algorithms for the theory of real closed fields, although other aspects of the decomposition become important there (the degrees of polynomial defining the elementary cells in the decomposition, etc.).

Even if the problem of decomposition can be solved satisfactorily, it is not clear whether the algorithms obtained using the techniques discussed above are close to optimal. For instance, consider the problem of range searching with cones in \mathbb{R}^3 of the form $z \geq (x-a)^2 + (y-b)^2$, a, b parameters determining the cone. (This problem is equivalent to preprocessing a set of circles in the plane so that, given a query point, the circles enclosing it can be found quickly.) Here it is only known how to get a roughly $O(n^{2/3})$ query time with linear space; in the arithmetic model, however, query complexity can be made only $O(\sqrt{n})$ with $O(n)$ generators (see Section 3). Currently it is unclear if a better algorithm also exists or if the arithmetic model is too unrealistic here and a strong lower bound holds in some other model of computation.

In concluding this section, we mention recent combinatorial results which may prove very significant for the theory of geometric range searching. Let $\sigma_1, \dots, \sigma_n \subset \mathbb{R}^d$ be graphs of algebraic functions $f_1, \dots, f_n: \mathbb{R}^{d-1} \rightarrow \mathbb{R}$, of a constant-bounded degree. The *lower envelope* of $\sigma_1, \dots, \sigma_n$ is the graph of the pointwise minimum of f_1, \dots, f_n . After a series of previous partial results,

Halperin and Sharir [1994] and Sharir [1993] recently proved nearly tight upper bound on the worst-case combinatorial complexity of the lower envelope— $O(n^{d-1+\epsilon})$, for an arbitrarily small constant $\epsilon > 0$. Some range-searching problems (typically range-emptiness problems) can be formulated to detect whether a query point lies below the lower envelope of a collection of algebraic surfaces in a small dimension d . Knowing the complexity of the lower envelope, one might hope to obtain algorithms with, say, polylogarithmic query time and storage $O(n^{d-1+\epsilon})$, or even some form of space/query time tradeoff. Currently no general algorithm (essentially point location below the lower envelope) taking advantage of the new combinatorial bound is known. The obstacle is the decomposition of the lower envelope into constant complexity cells (again!). A special case (with a particular type of surface) was solved by Mohabian and Sharir [1993]. They considered testing whether a query line lies above all spheres of a given collection; the analysis of their algorithm utilizes the new bound (for $d = 4$).

Related research aims at bounding the combinatorial complexity of a single cell in an arrangement of n algebraic surfaces of bounded degree in \mathbb{R}^d . One might hope that this complexity is close to n^{d-1} ; see Aronov and Sharir [1992] and Halperin and Sharir [1994].

6.5 Ray Shooting and Linear Optimization

In this section we consider another type of generalization of geometric range-searching problems. A good and important example is the *ray-shooting problem*, whose special case was already mentioned as Problem 2.3. In such problems we are given some set Γ of geometric objects (planes, triangles, balls), and the goal is to construct a data structure such that for a given point o and direction θ we can quickly determine the object of Γ hit by a ray ρ sent from the point o in the direction θ . (More formally, ρ is a semiline originating in o , and the question is which object of Γ is the one intersecting ρ closest to o .)

This problem is very popular in computer graphics, where it arises as an auxiliary problem in determining visibility of objects, hidden-surface elimination, ray tracing, and in other situations.

The ray-shooting problem has a somewhat different flavor than the range-searching problems, but it can be solved efficiently using data structures for suitable derived range-searching problems. We illustrate this relation on a specific problem, ray shooting in a convex polytope. Here the role of Γ is played by a set of H of hyperplanes in \mathbb{R}^d . For simplicity let us assume that none of the hyperplanes of H is vertical, and let U denote the upper unbounded cell of the arrangement of H , i.e., the set of points lying above all the hyperplanes. In our ray-shooting problem, we only permit rays originating in U .

This problem is useful in a number of applications. For instance, the problem of determining the closest point from a set $P \subset \mathbb{R}^{d-1}$ to a given point, the so-called *post office problem*, can be transformed into it. In dimensions higher than 3, this yields the most efficient algorithms for the post office problem, which can moreover be dynamized; see Matoušek and Schwarzkopf [1993].

Let us suppose that the given query ray ρ with origin $o \in U$ intersects the first hyperplane from H in a point x^* . Finding this point is our goal (together with the appropriate hyperplane). Given any point $x \in \rho$, we can decide whether x lies before x^* or after x^* : it suffices to test whether the segment ox intersects at least one hyperplane of H , and by the assumption $o \in U$ this happens iff $x \notin U$. The test whether x lies in U , that is, whether it lies above all hyperplanes of H , is the dual version of the halfspace emptiness problem, and thus the algorithms discussed in Section 6.1 are suitable for such tests.

Therefore, although we do not know x^* yet, we can efficiently determine whether some point $x \in \rho$ precedes or follows after x^* , and we would like to find x^* using tests of this type. The method of interval halving suggests itself; it can determine the position of x^* with some

Parametric search is a general strategy for algorithm design. Roughly speaking, we can say it produces algorithms for searching from algorithms for verification, under suitable assumptions.

Let us consider a problem in which the goal is to find a particular real number, t^* , which depends on some input objects. We consider these input objects fixed. Suppose that we have two algorithms at our disposal. First, an algorithm O , which for a given number t decides among the possibilities $t < t^*$, $t = t^*$, and $t > t^*$ (although it does not explicitly know t^* , only the input objects); let us call such an algorithm O the *oracle*. Second, an algorithm G (called the *generic* algorithm), whose computation depends on the input objects and on a real parameter t , and for which it is guaranteed that its computation for $t = t^*$ differs from the computation for any other $t \neq t^*$. We can also use algorithm O in the role of G , but often it is possible to employ a simpler algorithm for G . Under certain quite weak assumptions about algorithm G , the parametric search produces an algorithm for finding t^* .

The main idea is to simulate the computation of algorithm G for the (yet unknown) parameter value $t = t^*$. The computation of G of course depends on t , but we assume that all the required information about t is obtained by testing the signs of polynomials of small (constant bounded) degree in t . The coefficients in each such polynomial may depend on the input objects of the algorithm and on the outcomes of the previous tests, but not directly on t . The sign of a particular polynomial can also be tested in the unknown t^* : we find the roots t_1, \dots, t_k of the polynomial p ; we locate t^* among them using the algorithm O , and we derive the sign of $p(t^*)$ from it. In this way we can simulate the computation of the algorithm G at t^* .

If we record all tests involving t made by algorithm G during its computation, we can then find the (unique) value t^* giving appropriate results in all these tests, thereby solving the search problem.

In this version we need several calls to the oracle for every test performed by algorithm G . The second idea is to do many tests at once, whenever possible. If algorithm G executes a group of mutually independent tests with polynomials $p_1(t), \dots, p_m(t)$ (meaning that the polynomial p_i does not depend on the outcome of the test involving another polynomial p_j), we can answer all of them by $O(\log n)$ calls of the oracle: we compute the roots of all the polynomials p_1, \dots, p_m , and we locate the position of t^* among them by binary search. Parametric search will thus be particularly efficient for algorithms G implemented in parallel, with a small number of parallel steps, since the tests in one parallel step are necessarily independent in the above-mentioned sense.

Parametric search was formulated by Megiddo [1983]; the idea of simulating an algorithm at a generic value appears in Eisner and Severance [1976], Gusfield [1983], and Megiddo [1979]. A technical improvement, which sometimes reduces the running time by a logarithmic factor, was suggested by Cole [1987]. A generalization of parametric search to higher dimension, where the parameter t is a point in \mathbb{R}^d and the oracle can test the position of t^* with respect to a given hyperplane, appears in Cole et al. [1987], Cohen and Megiddo [1993], Norton et al. [1992], and Matoušek [1993a]. Currently, parametric search is a quite popular technique in computational geometry also; from numerous recent efforts, we select more or less randomly [Cole et al. 1989; Agarwal et al. 1993; 1992; Chazelle et al. 1993].

Algorithms based on parametric search, although theoretically elegant, appear quite complicated for implementation. In many specific problems, parametric search can be replaced by a randomized algorithm (see Dillencourt et al. [1992] and Matoušek [1991b]) or by other techniques (e.g., Chazelle et al. [1993] and Katz and Sharir [1993]) with a similar efficiency.

required numerical precision. This approach may not be bad in practice, but in our infinite-precision computation model it does not suffice, since an arbitrarily small interval along ρ may still contain many intersections with hyperplanes of H . Thus we cannot determine the first intersected hyperplane in a bounded number of steps. However, we can apply the method of *parametric search* (see box 5) which allows us to find x^* exactly, and usually quite efficiently. In connection with this method, data structures for the halfspace emptiness problem can be applied for ray shooting in a convex polytope as well. The space requirement

remains the same (since the data structure is identical), and the query time increases by a small power of $\log n$.

The passage described above, from halfspace emptiness testing to ray shooting in convex polytope, is illustrated in Figure 4. Figure 4(a) shows a half-space emptiness query schematically; 4(a') illustrates the dual version—testing whether a query point lies above all hyperplanes. Figure 4(b') depicts the ray-shooting problem; the dual version (b) means that we translate the bounding hyperplane of a given halfspace; we are interested in the first point hit by this hyperplane. In both the primal and the




Figure 4. Adding degrees of freedom to the halfspace emptiness query.

dual versions, we can interpret the passage from Figure 4(a) to (b) as adding one degree of freedom to the query object (point, halfspace) and looking for an extremal position within the given freedom.

The approach explained for ray shooting in convex polytopes is also applicable for other versions of the ray-shooting problem. First we formulate the derived range-searching problem, namely a test whether some initial segment ox of an admissible ray intersects at least one object from Γ . For this problem, we build a

data structure using the techniques for geometric range searching. (We usually construct a suitable multilevel data structure.) Then we apply parametric search and obtain a ray-shooting algorithm. The application of this strategy for particular ray-shooting problems does not need to be quite routine, however.

The idea of employing range-searching data structures for ray shooting is a natural one, and it appears in several papers, e.g., see De Berg et al. [1994]. The systematic approach (described above)

via parametric search was suggested in Agarwal and Matoušek [1993] and demonstrated in several examples. Some particular data structures for halfspace emptiness queries can be adapted directly for ray shooting in convex polytopes, which eliminates parametric search and speeds up query answering to some extent; see Matoušek and Schwarzkopf [1993].

Let us return to ray shooting in convex polytopes and its interpretation as membership queries for the convex polytope with one added degree of freedom (see Figure 4). From this point of view it is natural to add still more degrees of freedom to the query object, at most all d degrees of freedom, and look for some extremal position. Such a situation is shown in Figure 4(c'): instead of a single point or an (oriented) ray, we are given a linear function c (which can be interpreted as a direction in d -space, that is, a generalization of the orientation of the ray). We then want to find a vertex x_0 of the polytope U maximizing the function c .

The resulting problem is in fact a linear programming problem, but in a special situation, where the constraints are given in advance, while the optimized function is a query.¹¹ Matoušek [1993a] has shown that algorithms for halfspace emptiness queries can be transformed to solve this linear programming problem also, with query complexity increased by a polylogarithmic factor only (and with the same data structure). This result uses a multidimensional version of parametric search, and the construction proceeds by induction on the number of degrees of freedom of the query object. The resulting algorithm is quite complicated, and it would be useful to find a simpler (maybe randomized) variant.

¹¹The dual version is not so intuitive in this case—it is the finding of an empty halfspace maximizing a given linear function. Some optimal halfspace is always determined by a d -tuple of points from the given point set, or, in other words, it is supported by a facet of the convex hull; see Figure 4(c)

With a dynamic halfspace emptiness data structure as a basis, one may also insert and delete constraints. Using this machinery one can answer various queries concerning the convex hull of the set P , as we have mentioned at the end of Section 6.2.

Let us mention an application of this result to one classical computational geometry problem: that of finding extremal points. The input is an n -point set $P \subset \mathbb{R}^d$, and we want to detect which points of P are extremal, i.e., are vertices of the convex hull of P . This problem can be solved by computing a combinatorial representation of the convex hull of P , but for dimensions $d \geq 4$ this method is fairly inefficient, since the convex hull of an n -point set in \mathbb{R}^d can have combinatorial complexity of the order $n^{\lfloor d/2 \rfloor}$. Testing whether a point is extremal can be formulated as a linear programming problem in dimension d with n constraints. For different points these linear programs only differ by two constraints, so we can use linear programming in the preprocessing/query mode. We build a data structure for the appropriate constraints, and with its help we answer n queries, thereby determining the extremal points.¹² By suitably balancing the preprocessing time and query time, we obtain, for example, a total time $O(n^{4/3+\epsilon})$ in dimension 4, which is the most efficient known method.

6.6 More on Applications

The literature on computational geometry with applications for range searching is too extensive to be reviewed here. A survey paper by Agarwal [1991] contains many such applications. Let us, however, stress only one point, namely that most of the problems where range searching has been applied are not of the preprocessing/query type. An example is the

¹²It would seem that we must use a dynamic data structure, but in fact we can use a special treatment for this problem and get away with a static structure

computation of extreme points, mentioned at the end of the previous section.

Let us give one more (rather artificial) example to illustrate our point. Suppose that we want to count the number of pairwise intersections for a set S of n segments in the plane. An approach via range searching is to build a data structure which, given a query segment q , counts the number of segments of S intersected by it. Then we query the structure by each segment of S , sum the answers, and divide the sum by 2. (We must somehow deal with the singular cases of segments intersecting themselves, but this can be taken care of.) With the methods reviewed above, the required data structure is designed in a more or less routine manner (using the multilevel construction and halfplane decomposition schemes). When the preprocessing and query time are balanced appropriately, we obtain an $O(n^{4/3+\epsilon})$ algorithm.

The reader might suspect that this is a very sloppy way of solving the problem, and that we could do better with some “global” method, rather than imposing the preprocessing/query mode not present in the original problem. This is only partially true: there *do* exist somewhat better and simpler algorithms (those based on cuttings, e.g., see Agarwal [1990]), but their complexity is only slightly better, about $O(n^{4/3} \log^c n)$ for a small constant c . (We do not give a specific value, since by using some newer results on cuttings one can improve the published algorithms to some extent.) And it is believed that this might be roughly the actual complexity of the problem.¹³ Thus, the application of range-searching tools gave us a theoretically good algorithm very quickly.

This situation does not seem to be an exception. For many problems, the algorithm based on range searching indicates

¹³No lower bound larger than $n \log n$ is known, and problems of this type seem much less tractable than the query-type problems as far as lower bounds are concerned.

the complexity we can expect for other problem-specific and hopefully simpler and more practical algorithms also. And for many problems that are more complicated, no such specific algorithms have been elaborated, and the range-searching approach has remained the best published solution. Examples are hidden-surface removal problems (e.g., see De Berg et al. [1994]), counting circular arc intersections [Agarwal et al. 1993], problems concerning lines in space (e.g., Pellegrini [1992]), and the above-mentioned extreme point detection, to quote only a few.

7. A NOTE IN CONCLUSION

This text was intended as a survey article on geometric range searching. Although its extent is much larger than originally planned, it does not include everything that it perhaps should include. The choice of the material necessarily reflects the taste and knowledge of the author, and in topics mentioned only briefly (such as the broad area of applications) we can only refer to the literature.

ACKNOWLEDGMENT

I would like to thank Pankaj K. Agarwal and Raimund Seidel for reading preliminary versions of this survey and for numerous useful comments.

REFERENCES

- AGARWAL, P. K. 1993. On stabbing lines for polyhedra in 3d. Tech. Rep. CS-1993-09, Dept. Computer Science, Duke University, Durham, N.C.
- AGARWAL, P. K. 1991. Geometric partitioning and its applications. In *Computational Geometry: Papers from the DIMACS Special Year*, J. E. Goodman, R. Pollack, and W. Steiger, Eds. American Mathematical Society.
- AGARWAL, P. K. 1990. Partitioning arrangements of lines: II. Applications. *Discrete Comput. Geom.* 5, 533–573.
- AGARWAL, P. K. AND MATOUŠEK, J. 1994. On range searching with semialgebraic sets. *Discrete Comput. Geom.* 11, 393–418.
- AGARWAL, P. K. AND MATOUŠEK, J. 1993. Ray shooting and parametric search. *SIAM J. Comput.* 22, 4, 794–806.
- AGARWAL, P. K. AND MATOUŠEK, J. 1991. Dy-

- namical half-space range reporting and its applications. Tech. Rep. CS-1991-43, Duke University, Durham, N.C. Extended abstract, including also results of D. Eppstein: *Proc. 33. IEEE Symposium on Foundations of Computer Science* (1992), pp. 51–60.
- AGARWAL, P. K. AND SHARIR, M. 1993. Applications of a new partitioning scheme. *Discrete Comput. Geom.* 9, 11–38.
- AGARWAL, P. K., ARONOV, B., SHARIR, M., AND SURI, S. 1993. Selecting distances in the plane. *Algorithmica* 9, 495–513.
- AGARWAL, P. K., SHARIR, M., AND TOLEDO, S. 1992. Applications of parametric searching in geometric optimization. In *Proceedings of the 3rd ACM-SIAM Symposium on Discrete Algorithms*. ACM, New York, 72–82.
- AGARWAL, P. K., PELLEGRINI, M., AND SHARIR, M. 1993. Counting circular arc intersections. *SIAM J. Comput.* 22, 778–793.
- AGGARWAL, A., HANSEN, M., AND LEIGHTON, T. 1990. Solving query-retrieval problems by compacting Voronoi diagrams. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*. ACM, New York, 331–340.
- ALON, N., HAUSSLER, D., AND WELZL, E. 1987. Partitioning and geometric embedding of range spaces of finite Vapnik-Chervonenkis dimension. In *Proceedings of the 3rd Annual ACM Symposium on Computational Geometry*. ACM, New York, 331–340.
- ARONOV, B. AND SHARIR, M. 1992. Castles in the air revisited. In *Proceedings of the 8th Annual ACM Symposium on Computational Geometry*. ACM, New York, 146–156. Also to appear in *Discrete and Computational Geometry*.
- ARONOV, B., PELLEGRINI, M., AND SHARIR, M. 1993. On the zone of a surface in a hyperplane arrangement. *Discrete Comput. Geom.*, 9, 2, 177–186.
- AVIS, D. 1984. Non-partitionable point sets. *Inf. Process. Lett.* 19, 125–129.
- BENTLEY, J. L. 1980. Multidimensional divide-and-conquer. *Commun. ACM* 23, 4, 214–229.
- BENTLEY, J. L. AND SAXE, J. B. 1980. Decomposable searching problems I: Static-to-dynamic transformation. *J. Alg.* 1, 301–358.
- BRÖNNIMANN, H., CHAZELLE, B., AND PACH, J. 1993. How hard is halfspace range searching. *Discrete Comput. Geom.* 10, 143–155.
- BURNIKEL, C., MEHLHORN, K., AND SCHIRRA, S. 1994. On degeneracy in geometric computations. In *Proceedings of the 5th Annual Symposium on Discrete Algorithms*. ACM, New York, 16–23.
- CHAZELLE, B. 1993a. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.* 9, 2, 145–158.
- CHAZELLE, B. 1993b. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.* 10, 377–409.
- CHAZELLE, B. 1990a. Lower bounds for orthogonal range searching. I: The reporting case. *J. ACM* 37, 200–212.
- CHAZELLE, B. 1990b. Lower bounds for orthogonal range searching, II: The arithmetic model. *J. ACM* 37, 439–463.
- CHAZELLE, B. 1989. Lower bounds on the complexity of polytope range searching. *J. Am. Math. Soc.* 2, 637–666.
- CHAZELLE, B. 1988. A functional approach to data structures and its use in multidimensional searching. *SIAM J. Comput.* 17, 427–462.
- CHAZELLE, B. 1986. Filtering Search: A new approach to query-answering. *SIAM J. Comput.* 15, 703–724.
- CHAZELLE, B. AND FRIEDMAN, J. 1994. Point location among hyperplanes and unidirectional ray-shooting. *Comput. Geom. Theor. Appl.* 4, 53–62.
- CHAZELLE, B. AND FRIEDMAN, J. 1990. A deterministic view of random sampling and its use in geometry. *Combinatorica* 10, 3, 229–249.
- CHAZELLE, B. AND MATOUŠEK, J. 1993. On linear-time deterministic algorithms for optimization problems in fixed dimension. In *Proceedings of the 4th ACM-SIAM Symposium on Discrete Algorithms*. ACM, New York, 281–290.
- CHAZELLE, B. AND PREPARATA, F. P. 1986. Half-space range search: An algorithmic application of k -sets. *Discrete Comput. Geom.* 1, 83–93.
- CHAZELLE, B. AND ROSENBERG, B. 1992. Lower bounds on the complexity of simplex range reporting on a pointer machine. In *Proceedings of the 19th International Colloquium on Automata, Languages, and Programming. Lecture Notes in Computer Science*, vol. 623. Springer-Verlag, Berlin, 439–449. Also to appear in *Comput. Geom. Theory Appl.*
- CHAZELLE, B. AND ROSENBERG, B. 1991. The complexity of computing partial sums off-line. *Int. J. Comput. Geom. Appl.* 1, 1, 33–45.
- CHAZELLE, B. AND WELZL, E. 1989. Quasi-optimal range searching in spaces of finite VC-dimension. *Discrete Comput. Geom.* 4, 467–489.
- CHAZELLE, B., COLE, R., PREPARATA, F. P., AND YAP, C. K. 1986. New upper bounds for neighbor searching. *Inf. Control* 68, 105–124.
- CHAZELLE, B., EDELSBRUNNER, H., GUIBAS, L., AND SHARIR, M. 1993. Diameter, width, closest line pair and parametric searching. *Discrete Comput. Geom.* 10, 183–196.
- CHAZELLE, B., EDELSBRUNNER, H., GUIBAS, L., AND SHARIR, M. 1991. A singly-exponential stratification scheme for real semi-algebraic varieties and its applications. *Theor. Comput. Sci.* 84, 77–105.
- CHAZELLE, B., EDELSBRUNNER, H., GUIBAS, L. J., AND SHARIR, M. 1989. Lines in space: Combinatorics, algorithms, and applications. In *Proceedings of the 21st Annual ACM Symposium on the Theory of Computation*. ACM, New York, 382–393.

- CHAZELLE, B., GUIBAS, L. J., AND LEE, D. T. 1985. The power of geometric duality. *BIT* 25, 76–90.
- CHAZELLE, B., SHARIR, M., AND WELZL, E. 1992. Quasi-optimal upper bounds for simplex range searching and new zone theorems. *Algorithmica* 8, 407–429.
- CHEW, L. P. 1985. Building Voronoi diagrams for convex polygons in linear expected time. Tech. Report, Dept. of Mathematical and Computer Science, Dartmouth College, Hanover, N.H.
- CLARKSON, K. L. 1994. An algorithm for approximate closest-point queries. In *Proceedings of the 10th Annual ACM Symposium on Computational Geometry*. ACM, New York, 160–164.
- CLARKSON, K. L. 1988a. Applications of random sampling in computational geometry, II. In *Proceedings of the 4th Annual ACM Symposium on Computational Geometry*. ACM, New York, 1–11.
- CLARKSON, K. L. 1988b. A randomized algorithm for closest-point queries. *SIAM J. Comput.* 17, 830–847.
- CLARKSON, K. L. 1987. New applications of random sampling in computational geometry. *Discrete Comput. Geom.* 2, 195–222.
- CLARKSON, K., EDELSBRUNNER, H., GUIBAS, L., SHARIR, M. AND WELZL, E. 1990. Combinatorial complexity bounds for arrangements of curves and spheres. *Discrete Comput. Geom.* 5, 99–160.
- CLARKSON, K. L., EPPSTEIN, D., MILLER, G. L., STURTIVANT, C., AND TENG, S.-H. 1993. Approximating center points with iterated Radon points. In *Proceedings of the 9th Annual ACM Symposium on Computational Geometry*. ACM, New York, 91–98.
- COHEN, E. AND MEGIDDO, N. 1993. Strongly polynomial-time and NC algorithms for detecting cycles in dynamic graphs. *J. ACM* 40, 791–832.
- COLE, R. 1987. Slowing down sorting networks to obtain faster sorting algorithms. *J. ACM* 34, 200–208.
- COLE, R. 1985. Partitioning point sets in 4 dimensions. In *Proceedings of the 12th International Colloquium on Automata Language Programming*. Lecture Notes in Computer Science, vol. 194. Springer-Verlag, Berlin, 111–119.
- COLE, R. AND YAP, C. K. 1985. Geometric retrieval problems. *Inf. Control* 63, 39–57.
- COLE, R., SALOWE, J., STEIGER, W., AND SZEMERÉDI E. 1989. An optimal-time algorithm for slope selection. *SIAM J. Comput.* 18, 792–810.
- COLE, R., SHARIR, M. AND YAP, C. 1987. On k -hulls and related problems. *SIAM J. Comput.* 16, 1, 61–67.
- DE BERG, M., HALPERIN, D., OVERMARS, M., SNOEYINK, J., AND VAN KREVELD, M. 1994. Efficient ray shooting and hidden surface removal. *Algorithmica* 12, 30–53.
- DILLENCOURT, M. B., MOUNT, D. M., AND NETANYAHU, N. S. 1992. A randomized algorithm for slope selection. *Int. J. Comput. Geom. Appl.* 2, 1–27.
- DOBKIN, D. P. AND EDELSBRUNNER, H. 1987. Space searching for intersecting objects. *J. Alg.* 8, 348–361.
- DOBKIN, D. AND EDELSBRUNNER, H. 1984. Organizing point sets in two and three dimensions. Tech. Rep. F130, Technische Univ. Graz, Graz, Austria.
- EDELSBRUNNER, H. 1987. *Algorithms in Combinatorial Geometry*. *EATCS Monographs on Theoretical Computer Science*, vol. 10. Springer-Verlag, Heidelberg, Germany.
- EDELSBRUNNER, H. AND HUBER, F. 1984. Dissecting sets of points in two and three dimensions. Rep. F138, Inst. Informationsverarb., Tech. Univ. Graz, Graz, Austria.
- EDELSBRUNNER, H. AND MÜCKE, E. P. 1990. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graph.* 9, 66–104.
- EDELSBRUNNER, H. AND WELZL, E. 1986. Halfplanar range search in linear space and $O(n^{0.695})$ query time. *Inf. Process. Lett.* 23, 289–293.
- EDELSBRUNNER, H., GUIBAS, L., HERSHBERGER, J., SEIDEL, R., SHARIR, M., SNOEYINK, J., AND WELZL, E. 1989. Implicitly representing arrangements of lines or segments. *Discrete Comput. Geom.* 4, 433–466.
- EDELSBRUNNER, H., KIRKPATRICK, D. G., AND MAURER, H. A. 1982. Polygonal intersection searching. *Inf. Process. Lett.* 14, 74–79.
- EISNER, M. AND SEVERANCE, D. 1976. Mathematical techniques for efficient record segmentation in large shared database. *J. ACM* 23, 619–635.
- EMIRIS, I. AND CANNY, J. 1991. A general approach to removing degeneracies. In *Proceedings of the 32nd Annual IEEE Symposium on Foundations of Computer Science*. IEEE, New York, 405–413.
- EMIRIS, I. AND CANNY, J. 1992. An efficient approach to removing geometric degeneracies. In *Proceedings of the 8th Annual ACM Symposium on Computational Geometry*. ACM, New York, 74–82.
- FORTUNE, S. AND VAN WYK, C. J. 1993. Efficient exact arithmetic for computational geometry. In *Proceedings of the 9th Annual ACM Symposium on Computational Geometry*. ACM, New York, 163–172.
- FREDMAN, M. L. 1981. Lower bounds on the complexity of some optimal data structures. *SIAM J. Comput.* 10, 1–10.
- GABOW, H. N., BENTLEY, J. L., AND TARJAN, R. E. 1984. Scaling and related techniques for geometry problems. In *Proceedings of the 16th Annual ACM Symposium on Theory of Computing*. ACM, New York, 135–143.
- GOODMAN, J. E., POLLACK, R., AND STEIGER, W., EDS. 1991. *Computational Geometry: Papers from*

- the DIMACS Special Year. American Mathematical Society, Providence, R.I.
- GUSFIELD, D. 1983. Parametric combinatorial computing and a problem of program module distribution. *J. ACM* 30, 551–563
- HALPERIN, D. AND SHARIR, M. 1994. Almost tight upper bounds for the single cell and zone problems in three dimensions. In *Proceedings of the 10th Annual ACM Symposium on Computational Geometry*. ACM, New York.
- HALPERIN, D. AND SHARIR, M. 1994. New bounds for lower envelopes in three dimension, with applications to visibility in terrains. *Discrete Comput. Geom.* 12, 313–326.
- HARRIS, J. 1992. *Algebraic Geometry (A First Course)*. Springer-Verlag, Berlin.
- HAUSSLER, D. 1991. Sphere packing numbers for subsets of the Boolean n -cube with bounded Vapnik-Chervonenkis dimension. Tech Report UCSC-CRL-91-41, Univ. of California at Santa Cruz.
- HAUSSLER, D. AND WELZL, E. 1987. Epsilon-nets and simplex range queries. *Discrete Comput. Geom.* 2, 127–151.
- KATZ, M. J. AND SHARIR, M. 1993. An expander-based approach to geometric optimization. In *Proceedings of the 9th Annual ACM Symposium on Computational Geometry*. ACM, New York, 198–207.
- KOMLÓS, J., SZEMERÉDI, E., AND PINTZ, J. 1982. A lower bound for Heilbronn's problem. *J. London Math. Soc.* 25, 2, 13–24.
- LIPTON, R. J. AND TARJAN, R. E. 1980. Applications of a planar separator theorem. *SIAM J. Comput.* 9, 615–627.
- MATOUŠEK, J. 1994. Tight upper bounds for the discrepancy of halfspaces. KAM Series in Discrete Mathematics, (Tech. Rep.), Charles Univ., Prague. To appear in *Discrete Comput. Geom.*
- MATOUŠEK, J. 1993a. Linear optimization queries. *J. Alg.* 14, 432–448. The results combined with results of O. Schwarzkopf also appear in *Proceedings 8th ACM Symp. Comput. Geom.*, 1992, pp. 16–25.
- MATOUŠEK, J. 1993b. On vertical ray shooting in arrangements. *Comput. Geom. Theor Appl* 2, 5, 279–285.
- MATOUŠEK, J. 1993. Range searching with efficient hierarchical cuttings. *Discrete Comput. Geom.* 10, 2, 157–182.
- MATOUŠEK, J. 1992a. Reporting points in halfspaces. *Comput. Geom. Theor. Appl* 2, 3, 169–186.
- MATOUŠEK, J. 1992b. Efficient partition trees. *Discrete Comput. Geom.* 8, 315–334.
- MATOUŠEK, J. 1991a. Approximations and optimal geometric divide-and-conquer In *Proceedings of the 23rd Annual ACM Symposium on the Theory of Computing*. ACM, New York, 505–511. Also to appear in *J. Comput. Syst. Sci.*
- MATOUŠEK, J. 1991b. Randomized optimal algorithm for slope selection. *Inf. Process. Lett.* 39, 183–187.
- MATOUŠEK, J. 1991c. Cutting hyperplane arrangements. *Discr. Comput. Geom.* 6, 5, 385–406.
- MATOUŠEK, J. 1990. Construction of ϵ -nets. *Discr. Comput. Geom.* 5, 427–448.
- MATOUŠEK, J. AND SCHWARZKOPF, O. 1993. On ray shooting in convex polytopes. *Discr. Comput. Geom.* 10, 2, 215–232
- MATOUŠEK, J. AND SPENCER, J. 1994. Discrepancy in arithmetic progressions. KAM Series (Tech. Report), Charles Univ., Prague. Also to appear in *Am. J. Math.*
- MATOUŠEK, J. AND WELZL, E. 1992. Good splitters for counting points in triangles. *J. Alg.* 13, 307–319.
- MATOUŠEK, J., WELZL, E., AND WERNISCH, L. 1993. Discrepancy and ϵ -approximations for bounded VC-dimension. *Combinatorica* 13, 455–466.
- MEGIDDO, N. 1983. Applying parallel computation algorithms in the design of serial algorithms. *J. ACM* 30, 852–865.
- MEGIDDO, N. 1979. Combinatorial optimization with rational objective functions. *Math. Oper. Res.* 4, 414–424.
- MEHLHORN, K. 1984. *Data Structures and Algorithms Vol. 3. Multi-dimensional Searching and Computational Geometry*. Springer-Verlag, Heidelberg, Germany.
- MOHABAN, S AND SHARIR, M. 1993. Ray shooting amidst spheres in three dimensions and related problems. Manuscript, Tel Aviv Univ.
- MULMULEY, K. 1993 *Computational Geometry: An Introduction Through Randomized Algorithms*. Prentice-Hall, Englewood Cliffs, N.J
- MULMULEY, K. 1992. Randomized geometric algorithms and pseudo-random generators. In *Proceedings of the 33rd Annual IEEE Symposium on the Foundation of Computer Science*. IEEE, New York. 90–100.
- MULMULEY, K. 1991a. A fast planar partition algorithm, II. *J. ACM* 38, 74–103.
- MULMULEY, K. 1991b. Randomized multidimensional search trees: Dynamic sampling. In *Proceedings of the 7th Annual ACM Symposium on Computational Geometry*. ACM, New York, 121–131.
- MULMULEY, K. 1991c. Randomized multidimensional search trees: Further results in dynamic sampling. In *Proceedings of the 32nd Annual IEEE Symposium on the Foundations of Computer Science*. IEEE, New York, 216–227.
- MULMULEY, K. 1991d. Randomized multidimensional search trees: Lazy balancing and dynamic shuffling. In *Proceedings of the 32nd Annual IEEE Symposium on the Foundations of Computer Science*. IEEE, New York, 180–196.

- MULMULEY, K. 1988. A fast planar partition algorithm, I. In *Proceedings of the 29th Annual IEEE Symposium on the Foundations of Computer Science*. IEEE, New York, 580–589.
- NORTON, C. H., PLOTKIN, S. A. AND TARDOS, É. 1992. Using separation algorithms in fixed dimensions. *J. Alg.* 13, 79–98.
- OVERMARS, M., SCHIPPER, H., AND SHARIR, M. 1990. Storing line segments in partition trees. *BIT* 30, 385–403.
- OVERMARS, M. H. 1988. Efficient data structures for range searching on a grid. *J. Alg.* 9, 254–275.
- OVERMARS, M. H. 1983. *The Design of Dynamic Data Structures. Lecture Notes in Computer Science*, vol. 156. Springer-Verlag, Berlin.
- PACH, J. ED. 1993. *New Trends in Discrete and Computational Geometry. Algorithms and Combinatorics*, vol. 10. Springer-Verlag, Berlin.
- PATERSON, M. S. AND YAO, F. F. 1986. Point retrieval for polygons. *J. Alg.* 7, 441–447.
- PELLEGRINI, M. 1992. Incidence and nearest-neighbor problems for lines in 3-space. In *Proceedings of the 8th Annual ACM Symposium on Computational Geometry*. ACM, New York, 130–137.
- PELLEGRINI, M. AND SHOR, P. 1992. Finding stabbing lines in 3-space. *Discrete Comput. Geom.* 8, 191–208.
- PREPARATA, F. P. AND SHAMOS, M. I. 1985. *Computational Geometry: An Introduction*. Springer-Verlag, New York.
- PREPARATA, F. P., AND TAMASSIA, R. 1992. Efficient point location in a convex spatial cell-complex. *SIAM J. Comput.* 21, 267–280.
- RABIN, M. O. 1976. Probabilistic algorithms. In *Algorithms and Complexity*, J. F. Traub, Ed. Academic Press, New York, 21–30.
- ROTH, K. 1976. Developments in the triangle problem. *Adv. Math.* 22, 364–385.
- SCHIPPER, H. AND OVERMARS, M. H. 1991. Dynamic partition trees. *BIT* 31, 421–436.
- SCHWARZKOPF, O. 1992. Ray shooting in convex polytopes. In *Proceedings of the 8th Annual ACM Symposium on Computational Geometry*. ACM, New York, 286–295.
- SCHWARZKOPF, O. 1991. Dynamic maintenance of geometric structures made easy. In *Proceedings of the 32nd Annual IEEE Symposium on the Foundation of Computer Science*. IEEE, New York, 197–206.
- SEIDEL, R. 1994. The nature and meaning of perturbations in geometric computing. In *Proceedings of Symposium on Theoretical Aspects of Computer Science, Lecture Notes in Computer Science*. Springer-Verlag, Berlin.
- SEIDEL, R. 1991a. A simple and fast incremental randomized algorithm for computing trapezoidal decompositions and for triangulating polygons. *Comput. Geom. Theor. Appl.* 1, 51–64.
- SEIDEL, R. 1991b. Small-dimensional linear programming and convex hulls made easy. *Discrete Comput. Geom.* 6, 423–434.
- SHARIR, M. 1993. Almost tight upper bounds for lower envelopes in higher dimensions. In *Proceedings of the 34th Annual IEEE Symposium on the Foundations of Computer Science (FOCS 93)*. IEEE, New York, 498–507.
- WELZL, E. 1992. On spanning-trees with low crossing numbers. In *Lecture Notes in Computer Science*. Vol. 594. Springer-Verlag, Berlin, 233–249.
- WELZL, E. 1988. Partition trees for triangle counting and other range searching problems. In *Proceedings of the 4th Annual ACM Symposium on Computational Geometry*. ACM, New York, 23–33.
- WILLARD, D. E. 1985. New data structures for orthogonal range queries. *SIAM J. Comput.* 14, 232–253.
- WILLARD, D. E. 1982. Polygon retrieval. *SIAM J. Comput.* 11, 149–165.
- WILLARD, D. E. AND LUEKER, G. S. 1985. Adding range restriction capability to dynamic data structures. *J. ACM* 32, 597–617.
- YAO, F. F. 1983. A 3-space partition and its applications. In *Proceedings of the 15th Annual ACM Symposium on the Theory of Computation*. ACM, New York, 258–263.
- YAO, A. C. AND YAO, F. F. 1985. A general approach to D -dimensional geometric queries. In *Proceedings of the 17th Annual ACM Symposium on the Theory of Computation*. ACM, New York, 163–168.
- YAO, F. F., DOBKIN, D. P., EDELSBRUNNER, H., AND PATERSON, M. S. 1989. Partitioning space for range queries. *SIAM J. Comput.* 18, 371–384.