18. Davenport-Schinzel Sequences

Lecture on Monday 30th November, 2009 by Michael Hoffmann <hoffmann@inf.ethz.ch>

Definition 18.1 *A (n, s)-Davenport-Schinzel sequence 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 $\ldots a \ldots b \ldots a \ldots b \ldots$ 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 $bcbcb$.

Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a collection of real-valued continuous functions 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 $I_1, \ldots, I_\ell$ of (maximal connected) pieces on each of which a single function from $\mathcal{F}$ defines $\mathcal{L}_{\mathcal{F}}$. Define the sequence $\phi(\mathcal{F}) = (\phi_1, \ldots, \phi_\ell)$, where f_{ϕ_i} is the function from $\mathcal{F}$ which defines $\mathcal{L}_{\mathcal{F}}$ on I_i.

Observation 18.2 *$\phi(\mathcal{F})$ 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 18.3 *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.

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.

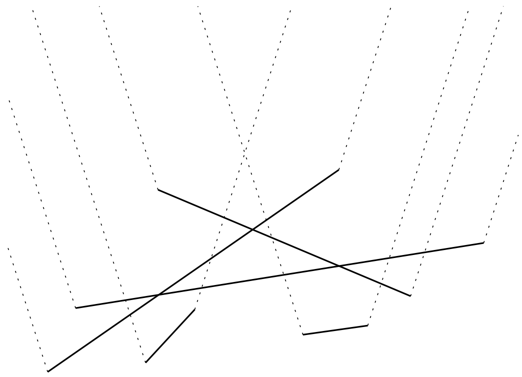

If $a_f \in [a_g, b_g]$ then r may create a new intersection with g, if $a_f > b_g$ then r creates a new intersection with the right extension of g from b_g, and if $a_f < a_g$ then r does not create any new intersection with g.)

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. □

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

Lemma 18.4 *$\lambda_1(n) = n$, $\lambda_2(n) = 2n - 1$, 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 which 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 $\sigma = \ldots bab \ldots$.

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 $\sigma = \ldots a \ldots bab \ldots a \ldots$, 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

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

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

Remarks. The upper bound is not tight. It can be shown [2] that $\lambda_3(n) = \Theta(n\alpha(n))$, where $\alpha(n)$ is the inverse Ackermann Function. More precisely, it is known that $2n\alpha(n) - O(n) \leq \lambda_3(n) \leq 2n\alpha(n) + O(n\sqrt{\alpha(n)})$, where the upper bound is due to Klazar [3] and the lower bound was recently shown by Nivasch [4].

The Ackermann function is defined on $\mathbb{N} \times \mathbb{N}$ as follows[1]. $A(1,n) = 2n$, $A(k,1) = 2$, for $k \geq 2$, and $A(k,n) = A(k-1, A(k, n-1))$, for $k, n \geq 2$. The inverse Ackermann Function $\alpha(n)$ is then given by $\alpha(n) = \min\{k \in \mathbb{N} \mid A(k,k) \geq n\}$. As $A(4,4)$ is a tower with 65536 2's, $\alpha(n) \leq 4$ for all practical purposes.

$\lambda_s(n)$ is almost linear even for larger values of s. For example, $\lambda_4(n) = \Theta(n2^{\alpha(n)})$.

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 $\mathrm{Ex}(\sigma, n)$ to be the maximum length of a word over $A = \{1, \ldots, n\}^*$ that does not contain a subsequence of the form σ. For example, $\mathrm{Ex}(ababa, n) = \lambda_3(n)$. If σ consists of two letters only, say a and b, then $\mathrm{Ex}(\sigma, n)$ 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.

18.1 Constructing lower envelopes

Theorem 18.5 *Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ 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 almost 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 $2\lambda_s(n/2) + \lambda_s(n) \leq 3\lambda_s(n)$ events. (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 $c \in \mathbb{N}$. Observe that $k\lambda_s(n/k) \leq \lambda_s(n)$, for $k \mid n$, because any k DS-sequences on an alphabet of size n/k can be concatenated to a single DS-sequence on an alphabet of size n by using pairwise disjoint (parts of the) alphabets for each of the k sequences. It follows that $T(n) \leq c \sum_{i=1}^{\log n} 2^i \lambda_s(n/2^i) = c \sum_{i=1}^{\log n} \lambda_s(n) = O(\lambda_s(n) \log n)$. □

[1] In fact, several different definitions are used in the literature, all of which exhibit the same asymptotic growth.

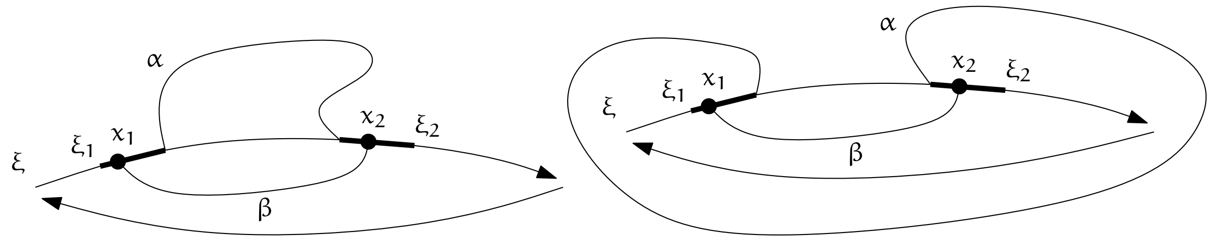

(a) f lies on the unbounded side of $\alpha \cup \beta$. (b) f lies on the bounded side of $\alpha \cup \beta$.

Figure 18.1: *Cases in the Consistency Lemma.*

18.2 Complexity of a single face

Theorem 18.6 *Let $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ be a collection of Jordan arcs in $\mathbb{R}^2$ such that each pair intersects in at most s points, for some $s \in \mathbb{N}$. Then the combinatorial complexity of any single face in the arrangement $\mathcal{A}(\Gamma)$ is $O(\lambda_{s+2}(n))$.*

Proof. Consider a face f of $\mathcal{A}(\Gamma)$. In general, the boundary of f might consist of several connected components. But as any single curve can appear in at most one component we may as well suppose that the boundary consists of one component only. (The complexity we are heading for is super-linear.)

Replace each γ_i by two directed arcs γ_i^+ and γ_i^- that together form a closed curve that is infinitesimally close to γ_i. Denote by S the circular sequence of these oriented curves, in their order along the (oriented) boundary ∂f of f.

Consistency Lemma. Let ξ be one of the oriented arcs γ_i^+ or γ_i^-. The order of portions of ξ that appear in S is consistent with their order along ξ. (That is, for each ξ we can break up the circular sequence S into a linear sequence $S(\xi)$ such that the portions of ξ that correspond to appearances of ξ in $S(\xi)$ appear in the same order along ξ.)

Consider two portions ξ_1 and ξ_2 of ξ that appear consecutively in S. Choose points $x_1 \in \xi_1$ and $x_2 \in \xi_2$ and connect them in two ways: first by the arc α following ∂f as in S, and second by an arc β inside the closed curve formed by γ_i^+ or γ_i^-. The curves α and β do not intersect except at their endpoints and they are both contained in the complement of the interior of f. In other words, $\alpha \cup \beta$ forms a closed Jordan curve and f lies either in the interior of this curve or in its exterior. In either case, the part of ξ between ξ_1 and ξ_2 is separated from f by $\alpha \cup \beta$ and, therefore, no point from this part can appear anywhere along ∂f. In other words, ξ_1 and ξ_2 are also consecutive boundary parts in the order of boundary portions along ξ, which proves the lemma.

Break up S into a linear sequence $S' = (s_1, \ldots, s_t)$ arbitrarily. For each oriented arc ξ, consider the sequence $s(\xi)$ of its portions along ∂f in the order in which they appear along ξ. By the Consistency Lemma, $s(\xi)$ corresponds to a subsequence of S, starting at s_k, for some $1 \leq k \leq t$. In order to consider $s(\xi)$ as a subsequence of S', break up the symbol for ξ into two symbols ξ and ξ' and replace all occurrences of ξ in S' before s_k

by ξ'. Doing so for all oriented arcs results in a sequence S^* on at most $4n$ symbols.

Claim: S^* is a $(4n, s+2)$-Davenport-Schinzel sequence.

Clearly no two adjacent symbols in S^* are the same. Suppose S^* contains an alternating subsequence $\xi \ldots \eta \ldots \xi \ldots \eta$ of length $s+4$. Consider any four consecutive elements of this subsequence. Choose points $x, y \in \xi$ and $z, w \in \eta$ such that they appear in the order x, z, y, w along ∂f. Connect x and y by a Jordan arc β_{xy} within the closed curve formed by ξ and its counterpart. Similarly, connect z and w by a Jordan arc β_{zw} within the closed curve formed by η and its counterpart. Then connect x, z, y, w along ∂f by curves β_{xz}, β_{zy}, β_{yw}, and β_{wx}. Observe that the last four curves are pairwise disjoint except for common endpoints. Moreover, none of them intersects β_{xy} or β_{zw}, except at a common endpoint.

We claim that β_{xy} and β_{zw} intersect. Suppose they do not. Then the six curves $\beta_{\cdot}$ form a plane graph on x, y, z, w which together with a point u chosen somewhere inside f and curves/edges connecting u to all of x, y, z, w within f form a plane embedding of K_5, contradiction.

In other words, any quadruple of consecutive elements from the alternating subsequence induces an intersection between the corresponding arcs ξ and η. Clearly these intersection points are pairwise distinct for any pair of distinct quadruples which altogether provides $s+4-3 = s+1$ points of intersection between ξ and η, in contradiction to the assumption that they intersect in at most s points. □

Corollary 18.7 *The combinatorial complexity of a single face in an arrangement of n line segments in $\mathbb{R}^2$ is $O(\lambda_3(n)) = O(n\alpha(n))$.*

Questions

46. *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.

47. *What is the order of magnitude of $\lambda_i(n)$, for $i \in \{1, 2, 3\}$?* Prove the bounds given in Lemma 18.4. You should also know the asymptotics of $\lambda_3(n)$, without proof.

48. *What is the combinatorial complexity of the lower envelope of a set of n lines/parabolas/line segments?*

49. *What is the combinatorial complexity of a single face in an arrangement of n line segments?* State the result and sketch the proof (Theorem 18.6).

References

[1] R. Adamec, M. Klazar, and P. Valtr, Generalized Davenport-Schinzel sequences with linear upper bound, *Discrete Math.* **108** (1992), 219–229.

[2] S. Hart and Micha Sharir, Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes, *Combinatorica* **6** (1986), 151–177.

[3] M. Klazar, On the maximum lengths of Davenport-Schinzel sequences, in: *Contemporary Trends in Discrete Mathematics* (R. Graham et al., ed.), volume 49 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, Amer. Math. Soc., Providence, RI, 1999, 169–178.

[4] Gabriel Nivasch, Improved bounds and new techniques for Davenport-Schinzel sequences and their generalizations, in: *Proc. 20th ACM-SIAM Sympos. Discrete Algorithms*, 2009, 1–10.