Mathematical Snapshots from the Computational Geometry Landscape

Jiří Matoušek

Abstract. We survey some mathematically interesting techniques and results that emerged in computational geometry in recent years.

1991 Mathematics Subject Classification: 068R99, 52C99, 52C10
Keywords and Phrases: computational geometry, combinatorial geometry, arrangement, Davenport-Schinzel sequence

We survey some mathematically interesting notions, techniques, and results that emerged in the field of computational geometry in recent years.

Computational geometry is a branch of theoretical computer science which constituted sometimes around the year 1980. It considers the design of efficient algorithms for computing with geometric objects in the Euclidean space $\mathbf{R}^d$. The objects are simple, like points, lines, spheres, etc., but there are many of them. The space dimension d is usually considered constant—many problems are studied mainly in the plane or in $\mathbf{R}^3$. As for general references, there is one fresh handbook [20] and another one pending [31]. A recent introductory textbook is [16]. Some mathematical spinoffs are nicely treated in [29].

Although this field mainly emphasizes algorithms, it has many fine purely mathematical results. I have selected a few of them for this overview quite subjectively (with many other, perhaps even nicer things omitted). Since they include the ideas of many researchers (my results being a tiny part only), it is not possible to give explicit credits to all of the contributors and to always refer to original sources (rather than surveys) in the limited space.

Combinatorial complexity of arrangements

The *arrangement* of a finite set of lines in the plane is a partition of the plane into cells of dimension 0, 1, and 2. The 0-cells (vertices) are the intersections of the lines, the 1-cells (edges) are the portions of the lines between vertices, and the 2-cells are the open convex polygons left after removing the lines from the plane. More generally, for a collection $H = \{h_1, h_2, \ldots, h_n\}$ of sets in $\mathbf{R}^d$, the arrangement of H is a decomposition of $\mathbf{R}^d$ into connected cells, where each cell is a connected component of the set of points lying in all of the sets h_i with $i \in I$ and in no h_j with $j \notin I$, for some index set $I \subseteq \{1, 2, \ldots, n\}$. In computational geometry, the most general sets considered in the role of the h_i's are usually the so-called *surface patches*, which means $(d-1)$-dimensional closed semialgebraic sets defined by Boolean combinations of polynomial inequalities; moreover, both the number of inequalities and the degree of the polynomials are bounded by some constant.

Arrangements, especially arrangements of hyperplanes, have been investigated for a long time from various points of view. In the direction of research reflected, e.g., by the recent book [28], one is mainly interested in topological and algebraic properties of the whole arrangement. Computational geometers have mostly studied different aspects, primarily asymptotic bounds on the combinatorial complexity of various parts of arrangements,[1] and while the number n of sets in H is considered large, d is fixed (and small). Some important problems also lead to considering arrangements of less "regular" objects than hyperplanes, such as segments in the plane, triangles in space, or even pieces of complicated algebraic surfaces in $\mathbf{R}^d$. Two thorough and up-to-date surveys by Agarwal and Sharir in [31] complement our sketchy exposition here and in the next section.

The total complexity, i.e. the total number of cells, of an arrangement is quite well understood. Exact formulas are known for hyperplane arrangements, and fairly precise estimates exist for arrangements of surface patches (rough bounds for surface patches come from old papers in real-algebraic geometry by Petrov and Oleinik, Milnor, and Thom, and there are some recent refinements, such as [7]). The complexity is always at most $O(n^d)$.[2] More challenging problems concern the complexity of certain portions of the arrangements; some of them are schematically illustrated in Fig. 1.

The *zone* of a set $X \subseteq \mathbf{R}^d$ in an arrangement consists of the cells intersecting X. For hyperplane arrangements, the complexity of the zone of any hyperplane is $O(n^{d-1})$ [17]. The zone of a low-degree algebraic surface, or of an arbitrary convex surface, in a hyperplane arrangement has at most $O(n^{d-1}\log n)$ complexity [5].

The *level* k in a hyperplane arrangement consists of the $(d-1)$-dimensional cells, i.e. edges in the case of lines in $\mathbf{R}^2$, with exactly k of the hyperplanes below them (where the x_d-axis is considered vertical, say). The maximum complexity of the k-level is a tantalizing open problem even for lines in the plane; we refer to the paper by Welzl in this volume for more information.

Next, we discuss the *lower envelope* of an arrangement. Informally, this is the part of the arrangement that can be seen by an observer sitting at $(0,0,\ldots,0,-\infty)$. The lower envelope in an arrangement of hyperplanes is the surface of a convex polyhedron with at most n facets, whose maximum complexity, of the order $n^{\lfloor d/2\rfloor}$, is known precisely (since McMullen's paper in 1970). This bound is trivial in the plane, but already for planar arrangements of segments, the lower envelope question is hard.

If we number the segments 1 through n and write down the numbers of the segments as they are encountered along the lower envelope from left to right, we get a

[1] If X is a set of cells in an arrangement, the *(combinatorial) complexity* of X is the number of cells of the arrangement that are contained in the closure of X. Typically, this complexity is asymptotically dominated by the number of vertices of the arrangement in the closure of X.

[2] Here and in the sequel, the constants hidden in the $O(.)$ and $\Omega(.)$ notations generally depend on d, and, in some cases, on other parameters declared fixed. For instance, here the constant also depends on the degree and formula size of the surface patches forming the arrangement.

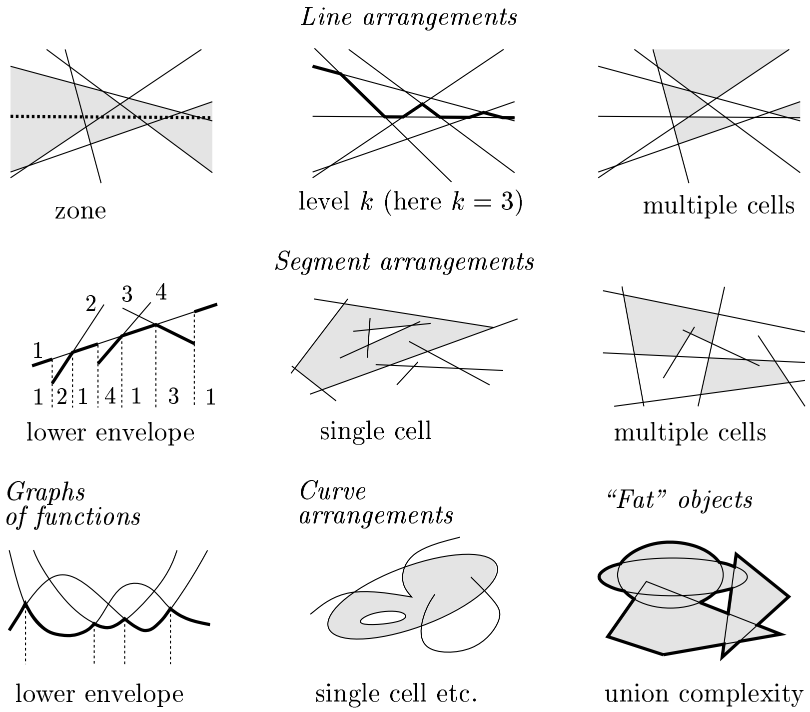

Figure 1: A bestiary of planar arrangement problems

sequence $a_1a_2a_3 \ldots a_m$, for which the following conditions hold: $a_i \in \{1, 2, \ldots, n\}$, $a_i \neq a_{i+1}$, and there is no (not necessarily contiguous) subsequence of the form *ababa*, where $a \neq b$. Any finite sequence satisfying these conditions is called a *Davenport-Schinzel sequence* (or DS-sequence for short) of order 3 over the symbols $1, 2, \ldots, n$. For DS-sequences of order s, the forbidden pattern is $abab \ldots$ with $s+2$ letters. Such sequences are obtained, e.g., from lower envelopes of x-monotone curves (i.e. graphs of univariate functions), such that any two of the curves intersect in at most s points (a typical example are graphs of degree-s polynomials). Davenport and Schinzel started investigating $\lambda_s(n)$, the maximum possible length of a DS-sequence of order s over n symbols, in 1965. Fairly precise estimates (asymptotically tight for many s's) were proved by Sharir, Hart, Agarwal, and Shor in the late 1980s (see [33] for an account). The results are remarkable: while $\lambda_1(n)$ and $\lambda_2(n)$ are easily seen to be linear, for any fixed $s \geq 3$, $\lambda_s(n)/n$ grows to infinity with $n \to \infty$, but incredibly slowly. For example, $\lambda_3(n)$ is asymptotically bounded by constant multiples of $n\alpha(n)$ from both above and below, where $\alpha(n)$ is the inverse of the Ackermann function.[3] For all practical purposes, for each fixed

[3] If we define a hierarchy of functions by $f_1(n) = 2n$ and $f_{k+1}(n) = f_k \circ f_k \circ \cdots \circ f_k(2)$ ($(n-1)$-fold composition), then the Ackermann function of n is $A(n) = f_n(n)$, and $\alpha(n) = \min\{k \geq 1 \colon A(k) \geq n\}$. For example, $A(4)$ is an exponential tower of 2s of height 2^{16}.

$s \geq 3$, $\lambda_s(n)$ behaves like a linear function, but it is nonlinear in a very subtle manner, and hence any proofs of the correct bounds must be quite complicated.

The maximum complexity of the lower envelope for segments is at most $\lambda_3(n) = O(n\alpha(n))$, and a construction by Wiernik and Sharir, later simplified by Shor, provides an arrangement of segments with lower envelope of complexity $\Omega(n\alpha(n))$. Thus, similar to DS-sequences, lower envelopes of segments are no laughing matter.

Before proceeding with the discussion of lower envelopes, we mention recent developments in generalized DS-sequences. In the original definition, the forbidden pattern *ababa*... is made of two letters. Klazar, Valtr, and Adamec studied forbidden patterns consisting of more letters, such as *abccbaabc* (for a forbidden pattern with k distinct letters, an analogue of the condition $a_i \neq a_{i+1}$ for DS-sequences is that any k consecutive symbols in the sequences be all distinct). They proved that for any fixed forbidden pattern, the maximum length of a sequence in n symbols is near-linear in n, and they characterized numerous cases where a linear bound holds (see e.g. [22, 23]). One forbidden pattern of the latter type is *abcdedcbabcde* (or analogous with more letters); this result was used by Valtr [35] for solving interesting problems concerning geometric graphs. A geometric graph is a drawing of a graph in the plane with edges drawn as straight segments (possibly crossing); they have recently been studied by Pach, Katchalski, Last, Károlyi, Tóth, and others.

The main result for lower envelopes in higher dimensions is quite recent, due to Sharir and Halperin [21, 32]. For an arrangement of surface patches in $\mathbf{R}^d$, with some mild additional technical assumptions, they prove lower envelope complexity bound of $O(n^{d-1+\varepsilon})$ for any fixed $\varepsilon > 0$, which is nearly tight (there is an $\Omega(n^{d-1}\alpha(n))$ lower bound). As a sample of techniques in the area, we demonstrate this proof in the planar case. This is a ridiculous setting, since here much better results are obtained via DS-sequences, but the higher-dimensional case is too complicated to fit here.

So let us consider a set H of n x-monotone curves (such as in Fig. 1 bottom left), any two intersecting in at most s points (s fixed). Moreover, assume for convenience that no 3 curves have a common intersection. Let $L = L(H)$ be the set of vertices on the lower envelope and let $f(n)$ denote the maximum possible cardinality of L in this situation. We aim at proving $f(n) = O(n^{1+\varepsilon})$.

First, let k be an auxiliary parameter, $2 \leq k \leq \frac{n}{2}$, let $L^{<k}$ be the set of vertices in the arrangement of H at level smaller than k (i.e. with fewer than k curves below them), and let $f^{<k}(n)$ be the maximum possible cardinality of $L^{<k}$.

Lemma. $f^{<k}(n) = O\big(k^2 f(\lfloor n/k \rfloor)\big)$.

Here is a beautiful probabilistic argument of Clarkson and Shor [15]. Suppose that $f^{<k}(n)$ is attained for H, set $r = \lfloor n/k \rfloor$, and let $R \subset H$ be an r-element subset of H picked uniformly at random. First, we lower-bound the expected size of $L(R)$. Consider a vertex $v \in L^{<k}(H)$ at a level $j < k$. Such a v appears in $L(R)$ iff both the curves defining v fall in R and none of the j curves below v does, and so $\mathrm{Prob}\,[v \in L(R)] = \binom{n-2-j}{r-2}/\binom{n}{r}$. Calculation shows that this probability is $\Omega(k^{-2})$, and so the expected size of $L(R)$ is $\Omega(k^{-2} f^{<k}(n))$. At the same time, $|L(R)| \leq f(r)$ for all R, and the lemma follows by comparing these two bounds.

Next, we partition the set $L = L(H)$ into subsets $L_1, \ldots, L_s$, with L_i consisting of the vertices of L that are the ith leftmost intersections of their two curves. Divide $L^{<k}$ similarly, and let $f_i(n)$ and $f_i^{<k}(n)$ be the corresponding maximum possible cardinalities.

The strategy of the proof is "there shouldn't better be any vertices on the lower envelope, and if there are, someone is going to pay for it". To find out who pays for a vertex $v \in L_i$, we start walking from v to the left along the curve h passing through v and not being on the lower envelope on the left of v. We charge every vertex encountered $\frac{1}{k_i}$ units, where k_i is an integer parameter (to be fixed later). If k_i vertices are encountered without returning to the lower envelope or escaping to $-\infty$ then the charging is complete. Otherwise, if we end up at $-\infty$, we charge 1 to the curve h itself. Finally, if we are back at the lower envelope without having passed at least k_i vertices then, crucially, we must have crossed the second curve h' defining the vertex v again, at a vertex $v' \in L_{i-1}^{<k_i}$, and this v' pays 1 for v. A picture illustrates these three cases of charging:

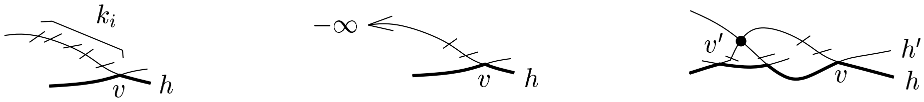

If we do this charging for all vertices $v \in L_i$ then, altogether, each curve was charged at most 1 and each vertex of $L^{<k_i}$ was charged at most $\frac{2}{k_i}$, except possibly for vertices of $L_{i-1}^{<k}$, which could each be charged 1 extra. Since at least 1 unit was paid for each vertex of L_i, we obtain $f_i(n) \leq n + \frac{2}{k_i} f^{<k_i}(n) + f_{i-1}^{<k_i}(n)$.

By substituting for $f^{<k_i}$ and $f_{i-1}^{<k_i}$ the bound from the lemma, we arrive at the system of inequalities $f_i(n) \leq n + O\big(k_i f(\lfloor n/k_i \rfloor) + k_i^2 f_{i-1}(\lfloor n/k_i \rfloor)\big)$, $i = 1, 2, \ldots, s$ (where we put $f_0 = 0$), and we also have $f \leq f_1 + \cdots + f_s$. If one sets $k_i = n^{\varepsilon_i}$ with $0 < \varepsilon_1 \ll \varepsilon_2 \ll \cdots \ll \varepsilon_s \ll \varepsilon$, a not too difficult calculation shows that $f(n) = O(n^{1+\varepsilon})$ as claimed. □

Bounding the maximum complexity of a single cell is usually considerably more demanding than the lower envelope question, mainly because a cell can have a complicated topology (cells in hyperplane arrangements, no more complicated than the lower envelope, are a honorable exception). In the plane, these obstacles are not too formidable, and by a reduction to DS-sequences, it can be shown that the single-cell complexity for segments is $O(n\alpha(n))$, and for pieces of algebraic curves it can be bounded by some $\lambda_s(n)$, with s depending on the maximum degree of the curves. In $\mathbf{R}^3$, a general near-tight bound of $O(n^{2+\varepsilon})$ was proved in [21]. Some more special results are known for all d, such as an $O(n^{d-1} \log n)$ bound for a single cell in an arrangement of $(d-1)$-dimensional simplices in $\mathbf{R}^d$ [6]. Very recently, Basu proved, in an unpublished manuscript, that the sum of the Betti numbers (i.e. "topological complexity") of a single cell in an arrangement of surface patches in $\mathbf{R}^d$ is $O(n^{d-1})$. This might be helpful in getting good bounds on the combinatorial complexity too.

Concerning the union of "fat" objects (Fig. 1 bottom right), let us consider n convex sets in the plane, and let us ask what is the combinatorial complexity of the complement of their union. To get a meaningful problem, we assume that the boundaries of any two sets intersect in at most s points for some fixed $s \geq 4$ ($s = 2$ is easy). Long and skinny sets can form a grid pattern and have union complexity about n^2, but if we also require that the sets be "fat" (the ratio of the circumradius and inradius is bounded by some constant K), then a recent result of Efrat and Sharir [18] shows that the union complexity is near-linear, at most $O(n^{1+\varepsilon})$, with the constant of proportionality depending on s, K, ε ([26] gives a simpler and more precise bound for fat triangles). Various extensions to non-convex cases or to higher dimensions seem easy to conjecture but quite hard to prove.

There are still many open problems in the above-discussed areas, but what seems to be needed most at the moment is a simplification and streamlining, since building up on the existing proofs is getting more and more cumbersome.

Here is an annoying open problem concerning arrangements of n algebraic surfaces in $\mathbf{R}^d$. If the degrees of the surfaces are bounded, the complexity of the arrangement is $O(n^d)$. But the cells can be combinatorially very complicated, while for many applications, one needs to work with cells definable by constant-size formulas, the so-called *Tarski cells* (curved analogues of simplices, so to speak). Can each of the cells of the arrangement be subdivided into Tarski cells, in such a way that altogether $O(n^d)$ Tarski cells result? The best known upper bound for $d \geq 3$ is a bit larger than $O(n^{2d-3})$ [11].

Multiple cells, incidences, cuttings

Besides a single cell, also the total complexity of several cells in an arrangement has been studied, and this has interesting connections to some old combinatorial-geometric problems. Let us consider some m 2-cells in a planar arrangement of n lines (call them *marked cells*), and let us denote the maximum possible total number of vertices of these cells by $K(n,m)$. While $K(n,1) = n$, $K(n,m)$ is considerably smaller than mn for large m.

To get a nontrivial upper bound on $K(n,m)$, we define a bipartite graph with the lines and the marked cells as vertices and with edges connecting each cell to the lines forming its sides. There cannot be 5 lines simultaneously connected to the same two cells, and the Kővári-Sós-Turán theorem in extremal graph theory implies that there are $O(m\sqrt{n} + n)$ edges; thus $K(n,m) = O(m\sqrt{n} + n)$. In particular, $K(n, \sqrt{n}) = O(n)$, (this is a result of Canham from 1969), which is obviously tight. But the bound is not tight for $n = m$, say, and the right bound is $K(n,m) = O(n^{2/3}m^{2/3} + n + m)$. This was proved by Clarkson et al. [14], using a general technique that emerged in previous work on geometric algorithms. We give the proof for $m = n$. The basic idea is this: since the bound we already have is good if there are many more lines than points, we subdivide the problem with n lines and n points into smaller subproblems, most of them with many more lines than points. The device for this subdivision is the so-called $\frac{1}{r}$-cutting.

For a parameter $r \geq 1$ and a set L of n lines in the plane, a $\frac{1}{r}$-*cutting* for L is a finite set of triangles[4] with disjoint interiors covering the plane, such that

[4]Where unbounded triangles are admitted too, i.e. a triangle means an intersection of 3 halfplanes here.

the interior of each triangle is intersected by no more than $\frac{n}{r}$ lines of L. A basic existence result says that for any L and r, a $\frac{1}{r}$-cutting exists consisting of $O(r^2)$ triangles (note that the bound is independent of n). Three proofs are known: a very elementary one [24], and two probabilistic ones which generalize to higher dimensions [12, 10].

For bounding $K(n,n)$, let L be the n considered lines, set $r = n^{1/3}$, and consider a $\frac{1}{r}$-cutting $\{\Delta_1, \ldots, \Delta_q\}$ for L, $q = O(r^2)$. Let $L_i \subset L$ be the set of lines intersecting the interior of Δ_i and suppose that there are m_i marked cells completely contained in Δ_i. The total complexity of these marked cells, over all Δ_i, is at most $\sum_{i=1}^{q} K(|L_i|, m_i) \leq \sum_{i=1}^{q} O(m_i\sqrt{n/r} + \frac{n}{r}) = O(n^{3/2}r^{-1/2} + nr)$, using the above-derived bound for $K(n,m)$ and $\sum m_i \leq n$. It remains to account for the marked cells intersecting boundaries of some of the Δ_i's. But each vertex of such a marked cell lies in the zone of a side of some Δ_i in the arrangement of L_i, and the total complexity of these zones is at most $3\sum_{i=1}^{q} O(|L_i|) = O(nr)$. Altogether we get $K(n,n) = O(n^{4/3})$. $\Box$

An easy consequence of the bound $K(n,m) = O(n^{2/3}m^{2/3} + m + n)$ is the same (and also tight) bound for the maximum number of incidences between n lines and m points in the plane. This bound for incidences was proved earlier by Szemerédi and Trotter, and the new proof via $\frac{1}{r}$-cuttings [14] was a considerable simplification. A still much simpler proof was found later by Székely [34] via geometric graphs, but so far his technique seems mainly applicable for problems in the plane, while with $\frac{1}{r}$-cuttings, various higher-dimensional problems can be handled too (see, e.g., [14, 29] or a survey by Agarwal and Sharir in [31] for more results and references).

The perhaps most challenging related problem is Erdős' question on unit distances: given n points in the plane, what is the maximum possible number of pairs of points at distance 1? By drawing a unit circle around each point, the question can be reduced to the maximum number of incidences between n points and n unit circles. Both Székely's technique and the one with $\frac{1}{r}$-cuttings yield the same $O(n^{4/3})$ bound as for line-point incidences, but while for lines this is tight, the best known lower bound for unit circles is only slightly superlinear. To decrease the upper bound for the unit-distance problem, a radically new approach seems to be needed, because the $n^{4/3}$ bound is tight for *pseudocircles*, i.e. collections of Jordan curves that combinatorially behave "like unit circles", and none of the known methods can take advantage of "true circularity" of the unit circles.

In this connection, a recent result of Elekés and Rónyai [19] should be mentioned. They characterized bivariate polynomials and rational functions that attain only $O(n)$ distinct values on $X \times Y$ for some n-element sets $X, Y \subset \mathbf{R}$. As a special case, they settled a conjecture of Purdy: if u and v are lines and $P \subset u$ and $Q \subset v$ are n-point sets such that the distance $|p-q|$ attains only $O(n)$ distinct values for $p \in P$ and $q \in Q$, then u and v must be parallel or perpendicular (provided n is large enough). The proof is in part algebraic and it strongly uses the "straightness" of the lines u and v.

Range searching, partitions, Heilbronn's problem

Let us consider the following algorithmic problem. Given an n-point set $P \subset \mathbf{R}^2$, we want to build some data structure for storing information about P, in such a way that if we get a stripe σ (bounded by two parallel lines) as a query, the number of points of P lying in σ can be determined quickly, hopefully much faster than by examining all points of P. Moreover, we insist that the space occupied by the data structure is at most proportional to n.

Questions of this type, the so-called *range searching problems*, have been studied quite intensively and in a much more general form—in higher dimensions, with different query shapes, with more space allowed, etc. (there is a survey by Agarwal in [20], and another survey is [25]). But many interesting aspects can be demonstrated on the particular problem formulated above. In this case, it is possible to answer the query in $O(\sqrt{n})$ time, and with some restriction on the type of algorithm used, this is asymptotically optimal. Ironically, while the known data structures for this problem are not very useful in practice, the underlying theory involves some of the nicest mathematics in computational geometry.

At first sight (and probably at many subsequent sights too), it is not clear how to achieve any sublinear query time. Willard discovered in 1981 that the following type of geometric construction can be used: given the point set P, partition the plane into some number r of regions, each containing roughly $\frac{n}{r}$ points of P, in such a way that no line intersects more than κ of these regions, where κ should be considerably smaller than r. How can this help with a query? We store the number of points in each of the regions. Given a query stripe σ, the boundary of σ intersects at most 2κ regions. These must be further examined, but each of the other regions can be processed in unit time using the stored point counts. The actual algorithms are more complicated but this is the basic idea.

Finding an optimal construction of such a partition took a long time. (Looking for good partitions stimulated, for instance, research in equipartitioning masses by hyperplanes—see e.g. [30]—although other approaches were used in the subsequent development.) One of the most important steps was the following result, essentially invented by Welzl, with a slight improvement in [13]: any $2n$-point set in the plane can be divided into pairs of points in such a way that any line crosses only $O(\sqrt{n})$ of the segments connecting the pairs. One almost wouldn't believe that after thousands of years of geometry, it is still possible to discover such pretty theorems about points in the plane. This was later generalized to a partition of an n-point set into r parts of size roughly $\frac{n}{r}$, with any line crossing $O(\sqrt{r})$ parts only (see [25]). Both these results are asymptotically optimal. The research in range searching also initiated a fruitful theory related to the so-called Vapnik-Chervonenkis dimension of set systems, with applications, e.g., in discrepancy theory; this is surveyed in [27].

Lower bounds for range searching were proved mainly by Chazelle; a key paper is [9]. In the proof, some integral-geometric considerations appear, and, interestingly, the lower bounds are related to a generalization of Heilbronn's problem from discrete geometry. For an n-point set $P \subset [0,1]^2$ and $3 \leq k \leq n$, let $a_k(P)$ denote the minimum area of the convex hull of a k-point subset of P. Heilbronn's problem asks for determining $a_3(P)$, and although the answer is unknown, it is known that $a_3(P)$ is of much smaller order than $\frac{1}{n}$ (which is what one might perhaps expect at first). In Chazelle's proof, one needs a set P with $a_k(P) = \Omega(\frac{k}{n})$ for all $k \in [k_0, n]$, with k_0 as small as possible. He achieves this with $k_0 \approx \log n$, and this causes the presence of an $\log n$ factor in the range-searching lower bound in $\mathbf{R}^3$ which probably shouldn't be there. From Heilbronn's problem, we know that $k_0 = 3$ is impossible to reach, but perhaps it might be possible to decrease k_0 to something smaller than $\log n$, which would improve the range-searching bound. For a more recent progress in range-searching lower bounds, and some nice geometric problems, see [8].

Many other areas and results would deserve to be mentioned, such as the developments related to linear programming algorithms (see the survey [1]) which also led to a nice purely mathematical application by Amenta [3] (a short proof of a Helly-type result), or the story of *weak ε-nets*, born in computational geometry and later used by Alon and Kleitman [2] in their solution of the long-open Hadwiger-Debrunner problem in convex geometry, or an interesting question of algebraic-topological nature arising in motion planning of multiple robots [4]. But it's really time to finish.

ACKNOWLEDGMENT. While working in computational geometry, I met many very nice people, and a number people helped me in various ways. I would like to use this opportunity to thank them all—I will not try to list names, in order not to omit anyone. I also thank Pankaj K. Agarwal, Eva Matoušková, Micha Sharir, Ricky Pollack, and Pavel Valtr for constructive comments on a draft of the present paper.

REFERENCES

[1] P. K. Agarwal and M. Sharir. Algorithmic techniques for geometric optimization. In *Computer Science Today: Recent Trends and Developments*, volume 1000 of *Lecture Notes Comput. Sci.*, pages 234–253. Springer-Verlag, 1995.

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

[3] N. Amenta. A short proof of an interesting Helly-type theorem. *Discr. Comput. Geom.*, 15:423–427, 1996.

[4] B. Aronov, M. de Berg, F. van der Stappen, and P. Švestka. Motion planning for multiple robots. In *Proc. 14th Annu. ACM Sympos. Comput. Geom.*, 1998.

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

[6] B. Aronov and M. Sharir. Castles in the air revisited. *Discrete Comput. Geom.*, 12:119–150, 1994.

[7] S. Basu, R. Pollack, and M.-F. Roy. On the number of cells defined by a family of polynomials on a variety. *Mathematika*, 43:120–126, 1996.

[8] H. Brönnimann, B. Chazelle, and J. Pach. How hard is halfspace range searching. *Discrete Comput. Geom.*, 10:143–155, 1993.

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

[10] B. Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(2):145–158, 1993.

[11] B. Chazelle, H. Edelsbrunner, L. Guibas, and M. Sharir. A singly-exponential stratification scheme for real semi-algebraic varieties and its applications. In *Proc. 16th Internat. Colloq. Automata Lang. Program.*, volume 372 of *Lecture Notes Comput. Sci.*, pages 179–192. Springer-Verlag, 1989.

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

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

[14] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir, and E. Welzl. Combinatorial complexity bounds for arrangements of curves and spheres. *Discrete Comput. Geom.*, 5:99–160, 1990.

[15] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

[16] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, 1997.

[17] H. Edelsbrunner, R. Seidel, and M. Sharir. On the zone theorem for hyperplane arrangements. *SIAM J. Comput.*, 22(2):418–429, 1993.

[18] A. Efrat and M. Sharir. On the complexity of the union of fat objects in the plane. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 104–112, 1997.

[19] Gy. Elekes and L. Rónyai. A combinatorial problem on polynomials and rational functions. *J. Comb. Thoery Ser. B*, 1998. To appear.

[20] J. E. Goodman and J. O'Rourke, editors. *Handbook of Discrete and Computational Geometry*. CRC Press LLC, Boca Raton, FL, 1997.

[21] D. Halperin and M. Sharir. Almost tight upper bounds for the single cell and zone problems in three dimensions. *Discrete Comput. Geom.*, 14:385–410, 1995.

[22] M. Klazar. A general upper bound in extremal theory of sequences. *Comment. Math. Univ. Carol.*, 33:737–746, 1992.

[23] M. Klazar and P. Valtr. Generalized Davenport–Schinzel sequences. *Combinatorica*, 14:463–476, 1994.

[24] J. Matoušek. Construction of ϵ-nets. *Discrete Comput. Geom.*, 5:427–448, 1990.

[25] J. Matoušek. Geometric range searching. *ACM Comput. Surv.*, 26:421–461, 1994.

[26] J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl. Fat triangles determine linearly many holes. *SIAM J. Comput.*, 23:154–169, 1994.

[27] J. Matoušek. Geometric set systems. In *Proceedings of the 2nd European Congress of Mathematicians.* Birkhäuser, Basel, 1998. In press.

[28] P. Orlik and H. Terao. *Arrangements of hyperplanes.* Springer-Verlag, Berlin etc., 1991.

[29] J. Pach and P. K. Agarwal. *Combinatorial Geometry.* John Wiley & Sons, New York, NY, 1995.

[30] E. A. Ramos. Equipartition of mass distributions by hyperplanes. *Discrete Comput. Geom.*, 15:147–167, 1996.

[31] J.-R. Sack and J. Urrutia, editors. *Handbook on Computational Geometry.* North-Holland, 1998. To appear.

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

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

[34] L. Székely. Crossing numbers and hard Erdős problems in discrete geometry. *Combinatorics, Probability, and Computing*, 6:353–358, 1997.

[35] P. Valtr. On geometric graphs with no k pairwise parallel edges. *Discrete Comput. Geom.*, 19:461–469, 1998.

Jiří Matoušek
Dept. of Applied Mathematics
Charles University
Malostranské nám. 25
118 00 Praha 1, Czech Republic