

Chapter 12

Davenport-Schinzel Sequences

The complexity of a simple arrangement of n lines in \mathbb{R}^2 is $\Theta(n^2)$ and so every algorithm that uses such an arrangement explicitly needs $\Omega(n^2)$ time. However, there are many scenarios in which we do not need the whole arrangement but only some part of it. For instance, to construct a ham sandwich cut for two sets of points in \mathbb{R}^2 one needs the median levels of the two corresponding line arrangements only. As mentioned in the previous section, the relevant information about these levels can actually be obtained in linear time. Similarly, in a motion planning problem where the lines are considered as obstacles we are only interested in the cell of the arrangement we are located in. There is no way to ever reach any other cell, anyway.

This chapter is concerned with analyzing the complexity—that is, the number of vertices and edges—of a single cell in an arrangement of n curves in \mathbb{R}^2 . In case of a line arrangement this is mildly interesting only: Every cell is convex and any line can appear at most once along the cell boundary. On the other hand, it is easy to construct an example in which there is a cell C such that every line appears on the boundary ∂C .

But when we consider arrangement of line segments rather than lines, the situation changes in a surprising way. Certainly a single segment can appear several times along the boundary of a cell, see the example in Figure 12.1. Make a guess: What is the maximal complexity of a cell in an arrangement of n line segments in \mathbb{R}^2 ?




Figure 12.1: A single cell in an arrangement of line segments.

You will find out the correct answer soon, although we will not prove it here. But my guess would be that it is rather unlikely that your guess is correct, unless, of course, you knew the answer already. :-)

For a start we will focus on one particular cell of any arrangement that is very easy to describe: the lower envelope or, intuitively, everything that can be seen vertically from below. To analyze the complexity of lower envelopes we use a combinatorial description using strings with forbidden subsequences, so-called Davenport-Schinzel sequences. These sequences are of independent interest, as they appear in a number of combinatorial problems [2] and in the analysis of data structures [7]. The techniques used apply not only to lower envelopes but also to arbitrary cells of arrangements.

Definition 12.1 *An (n, s) -Davenport-Schinzel sequence, for $n, s \in \mathbb{N}$, is a sequence over an alphabet A of size n in which*

- *no two consecutive characters are the same and*
- *there is no alternating subsequence of the form $\dots a \dots b \dots a \dots b \dots$ of $s + 2$ characters, for any $a, b \in A$.*

Let $\lambda_s(n)$ be the length of a longest (n, s) -Davenport-Schinzel sequence.

For example, $abcbacb$ is a $(3, 4)$ -DS sequence but not a $(3, 3)$ -DS sequence because it contains the subsequence $bcbcb$.

Proposition 12.2 $\lambda_s(m) + \lambda_s(n) \leq \lambda_s(m + n)$.

Proof. On the left hand side, we consider two Davenport-Schinzel sequences, one over an alphabet A of size m and another over an alphabet B of size n . We may suppose that $A \cap B = \emptyset$ (for each character $x \in A \cap B$ introduce a new character x' and replace all occurrences of x in the second sequence by x'). Concatenating both sequences yields a Davenport-Schinzel sequence over the alphabet $A \cup B$ of size $m + n$. \square

Let us now see how Davenport-Schinzel sequences are connected to lower envelopes. Consider a collection $\mathcal{F} = \{f_1, \dots, f_n\}$ of real-valued continuous functions that are defined on a common interval $I \subset \mathbb{R}$. The *lower envelope* $\mathcal{L}_{\mathcal{F}}$ of \mathcal{F} is defined as the pointwise minimum of the functions f_i , $1 \leq i \leq n$, over I . Suppose that any pair f_i, f_j , $1 \leq i < j \leq n$, intersects in at most s points. Then I can be decomposed into a finite sequence I_1, \dots, I_ℓ of (maximal connected) pieces on each of which a single function from \mathcal{F} defines $\mathcal{L}_{\mathcal{F}}$. Define the sequence $\phi(\mathcal{F}) = (\phi_1, \dots, \phi_\ell)$, where f_{ϕ_i} is the function from \mathcal{F} which defines $\mathcal{L}_{\mathcal{F}}$ on I_i .


Observation 12.3 $\phi(\mathcal{F})$ is an (n, s) -Davenport-Schinzel sequence.

In the case of line segments the above statement does not hold because a set of line segments is in general not defined on a common real interval.

Proposition 12.4 *Let \mathcal{F} be a collection of n real-valued continuous functions each of which is defined on some real interval. If any two functions from \mathcal{F} intersect in at most s points then $\phi(\mathcal{F})$ is an $(n, s + 2)$ -Davenport-Schinzel sequence.*

Proof. Let I denote the union of all intervals on which one of the functions from \mathcal{F} is defined. Consider any function $f \in \mathcal{F}$ defined on $[a, b] \subseteq I = [c, d]$. Extend f on I by extending it using almost vertical rays pointing upward, from a use a ray of sufficiently small slope, from b use a ray of sufficiently large slope. For all functions use the same slope on these two extensions such that no extensions in the same direction intersect. By *sufficiently small/large* we mean that for any extension ray there is no function endpoint nor an intersection point of two functions in the open angular wedge bounded by the extension ray and the vertical ray starting from the same source. (There may be such points *on* the vertical ray, but not in the open wedge between the two rays.)

Denote the resulting collection of functions totally defined on I by \mathcal{F}' . If the rays are sufficiently close to vertical then $\phi(\mathcal{F}') = \phi(\mathcal{F})$.



For any $f \in \mathcal{F}'$ a single extension ray can create at most one additional intersection with any $g \in \mathcal{F}'$. (Let $[a_f, b_f]$ and $[a_g, b_g]$ be the intervals on which the function f and g , respectively, was defined originally. Consider the ray r extending f from a_f to the left. If $a_f \in [a_g, b_g]$ then r may create a new intersection with g , if $a_f > b_g$ then r creates a new intersection with the right extension of g from b_g , and if $a_f < a_g$ then r does not create any new intersection with g .)

On the other hand, for any pair s, t of segments, neither the left extension of the leftmost segment endpoint nor the right extension of the rightmost segment endpoint can introduce an additional intersection. Therefore, any pair of segments in \mathcal{F}' intersects at most $s + 2$ times and the claim follows. \square

Next we will give an upper bound on the length of Davenport-Schinzel sequences for small s .

Lemma 12.5 $\lambda_1(n) = n$, $\lambda_2(n) = 2n - 1$, and $\lambda_3(n) \leq 2n(1 + \log n)$.

Proof. $\lambda_1(n) = n$ is obvious. $\lambda_2(n) = 2n - 1$ is given as an exercise. We prove $\lambda_3(n) \leq 2n(1 + \log n) = O(n \log n)$.

For $n = 1$ it is $\lambda_3(1) = 1 \leq 2$. For $n > 1$ consider any $(n, 3)$ -DS sequence σ of length $\lambda_3(n)$. Let a be a character that appears least frequently in σ . Clearly a appears at most $\lambda_3(n)/n$ times in σ . Delete all appearances of a from σ to obtain a sequence σ' on $n-1$ symbols. But σ' is not necessarily a DS sequence because there may be consecutive appearances of a character b in σ' , in case that $\sigma = \dots bab \dots$.

Claim: There are at most two pairs of consecutive appearances of the same character in σ' . Indeed, such a pair can be created around the first and last appearance of a in σ only. If any intermediate appearance of a creates a pair bb in σ' then $\sigma = \dots a \dots bab \dots a \dots$, in contradiction to σ being an $(n, 3)$ -DS sequence.

Therefore, one can remove at most two characters from σ' to obtain a $(n-1, 3)$ -DS sequence $\tilde{\sigma}$. As the length of $\tilde{\sigma}$ is bounded by $\lambda_3(n-1)$, we obtain $\lambda_3(n) \leq \lambda_3(n-1) + \lambda_3(n)/n + 2$. Reformulating yields

$$\underbrace{\frac{\lambda_3(n)}{n}}_{=: f(n)} \leq \underbrace{\frac{\lambda_3(n-1)}{n-1}}_{=: f(n-1)} + \frac{2}{n-1} \leq \underbrace{1}_{=: f(1)} + 2 \sum_{i=1}^{n-1} \frac{1}{i} = 1 + 2H_{n-1}$$

and together with $2H_{n-1} < 1 + 2 \log n$ we obtain $\lambda_3(n) \leq 2n(1 + \log n)$. \square

Bounds for higher-order Davenport-Schinzel sequences. As we have seen, $\lambda_1(n)$ (no aba) and $\lambda_2(n)$ (no $abab$) are both linear in n . It turns out that for $s \geq 3$, $\lambda_s(n)$ is slightly superlinear in n (taking s fixed). The bounds are known almost exactly, and they involve the inverse Ackermann function $\alpha(n)$, a function that grows extremely slowly.

To define the inverse Ackermann function, we first define a hierarchy of functions $\alpha_1(n)$, $\alpha_2(n)$, $\alpha_3(n)$, \dots where, for every fixed k , $\alpha_k(n)$ grows much more slowly than $\alpha_{k-1}(n)$:

We first let $\alpha_1(n) = \lceil n/2 \rceil$. Then, for each $k \geq 2$, we define $\alpha_k(n)$ to be the number of times we must apply α_{k-1} , starting from n , until we get a result not larger than 1. In other words, $\alpha_k(n)$ is defined recursively by:

$$\alpha_k(n) = \begin{cases} 0, & \text{if } n \leq 1; \\ 1 + \alpha_k(\alpha_{k-1}(n)), & \text{otherwise.} \end{cases}$$

Thus, $\alpha_2(n) = \lceil \log_2 n \rceil$, and $\alpha_3(n) = \log^* n$.

Now fix n , and consider the sequence $\alpha_1(n)$, $\alpha_2(n)$, $\alpha_3(n)$, \dots . For every fixed n , this sequence decreases rapidly until it settles at 3. We define $\alpha(n)$ (the inverse Ackermann function) as the function that, given n , returns the smallest k such that $\alpha_k(n)$ is at most 3:

$$\alpha(n) = \min\{k \mid \alpha_k(n) \leq 3\}.$$

We leave as an exercise to show that for every fixed k we have $\alpha_k(n) = o(\alpha_{k-1}(n))$ and $\alpha(n) = o(\alpha_k(n))$.

Coming back to the bounds for Davenport-Schinzel sequences, for $\lambda_3(n)$ (no ababa) it is known that $\lambda_3(n) = \Theta(n\alpha(n))$ [4]. In fact it is known that $\lambda_3(n) = 2n\alpha(n) \pm O(n\sqrt{\alpha(n)})$ [5, 6]. For $\lambda_4(n)$ (no ababab) we have $\lambda_4(n) = \Theta(n \cdot 2^{\alpha(n)})$ [3].

For higher-order sequences the known upper and lower bounds are almost tight, and they are of the form $\lambda_s(n) = n \cdot 2^{\text{poly}(\alpha(n))}$, where the degree of the polynomial in the exponent is roughly $s/2$ [3, 6].

Realizing DS sequences as lower envelopes. There exists a construction of a set of n segments in the plane whose lower-envelope sequence has length $\Omega(n\alpha(n))$. (In fact, the lower-envelope sequence has length $n\alpha(n) - O(n)$, with a leading coefficient of 1; it is an open problem to get a leading coefficient of 2, or prove that this is not possible.)

It is an open problem to construct a set of n parabolic arcs in the plane whose lower-envelope sequence has length $\Omega(n \cdot 2^{\alpha(n)})$.

Generalizations of DS sequences. Also generalizations of Davenport-Schinzel sequences have been studied, for instance, where arbitrary subsequences (not necessarily an alternating pattern) are forbidden. For a word σ and $n \in \mathbb{N}$ define $\text{Ex}(\sigma, n)$ to be the maximum length of a word over $A = \{1, \dots, n\}^*$ that does not contain a subsequence of the form σ . For example, $\text{Ex}(\text{ababa}, n) = \lambda_3(n)$. If σ consists of two letters only, say a and b , then $\text{Ex}(\sigma, n)$ is super-linear if and only if σ contains ababa as a subsequence [1]. This highlights that the alternating forbidden pattern is of particular interest.

Exercise 12.6 Prove that $\lambda_2(n) = 2n - 1$.

Exercise 12.7 Prove that $\lambda_s(n)$ is finite for all s and n .

12.1 Constructing lower envelopes

Theorem 12.8 Let $\mathcal{F} = \{f_1, \dots, f_n\}$ be a collection of real-valued continuous functions defined on a common interval $I \subset \mathbb{R}$ such that no two functions from \mathcal{F} intersect in more than s points. Then the lower envelope $\mathcal{L}_{\mathcal{F}}$ can be constructed in $O(\lambda_s(n) \log n)$ time. (Assuming that an intersection between any two functions can be constructed in constant time.)

Proof. Divide and conquer. For simplicity, assume that n is a power of two. Split \mathcal{F} into two equal parts \mathcal{F}_1 and \mathcal{F}_2 and construct $\mathcal{L}_{\mathcal{F}_1}$ and $\mathcal{L}_{\mathcal{F}_2}$ recursively. The resulting envelopes can be merged using line sweep by processing $2\lambda_s(n/2) + \lambda_s(n) \leq 2\lambda_s(n)$ events (the inequality $2\lambda_s(n/2) \leq \lambda_s(n)$ is by Proposition 12.2). Here the first term accounts for events generated by the vertices of the two envelopes to be merged. The second term accounts for their intersections, each of which generates a vertex of the resulting envelope. Observe that no sorting is required and the SLS structure is of constant size. Therefore, the sweep can be done in time linear in the number of events.

This yields the following recursion for the runtime $T(n)$ of the algorithm. $T(n) \leq 2T(n/2) + c\lambda_s(n)$, for some constant $c \in \mathbb{N}$. Using Proposition 12.2 it follows that $T(n) \leq c \sum_{i=1}^{\log n} 2^i \lambda_s(n/2^i) \leq c \sum_{i=1}^{\log n} \lambda_s(n) = O(\lambda_s(n) \log n)$. \square

Exercise 12.9 Show that every (n, s) -Davenport-Schinzel sequence can be realized as the lower envelope of n continuous functions from \mathbb{R} to \mathbb{R} , every pair of which intersect at most s times.

Exercise 12.10 Show that every Davenport-Schinzel sequence of order two can be realized as a lower envelope of n parabolas.

It is a direct consequence of the symmetry in the definition that the property of being a Davenport-Schinzel sequence is invariant under permutations of the alphabet. For instance, $\sigma = bcacba$ is a $(3, 3)$ -DS sequence over $A = \{a, b, c\}$. Hence the permutation $\pi = (ab)$ induces a $(3, 3)$ -DS sequence $\pi(\sigma) = acbcab$ and similarly $\pi' = (cba)$ induces another $(3, 3)$ -DS sequence $\pi'(\sigma) = abcbac$.

When counting the number of Davenport-Schinzel sequences of a certain type we want to count *essentially distinct* sequences only. Therefore we call two sequences over a common alphabet A *equivalent* if and only if one can be obtained from the other by a permutation of A . Then two sequences are *distinct* if and only if they are not equivalent. A typical way to select a representative from each equivalence class is to order the alphabet and demand that the first appearance of a symbol in the sequence follows that order. For example, ordering $A = \{a, b, c\}$ alphabetically demands that the first occurrence of a precedes the first occurrence of b , which in turn precedes the first occurrence of c .

Exercise 12.11 Let P be a convex polygon with $n+1$ vertices. Find a bijection between the triangulations of P and the set of pairwise distinct $(n, 2)$ -Davenport-Schinzel sequences of maximum length $(2n-1)$. It follows that the number of distinct maximum $(n, 2)$ -Davenport-Schinzel sequences is exactly $C_{n-1} = \frac{1}{n} \binom{2n-2}{n-1}$, which is the $(n-1)$ -st Catalan number.

Questions

58. What is an (n, s) Davenport-Schinzel sequence and how does it relate to the lower envelope of real-valued continuous functions? Give the precise definition and some examples. Explain the relationship to lower envelopes and how to apply the machinery to partial functions like line segments.
59. What is the value of $\lambda_1(n)$ and $\lambda_2(n)$?
60. What is the asymptotic value of $\lambda_3(n)$, $\lambda_4(n)$, and $\lambda_s(n)$ for larger s ?
61. What is the combinatorial complexity of the lower envelope of a set of n lines/parabolas/line segments?

References

- [1] Radek Adamec, Martin Klazar, and Pavel Valtr, Generalized Davenport-Schinzel sequences with linear upper bound. *Discrete Math.*, **108**, (1992), 219–229, URL [http://dx.doi.org/10.1016/0012-365X\(92\)90677-8](http://dx.doi.org/10.1016/0012-365X(92)90677-8).
- [2] Pankaj K. Agarwal and Micha Sharir, *Davenport-Schinzel sequences and their geometric applications*. Cambridge University Press, New York, NY, 1995.
- [3] Pankaj K. Agarwal, Micha Sharir, and Peter W. Shor, Sharp upper and lower bounds on the length of general Davenport-Schinzel sequences. *J. Combin. Theory Ser. A*, **52**, 2, (1989), 228–274, URL [http://dx.doi.org/10.1016/0097-3165\(89\)90032-0](http://dx.doi.org/10.1016/0097-3165(89)90032-0).
- [4] Sergiu Hart and Micha Sharir, Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes. *Combinatorica*, **6**, (1986), 151–177, URL <http://dx.doi.org/10.1007/BF02579170>.
- [5] Martin Klazar, On the maximum lengths of Davenport-Schinzel sequences. In R. Graham et al., ed., *Contemporary Trends in Discrete Mathematics*, vol. 49 of *DISCRETE MATHEMATICS AND THEORETICAL COMPUTER SCIENCE*, pp. 169–178, Amer. Math. Soc., Providence, RI, 1999.
- [6] Gabriel Nivasch, Improved bounds and new techniques for Davenport-Schinzel sequences and their generalizations. *J. ACM*, **57**, 3, (2010), Article No. 17, URL <http://dx.doi.org/10.1145/1706591.1706597>.
- [7] Seth Pettie, Splay trees, Davenport-Schinzel sequences, and the deque conjecture. In *Proc. 19th ACM-SIAM Sympos. Discrete Algorithms*, pp. 1115–1124, 2008, URL <http://doi.acm.org/10.1145/1347082.1347204>.