Sharp Bounds on Davenport-Schinzel Sequences of Every Order

SETH PETTIE, University of Michigan

One of the longest-standing open problems in computational geometry is bounding the complexity of the lower envelope of n univariate functions, each pair of which crosses at most s times, for some fixed s. This problem is known to be equivalent to bounding the length of an order-s Davenport-Schinzel sequence, namely, a sequence over an n-letter alphabet that avoids alternating subsequences of the form $a \cdots b \cdots a \cdots b \cdots$ with length $s+2$. These sequences were introduced by Davenport and Schinzel in 1965 to model a certain problem in differential equations and have since been applied to bound the running times of geometric algorithms, data structures, and the combinatorial complexity of geometric arrangements.

Let $\lambda_s(n)$ be the maximum length of an order-s DS sequence over n letters. What is λ_s asymptotically? This question has been answered satisfactorily [Hart and Sharir 1986; Agarwal et al. 1989; Klazar 1999; Nivasch 2010], when s is even or $s \leq 3$. However, since the work of Agarwal et al. in the mid-1980s, there has been a persistent gap in our understanding of the odd orders.

In this work, we effectively close the problem by establishing sharp bounds on Davenport-Schinzel sequences of every order s. Our results reveal that, contrary to one's intuition, $\lambda_s(n)$ behaves essentially like $\lambda_{s-1}(n)$ when s is odd. This refutes conjectures by Alon et al. [2008] and Nivasch [2010].

Categories and Subject Descriptors: G.2.1 [**Discrete Mathematics**]: Combinatorics

General Terms: Theory

Additional Key Words and Phrases: Combinatorics, sequences, forbidden patterns

ACM Reference Format:
Seth Pettie. 2015. Sharp bounds on davenport-Schinzel sequences of every order. J. ACM 62, 5, Article 36 (October 2015), 40 pages.
DOI: http://dx.doi.org/10.1145/2794075

1. INTRODUCTION

Consider the problem of bounding the complexity of the lower envelope of n continuous univariate functions $f_1, \ldots, f_n$, each pair of which cross at most s times. In other words, how many maximal connected intervals of the $\{f_i\}$ are contained in the graph of the function $f_{\min}(x) = \min\{f_1(x), \ldots, f_n(x)\}$? In the absence of any constraints on $\{f_i\}$, this problem can be completely stripped of its geometry by *transcribing* the lower envelope $f_{\min}$ as a *Davenport-Schinzel (DS) sequence* of order s, namely, a repetition-free sequence over the alphabet $\{1, \ldots, n\}$ that does not contain any alternating subsequences of the form $\cdots a \cdots b \cdots a \cdots b \cdots$ with length $s+2$, for any $a, b \in \{1, \ldots, n\}$.[1]

[1]If the sequence corresponding to the lower envelope contained an alternating subsequence $abab\cdots$ with length $s+2$, then the functions f_a and f_b must have crossed at least $s+1$ times, a contradiction.

This work is supported by NSF CAREER grant CCF-0746673, NSF grants CCF-1217338 and CNS-1318294, and a grant from the US-Israel Binational Science Foundation.
A preliminary version of this work appeared in *Proceedings of the 29th Annual Symposium on Computational Geometry (SoCG'13)*.
Author's address: S. Pettie, Department of Electrical Engineering and Computer Science, University of Michigan, 2260 Hayward St., Ann Arbor, MI, 48109; email: pettie@umich.edu.

DOI: http://dx.doi.org/10.1145/2794075

Although Davenport and Schinzel [1965a] introduced this problem nearly 50 years ago, it was only in the early 1980s that DS sequences became well known in the computational geometry community [Atallah 1985; Sharir et al. 1986]. Since then, DS sequences have found a startling number of geometric applications, with a growing number that are not overtly geometric [Alstrup et al. 1997; Burkard and Dollani 2001; Klawe 1992; López-de-los Mozos et al. 2013; Pettie 2008; Di Salvo and Proietti 2007].[2] In each of these applications, some quantity (e.g., running time, combinatorial complexity) is expressed in terms of $\lambda_s(n)$, the maximum length of an order-s DS sequence over an n-letter alphabet. To improve bounds on λ_s is, therefore, to improve our understanding of numerous problems in algorithms, data structures, and discrete geometry.

Davenport and Schinzel [1965a] established $n^{1+o(1)}$ upper bounds on $\lambda_s(n)$ for every order s. In order to properly survey the improvements that followed [Agarwal et al. 1989; Davenport 1971; Hart and Sharir 1986; Klazar 1999; Komjáth 1988; Nivasch 2010; Sharir 1987, 1988; Szemerédi 1974], we must define some notation for forbidden sequences and their extremal functions.

1.1. Sequence Notation and Terminology

We adopt and extend the sequence notation from the generalized Davenport-Schinzel literature (see, e.g., Klazar 2002; Pettie 2011b). Let $|\sigma|$ be the length of a sequence $\sigma = (\sigma(i))_{1\le i\le|\sigma|}$ and let $\|\sigma\|$ be the size of its alphabet $\Sigma(\sigma) = \{\sigma(i)\}$. Two equal length sequences are *isomorphic* if they are the same up to a renaming of their alphabets. We say σ is a *subsequence* of σ', written $\sigma \precsim \sigma'$, if σ can be obtained by deleting symbols from σ'. The predicate $\sigma \prec \sigma'$ asserts that σ is isomorphic to a subsequence of σ'. If $\sigma \not\prec \sigma'$ we say σ' is σ-*free*. If P is a *set* of sequences, $P \not\prec \sigma'$ holds if $\sigma \not\prec \sigma'$ for *every* $\sigma \in P$. The assertion that σ *appears in* or *occurs in* or *is contained in* σ' means either $\sigma \prec \sigma'$ or $\sigma \precsim \sigma'$, which one being clear from context. The *projection* of a sequence σ onto $G \subseteq \Sigma(\sigma)$ is obtained by deleting all non-G symbols from σ. A sequence σ is k-*sparse* if whenever $\sigma(i) = \sigma(j)$ and $i \ne j$, then $|i - j| \ge k$. A *block* is a sequence of distinct symbols. If σ is understood to be partitioned into a sequence of blocks, $[\![\sigma]\!]$ is the number of blocks. The predicate $[\![\sigma]\!] = m$ asserts that σ can be partitioned into at most m blocks. The extremal functions for *generalized* Davenport-Schinzel sequences are defined to be

$$\mathrm{Ex}(\sigma, n, m) = \max\{|S| \ : \ \sigma \not\prec S,\ \|S\| = n,\ \text{and } [\![S]\!] \le m\}$$
$$\mathrm{Ex}(\sigma, n) = \max\{|S| \ : \ \sigma \not\prec S,\ \|S\| = n,\ \text{and } S \text{ is } \|\sigma\|\text{-sparse}\},$$

[2]To cite a fraction of the literature, DS sequences/lower envelopes are routinely applied to problems related to geometric arrangements [Agarwal and Sharir 2002; Aronov and de Berg 2012; Aronov and Drusvyatskiy 2011; Bern et al. 1991; de Berg 2010; Dumitrescu et al. 2009; Edelsbrunner et al. 1987; Efrat 2005; Har-Peled 2000; Huttenlocher and Kedem 1990; Huttenlocher et al. 1993; Jaromczyk and Kowaluk 1988; Koltun 2004; Koltun and Sharir 2003; Matoušek et al. 1994; Pettie 2011c; Tagansky 1996], in kinetic data structures and dynamic geometric algorithms [Abam et al. 2010; Agarwal et al. 2008; Albers et al. 1998; Atallah 1985; Guibas et al. 1992; Huttenlocher et al. 1992; Kaplan et al. 2011; Wahid et al. 2010], in visibility [Cole and Sharir 1989; Moet et al. 2008; Sharir et al. 1986], motion planning [Leven and Sharir 1987; Sharir et al. 1986], and geometric containment problems [Aonuma et al. 1990; Augustine et al. 2010; Sharir et al. 1986; Sharir and Toledo 1994], as well as variations on classical problems such as computing shortest paths [Aronov et al. 2011; Bae and Okamoto 2012; Baltsan and Sharir 1988] and convex hulls [Aurenhammer and Jüttler 2012; Ezra and Mulzer 2012]. They have also been used in some industrial applications [Berretty et al. 2001; Ilushin et al. 2005]. Refer to Sharir and Agarwal [1995] for a survey of DS sequences and their early applications in computational geometry and to Klazar [2002] for a survey of DS sequences and related problems in extremal combinatorics.

where σ may be a single sequence or a set of sequences. The conditions $[\![S]\!] \le m$ and S is $\|\sigma\|$-sparse guarantee that the extremal functions are finite. For example, if $\|\sigma\| = 2$, the sparsity criterion forbids immediate repetitions and such infinite degenerate sequences as $aaaaa\cdots$. Blocked sequences, on the other hand, have no sparsity criterion. The extremal functions for (standard) Davenport-Schinzel sequences are defined to be

$$\lambda_s(n, m) = \mathrm{Ex}(\overbrace{ababa\cdots}^{\text{length } s+2}, n, m) \quad \text{and} \quad \lambda_s(n) = \mathrm{Ex}(\overbrace{ababa\cdots}^{\text{length } s+2}, n).$$

Bounds on generalized Davenport-Schinzel sequences are expressed as a function of the inverse-Ackermann function α, yet there is no universally agreed-upon definition of Ackermann's function or its inverse. All definitions in the literature differ by at most a constant, which usually obviates the need for more specificity. In this article, we use the following definition of Ackermann's function.

$$\begin{aligned} a_{1,j} &= 2^j && j \ge 1 \\ a_{i,1} &= 2 && i \ge 2 \\ a_{i,j} &= w \cdot a_{i-1,w} && i, j \ge 2 \\ &\quad \text{where } w = a_{i,j-1}. \end{aligned}$$

Note that in the table of $\{a_{i,j}\}$ values, the first column is constant ($a_{i,1} = 2$), and the second merely exponential ($a_{i,2} = 2^i$), so we have to look to the third column to find Ackermann-type growth. We define the double- and single-argument versions of the inverse-Ackermann function to be

$$\begin{aligned} \alpha(n, m) &= \min\{i \mid a_{i,j} \ge m, \ \text{ where } j = \max\{\lceil n/m \rceil, 3\}\} \\ \alpha(n) &= \alpha(n, n). \end{aligned}$$

We could have defined $\alpha(n, m)$ without direct reference to Ackermann's function. Note that $j = \log(a_{1,j})$. One may convince oneself that $j = \log^{\star}(a_{2,j}) - O(1)$, $j = \log^{\star\star}(a_{3,j}) - O(1)$, and in general, that $j = \log^{[i-1]}(a_{i,j}) - O(1)$, where $[i - 1]$ is short for $i - 1$ $\star$s.[3] Thus, up to $O(1)$ differences $\alpha(n, m)$ could be defined as $\min\{i \mid \log^{[i-1]}(m) \le \max\{\lceil n/m \rceil, 3\}\}$. We state previous results in terms of the single argument version of α. However, they all generalize to the two-argument version by replacing $\lambda_s(n)$ with $\lambda_s(n, m)$ and $\alpha(n)$ with $\alpha(n, m)$.

1.2. A Brief History of λ_s

After introducing the problem, Davenport and Schinzel [1965a] proved that $\lambda_1(n) = n$, $\lambda_2(n) = 2n - 1$, $\lambda_3(n) = O(n \log n)$, and for all $s \ge 4$, that $\lambda_s(n) = n \cdot 2^{O(\sqrt{\log n})}$, where the leading constant in the exponent depends on s. Shortly thereafter, Davenport [1971] improves the bound on $\lambda_3(n)$ to $O(n \log n / \log\log n)$. Szemerédi [1973] dramatically improves the upper bounds for all $s \ge 3$, showing that $\lambda_s(n) = O(n \log^{\star} n)$, where the leading constant depends on s.

From a purely numerical perspective Szemerédi's bound settled the problem for all values of n one might encounter in nature (the log-star function being at most 5 for n less than $10^{19,000}$), so why should any thoughtful mathematician continue to work on the problem? In our view, the problem of quantitatively estimating $\lambda_s(n)$ has always been a proxy for several qualitative questions: Is $\lambda_s(n)$ linear or nonlinear? What is the

[3]If $f : \mathbb{N}\backslash\{0\} \to \mathbb{N}$ is a decreasing function, $f^{\star}(m)$ is, by definition, $\min\{\ell \mid f^{(\ell)}(m) \le 1\}$, where $f^{(0)}(m) = m$ and $f^{(\ell)}(m) = f(f^{(\ell-1)}(m))$.

structure of extremal sequences realizing $\lambda_s(n)$? Does it even matter what s is? Hart and Sharir [1986] answer the first two questions for order-3 DS sequences. They gave a bijection between order-3 (*ababa*-free) DS sequences and so-called generalized postorder path compression schemes. Although these schemes resembled the path compressions found in set-union data structures, Tarjan's [1975] analysis did not imply any nontrivial upper or lower bounds on their length. Hart and Sharir [1986] prove that such path compression schemes have length $\Theta(n\alpha(n))$, thereby settling the asymptotics of $\lambda_3(n)$.

Agarwal, Sharir, and Shor [1989] (improving on Sharir [1987, 1988]) give asymptotically tight bounds on order-4 DS sequences and reasonably tight bounds on higher order sequences.

$$\lambda_4(n) = \Theta(n \cdot 2^{\alpha(n)})$$

$$\lambda_s(n) \begin{cases} > n \cdot 2^{(1+o(1))\alpha^t(n)\,/\,t!} \\ < n \cdot 2^{(1+o(1))\alpha^t(n)} \end{cases} \qquad \text{for even } s \geq 6,\, t = \lfloor \tfrac{s-2}{2} \rfloor.$$

$$\lambda_s(n) \begin{cases} > \lambda_{s-1}(n) \\ < n \cdot (\alpha(n))^{(1+o(1))\alpha^t(n)} \end{cases} \qquad \text{for odd } s \geq 5,\, t = \lfloor \tfrac{s-2}{2} \rfloor.$$

For even s, their bounds were tight up to the constant in the exponent: 1 for the upper bound and $1/t!$ for the lower bound. Moreover, their lower bound construction gave a qualitatively satisfying answer to the question of how extremal sequences are structured when s is even. For odd s, the gap between upper and lower bounds was wider, the base of the exponent being 2 at the lower bound and $\alpha(n)$ at the upper bound.

Remark 1.1. The results of Agarwal, Sharir, and Shor [1989] force us to confront another question, namely, when is it safe to declare victory and call the problem closed? As Nivasch [2010, §8] observes, the $+o(1)$ in the exponent necessarily hides a $\pm\Omega(\alpha^{t-1}(n))$ term if we express the bound in an *Ackermann-invariant* fashion, that is, in terms of the generic $\alpha(n)$, without specifying the precise variant of Ackermann's function for which it is the inverse. Furthermore, under any of the definitions in the literature $\alpha(n)$ is an integer-valued function, whereas $\lambda_s(n)/n$ must increase fairly smoothly with n, that is, an estimate of $\lambda_s(n)$ that is expressed as a function of any integer-valued $\alpha(n)$ must be off by at least a $2^{\Omega(\alpha^{t-1}(n))}$ factor. A reasonable definition of *sharp bound* (when dealing with generalized Davenport-Schinzel sequences) is an expression that cannot be improved, given $\pm\Theta(1)$ uncertainty in the definition of $\alpha(n)$. For example, $\lambda_4(n) = \Theta(n2^{\alpha(n)})$ is sharp in this sense since the constant hidden by Θ reflects this uncertainty. In contrast, $\lambda_3(n) = \Theta(n\alpha(n))$ is not sharp in an Ackermann-invariant sense. See the tighter bounds on $\lambda_3(n)$ cited next and in Theorem 1.3.

Remark 1.1 brings up several issues that demand their own remarks.

Remark 1.2. We are not forced by nature to work with a generic $\alpha(n)$. It is conceivable that there is a "right" definition of α (for each s) and that this definition could be succinctly defined, but this still leaves the issue of α being integer-valued. To fix this problem, one might attempt to define a *continuous* α that interpolates the integer-valued α, in the same way that the Γ function interpolates the factorial function. Defining such an α is a Herculean task, and one that would probably be appreciated by a number of researchers countable on one hand. The most tractable problem brought up in Remark 1.1 is to bound the smoothness of $\lambda_s(n)/n$. Davenport and Schinzel's $O(n \log n)$ bound on $\lambda_3(n)$ implies good bounds on the smoothness of $\lambda_3(n)/n$. For higher orders, $s > 3$, the smoothness problem has received little attention. Nonetheless, it is trivial to see that $\lambda_s(2n)/2n = O(\lambda_s(n)/n)$, which is smooth enough for our argument.

Nivasch [2010] presented a superior method for upper bounding $\lambda_s(n)$. In addition, he provides a new construction of order-3 DS sequences that matched an earlier upper bound of Klazar [1999] up to the leading constant.

$$\lambda_s(n) = \begin{cases} 2n\alpha(n) + O(n\sqrt{\alpha(n)}) & \text{for } s = 3\text{; upper bound due to Klazar [1999].} \\ \Theta(n \cdot 2^{\alpha(n)}) & \text{for } s = 4. \\ n \cdot 2^{(1+o(1))\alpha^t(n)\,/\,t!} & \text{for even } s \geq 6,\ t = \lfloor \frac{s-2}{2} \rfloor. \end{cases}$$

$$\lambda_s(n) \begin{cases} > \lambda_{s-1}(n) & \\ < n \cdot (\alpha(n))^{(1+o(1))\alpha^t(n)\,/\,t!} & \text{for odd } s \geq 5,\ t = \lfloor \frac{s-2}{2} \rfloor. \end{cases} \tag{Niv}$$

This closed the problem for even $s \geq 6$ (the leading constant in the exponent being precisely $1/t!$) but left the odd case open. Alon et al. [2008] conjecture that the upper bounds (Niv) for odd orders are tight, that is, the base of the exponent is, in fact, $\alpha(n)$. This conjecture was spurred by their discovery of similar functions that arose in an apparently unrelated combinatorial problem, *stabbing interval chains with j-tuples* [Alon et al. 2008].

1.3. New Results

We provide new upper and lower bounds on the length of Davenport-Schinzel sequences and, in the process, refute conjectures by Alon et al. [2008, Section 5], Nivasch [2010, Section 8], and Pettie [2011b, Section 7].

THEOREM 1.3. *Let $\lambda_s(n)$ be the maximum length of an order-s Davenport-Schinzel sequence. For any $s \geq 1$, λ_s satisfies*

$$\lambda_s(n) = \begin{cases} n & s = 1 \\ 2n - 1 & s = 2 \\ 2n\alpha(n) + O(n) & s = 3 \\ \Theta(n2^{\alpha(n)}) & s = 4 \\ \Theta(n\alpha(n)2^{\alpha(n)}) & s = 5 \\ n \cdot 2^{(1+o(1))\alpha^t(n)/t!} & \textit{for both even and odd } s \geq 6,\ t = \lfloor \frac{s-2}{2} \rfloor. \end{cases}$$

Theorem 1.3 is optimal in that it provides the tightest bounds that can be expressed in an Ackermann-invaraint fashion (see Remark 1.1), and in this sense closes the Davenport-Schinzel problem.[4] However, we believe our primary contributions are not the tight asymptotic bounds per se but the structural differences they reveal between even and odd s. We can now give a cogent explanation for why odd orders $s \geq 5$ behave essentially like the preceding even orders and yet why they are intrinsically more difficult to understand.

1.4. Generalizations of Davenport-Schinzel Sequences

The (Niv) bounds are actually corollaries of a more general theorem [Nivasch 2010] concerning the length of sequences avoiding *catenated permutations*[5], which were introduced by Klazar [1992]. Define Perm$(r, s + 1)$ to be the set of all sequences obtained by concatenating $s + 1$ permutations over an r-letter alphabet. For example,

[4]The exponent $(1 + o(1))\alpha^t(n)/t!$ is the Ackermann-invariant expression $\alpha^t(n)/t! + O(\alpha^{t-1}(n))$.
[5]Nivasch called these formation-free sequences.

abcd cbad badc abcd dcba $\in$ Perm(4, 5). Define the extremal function of Perm($r, s+1$)-free sequences to be

$$\Lambda_{r,s}(n) = \mathrm{Ex}(\mathrm{Perm}(r, s+1), n).$$

The $s+1$ here is chosen to highlight the parallels with order-s DS sequences. Every $\sigma \in \mathrm{Perm}(2, s+1)$ contains an alternating sequence $abab\cdots$ with length $s+2$,[6] so order-s DS sequences are also Perm($2, s+1$)-free, implying that $\lambda_s(n) \le \Lambda_{2,s}(n)$. Nivasch [2010] proves that $\Lambda_{r,s}(n)$ obeys all the upper bounds of (Niv), as well as its lower bounds when $s \ge 4$ is even or $s \le 3$.

There are other natural ways to generalize standard Davenport-Schinzel sequences. *Doubled* Davenport-Schinzel sequences are studied [Adamec et al. 1992; Davenport and Schinzel 1965b; Klazar and Valtr 1994; Pettie 2011b]. Define $\lambda_s^{\mathrm{dbl}}(n)$ to be the extremal function of dbl($abab\cdots$)-free sequences, where the alternating sequence has length $s+2$ and dbl(σ) is obtained by doubling every symbol in σ save the first and last. For example, dbl($abab$) $=$ $abbaab$.[7] Davenport and Schinzel [1965b] noted that $\lambda_1^{\mathrm{dbl}}(n) = O(\lambda_1(n)) = O(n)$ (see Klazar [2002]) and Adamec et al. [1992] proved that $\lambda_2^{\mathrm{dbl}}(n) = O(\lambda_2(n)) = O(n)$. Pettie proves that $\lambda_3^{\mathrm{dbl}}(n) = O(n\alpha^2(n))$ and that $\lambda_s^{\mathrm{dbl}}(n)$ obeys all the upper bounds of (Niv) for $s \ge 4$.

If one views alternating sequences as forming a zigzagging pattern, an obvious generalization is to extend the length of each zig and zag to include a larger alphabet. For example, the N-shaped sequences $N_k = 12\cdots k(k+1)\cdots 212\cdots k(k+1)$ generalize $abab = N_1$, and the M-shaped sequences $M_k = 12\cdots k(k+1)k\cdots 212\cdots k(k+1)k\cdots 21$ generalize $ababa = M_1$. Klazar and Valtr [1994] (see also [Pettie 2011c]) proved that $\mathrm{Ex}(\mathrm{dbl}(N_k), n) = O(\lambda_2(n)) = O(n)$, and Pettie [2011c] proved that $\mathrm{Ex}(\{M_k, ababab\}, n) = O(\lambda_3(n))$. Sequences avoiding N- and M-shaped sequences have proved very useful in bounding the complexity of geometric graphs [Fox et al. 2013; Pettie 2011c; Suk 2012; Valtr 1997].

In a companion paper [Pettie 2015], we provide new upper and lower bounds on doubled DS sequences, M_k-free sequences, and both Perm($r, s+1$)-free and dbl(Perm($r, s+1$))-free sequences. Let $\Lambda_{r,s}^{\mathrm{dbl}}(n)$ be the extremal function of dbl(Perm($r, s+1$))-free sequences. The strangest of these results is that $\Lambda_{r,s}$ is very sensitive to the alphabet size r, but only when s is odd and at least 5. In particular, $\Lambda_{2,s}(n) = \Theta(\lambda_s^{\mathrm{dbl}}(n)) = \Theta(\lambda_s(n))$, but this is not true for general $r \ne 2$.

THEOREM 1.4. *The following bounds hold for all $r \ge 2$, $s \ge 1$, where $t = \lfloor \frac{s-2}{2} \rfloor$.*

$$\Lambda_{r,s}(n),\ \Lambda_{r,s}^{\mathrm{dbl}}(n) = \begin{cases} \Theta(n) & \text{for } s \in \{1,2\} \text{ and all } r \ge 2 \\ \Theta(n\alpha(n)) & \text{for } s = 3 \text{ and all } r \ge 2 \\ \Theta(n2^{\alpha(n)}) & \text{for } s = 4 \text{ and all } r \ge 2 \\ \Theta(n\alpha(n)2^{\alpha(n)}) & \text{for } s = 5 \text{ and } r = 2 \\ n \cdot (\alpha(n))^{(1+o(1))\alpha(n)} & \text{for } s = 5 \text{ and all } r \ge 3 \\ n \cdot 2^{(1+o(1))\alpha^t(n)/t!} & \text{for even } s \ge 6 \text{ and all } r \ge 2 \\ n \cdot 2^{(1+o(1))\alpha^t(n)/t!} & \text{for odd } s \ge 7 \text{ and } r = 2 \\ n \cdot (\alpha(n))^{(1+o(1))\alpha^t(n)/t!} & \text{for odd } s \ge 7 \text{ and all } r \ge 3. \end{cases}$$

[6]The first permutation contributes two symbols and every subsequent permutation contributes at least one.
[7]Why not consider higher multiplicities? It is fairly easy to show that repeating symbols more than twice, or repeating the first and last at all, affects the extremal function by at most a constant factor. See Adamec et al. [1992].

Theorem 1.4 is rather surprising, even given Theorem 1.3 and even in retrospect. One consequence of Theorem 1.4 is that Cibulka and Kynčl's [2012] upper bounds on the size of sets of permutations with fixed VC-dimension are tight.

1.5. Organization

In Section 2, we present an informal discussion of the method of Agarwal, Sharir, and Shor [1989] and Nivasch [2010], its limitations for dealing with odd-order DS sequences, and the key ideas behind the proof of Theorem 1.3. Section 3 reviews Nivasch's recurrence for λ_s as well as some basic upper bounds on λ_s. The critical structure in our analysis is the *derivation tree* of a DS sequence. Its properties are analyzed in Section 4. In Section 5, we use the derivation tree to obtain a new recurrence for odd-order DS sequences. The recurrences for even- and odd-order DS sequences are solved in Section 5.1. In Section 5.2, we complete the proof of the upper bounds of Theorem 1.3 for all orders, with two exceptions. Order-3 and order-5 DS sequences require special attention. They are analyzed in Section 6. In Section 7, we establish Theorem 1.3's lower bound on order-5 DS sequences. We discuss several open problems in Section 8. Some proofs appear in Appendices A and B.

2. A TOUR OF THE PROOF

The proof of Theorem 1.3 diverges sharply from previous analyses [Agarwal et al. 1989; Nivasch 2010; Sharir 1987] in that it treats even and odd orders as fundamentally different beasts. To understand why all orders cannot be analyzed in a uniform fashion, we must review the method of Agarwal, Sharir, and Shor [1989] and Nivasch [2010].

The basic inductive hypothesis of Nivasch [2010] is that there are values $\{\mu_{s,i}\}$, increasing in both s and i, for which $\lambda_s(n, m) < \mu_{s,i}(n + m\,\mathrm{poly}(\log^{[i-1]}(m)))$, for any choice of i.[8] In other words, the multiplicity of symbols is at most $\mu_{s,i}$, up to an additive term that depends on the block count m, which may be the dominant term if i is too small. One would, ultimately, choose i to make $m\,\mathrm{poly}(\log^{[i-1]}(m)) = O(n)$ (and $\alpha(n, m) + O(1)$ is a good choice), but for the sake of simplifying the discussion, we ignore the dependence on m. Given a sequence S with parameters s, n, m, to *invoke the inductive hypothesis with parameter* i means to upper bound $|S|$ by $\mu_{s,i}n$, with the understanding that i is not chosen to be too small.

Suppose S is an order-s, m-block DS sequence over an n-letter alphabet. The analysis of Agarwal et al. [1989] and Nivasch [2010] begins by partitioning S into $\hat{m}$ intervals of consecutive blocks, $\hat{m}$ typically being much smaller than m. Write S as $S_1 \cdots S_{\hat{m}}$. We can put symbols into two categories: *local* symbols are those that appear exclusively in one interval S_q and *global* symbols are those that appear in multiple intervals. Let $\check{S} = \check{S}_1 \cdots \check{S}_{\hat{m}}$ be the subsequence of S consisting of local symbols and $\hat{S} = \hat{S}_1 \cdots \hat{S}_{\hat{m}}$ the subsequence of S consisting of global symbols, so $|S| = |\check{S}| + |\hat{S}|$. On each $\check{S}_q$ (for $1 \le q \le \hat{m}$), we invoke the inductive hypothesis with parameter i and deduce that $|\check{S}| \le \mu_{s,i}\|\check{S}\|$. What remains is to bound the length of $\hat{S}$.

The next step is to form a *contracted* sequence from $\hat{S}$ that has a much higher alphabet-to-block count ratio, thereby allowing us to invoke the inductive hypothesis with a smaller "i" parameter. Let $\hat{S}'$ be obtained from $\hat{S}$ by replacing each interval $\hat{S}_q$ with a single block β_q containing the first occurrence of each distinct symbol in $\hat{S}_q$. Thus, the alphabet of $\hat{S}'$ is the same as $\hat{S}$ but it consists of just $\hat{m}$ blocks $\beta_1\beta_2 \cdots \beta_{\hat{m}}$. On $\hat{S}'$, we invoke the inductive hypothesis with parameter $i - 1$ and conclude that $|\hat{S}'| \le \mu_{s,i-1}\|\hat{S}\|$. One cannot immediately deduce any bound on $\hat{S}$ from a bound on $\hat{S}'$,

[8] Recall that $\log^{[i-1]}(m)$ is the $\log^{\star\cdots\star}(m)$ function, with $i-1$ $\star$s.

since each interval $\hat{S}_q$ could contain numerous copies of a symbol, only one of which is retained in β_q.

Imagine reversing the contraction operation. We replace each block β_q with a sequence $\hat{S}_q$, thereby reconstructing $\hat{S}$. To bound the length of $\hat{S}_q$ in terms of $|\beta_q| = \|\hat{S}_q\|$, we will invoke the inductive hypothesis three more times. Put the symbols of β_q into three categories: those that make their first appearance (in $\hat{S}'$) in β_q, those that make their last appearance in β_q, and those that make a middle (non-first, non-last) appearance in β_q. Discard from $\hat{S}_q$ all symbols not classified as first in β_q and call the resulting sequence $\acute{S}_q$. Every symbol in $\acute{S}_q$ appears at least once after $\hat{S}_q$ (by virtue of being categorized as first in β_q), which implies that $\acute{S} = \acute{S}_1 \cdots \acute{S}_{\hat{m}}$ is an order-$(s-1)$ DS sequence. See the following diagram.

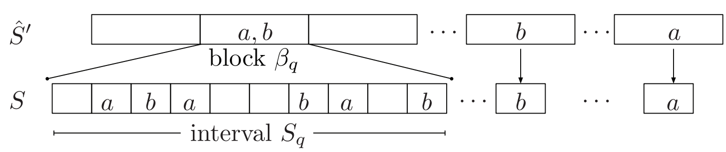

An occurrence of $\sigma_{s+1} = abab\cdots$ (length $s+1$) in $\acute{S}_q$, together with an a or b following $\hat{S}_q$ (depending on whether σ_{s+1} ends in b or a) gives an instance of σ_{s+2} in $\hat{S}$, contradicting the fact that it has order s. The same argument applies in a symmetric fashion to the subsequence of $\hat{S}_q$ formed by symbols making their last appearance in β_q, call it $\grave{S}_q$. By invoking the inductive hypothesis with parameter i on $\acute{S}_1 \cdots \acute{S}_{\hat{m}}$ and $\grave{S}_1 \cdots \grave{S}_{\hat{m}}$, we can conclude the contribution of first and last symbols to S is $2\mu_{s-1,i}\|\hat{S}\|$.

The length of the subsequence of middle symbols in $\hat{S}_q$, call it $\bar{S}_q$, is bounded with the same argument, except now there are, by definition of *middle*, occurrences of both $a, b \in \Sigma(\bar{S}_q)$ both before and after $\hat{S}_q$. That is, if $\sigma_s = baba\cdots$ (length s) appeared in $\bar{S}_q$, then, together with an a preceding $\hat{S}_q$ and either an a or b following $\hat{S}_q$ (depending on whether σ_s ends in b or a), there would be an instance of σ_{s+2} in $\hat{S}$ contradicting the fact that it has order s. We invoke the inductive hypothesis one last time, with parameter i, on *each* $\bar{S}_1, \ldots, \bar{S}_{\hat{m}}$, which implies that $|\bar{S}_q| \le \mu_{s-2,i}\|\bar{S}_q\|$.

Recall that each symbol in $\hat{S}'$ appeared $\mu_{s,i-1}$ times, $\mu_{s,i-1}-2$ times in blocks where it is categorized as middle. Thus, the contribution of middle symbols to $|\hat{S}|$ is $\mu_{s-2,i}(\mu_{s,i-1}-2)$. In order for every symbol, local and global alike, to appear in S with multiplicity at most $\mu_{s,i}$, we must have

$$\mu_{s,i} \ge 2\mu_{s-1,i} + \mu_{s-2,i}(\mu_{s,i-1} - 2). \tag{1}$$

When $s = 3$, we do not need to use an inductive hypothesis to determine $\mu_{1,i}$ and $\mu_{2,i}$. They are just 1 and 2; the i parameter does not come into play.[9] This leads to a bound of $\mu_{3,i} = 2i + O(1)$.[10] Although the contribution of first and last symbols is significant at $s = 3$, entertain the idea that their contribution becomes negligible at higher orders, so we can further simplify (1) as follows:

$$\mu_{s,i} \ge \mu_{s-2,i}\mu_{s,i-1}. \tag{2}$$

Inequality (2) is satisfied when $\mu_{s,i} = g^{\binom{i+t}{t}}$ for any base g; recall that $t = \lfloor \frac{s-2}{2} \rfloor$ by definition. By Pascal's identity, $g^{\binom{i+t}{t}} = g^{\binom{i+(t-1)}{(t-1)}} \cdot g^{\binom{(i-1)+t}{t}}$. The correct base depends on

[9]It is easy to show that $\lambda_1(n, m) < \underline{1} \cdot n + m$ and $\lambda_2(n, m) < \underline{2} \cdot n + m$. See Lemma 3.2.
[10]We have not said what to do in the base case when $i = 1$, which determines the $O(1)$ term.

where the inductively defined Inequality (2) bottoms out, at order 2 when $s \geq 4$ is even and at order 3 when $s \geq 5$ is odd. When s is even, the correct base is $2 = \mu_{2,i}$. When s is odd, the calculations are less clean, since $\mu_{3,i} = 2i + O(1)$ is not constant but depends on i. Nonetheless, the correct base is on the order of i, that is, $\mu_{s,i} = \Theta(i)^{\binom{i+t}{t}}$ satisfies (2) at the odd orders. Plugging in $\alpha(n, m)$ for i ultimately leads to Nivasch's bounds (Niv), since $\binom{i+t}{t} = i^t/t! + O(i^{t-1}) = (1 + o(1))i^t/t!$.

To obtain a construction of order-s sequences realizing the (Niv) bounds, one should start by attempting to reverse-engineer the preceding argument. To form an order-s sequence S with certain alphabet and block parameters, start by generating (inductively) local order-s sequences $\check{S}_1 \cdots \check{S}_{\hat{m}}$ over disjoint alphabets, and a single global order-s sequence $\hat{S}'$ having $\hat{m}$ blocks. Take some block β_q in $\hat{S}'$ and suppose for the sake of simplicity that β_q consists solely of middle symbols. We need to substitute for β_q an order-$(s-2)$ DS sequence $\bar{S}_q$ and then somehow merge it with $\check{S}_q$ in a way that does not introduce into S an alternating sequence with length $s + 2$. This is the point at which the even and odd orders diverge.

If s is even, the longest alternating sequence $baba \cdots b$ in $\bar{S}_q$ has length $s - 1$ and therefore begins and ends with b. We can only afford to introduce one alternation at each boundary of $\bar{S}_q$, so the pattern of as and bs on either side of β_q must look like $a^*\ b^*\ \ \beta_q\ \ b^*\ a^*$, as in the following diagram. We will call a and b *nested in* β_q if the sequence contains $a\,b\,\beta_q\,b\,a$ or the equivalent $b\,a\,\beta_q\,a\,b$.

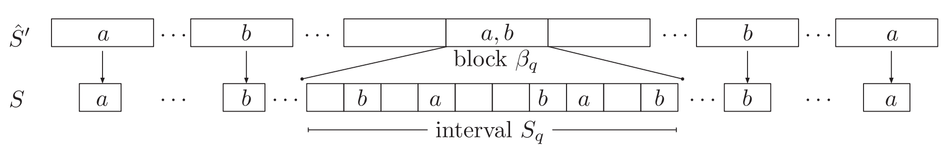

On the other hand, if s is odd, then the longest alternating sequence $baba \cdots a$ in $\bar{S}_q$ has length $s - 1$, begins with b and ends with a, so the pattern of as and bs in $\hat{S}'$ looks like $a^*\ b^*\ \beta_q\ a^*\ b^*$. A pair of middle symbols that are not nested in β_q are called *interleaved in* β_q.

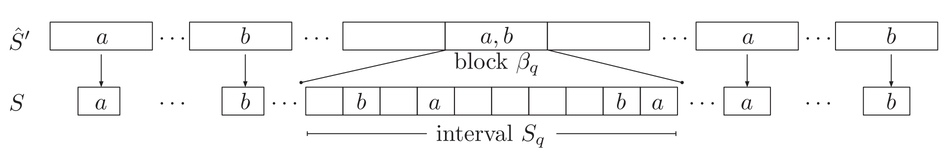

If the (Niv) bounds prove to be tight, there must be two systems for generating sequences: one where nesting is the norm, when s is even, and one where interleaving is the norm, when s is odd. If interleaving were somehow outlawed then to avoid creating an alternating sequence with length $s + 2$, the sequence $\bar{S}_q$ substituted for β_q would have to be an order-$(s-3)$ DS sequence rather than an order-$(s-2)$ one. However, it is clearly impossible to claim that interleaving simply cannot exist.

What makes the argument of Agarwal et al. [1989] and Nivasch [2010] brilliantly simple is how little it leaves to direct calculation. The length of every sequence ($\check{S}_q$, $\hat{S}'$, $\acute{S}_q$, $\bar{S}_q$, etc.) is bounded by delegation to an inductive hypothesis. *However, such useful notions as nearly all middle symbols in a block are mutually nested are difficult to capture in a strengthened inductive hypothesis*. We need to understand and characterize the phenomenon of nestedness to improve on Agarwal et al. [1989] and Nivasch [2010]. This requires a deeper understanding of the structure of Davenport-Schinzel sequences.

The Derivation Tree. Inductively defined objects can be apprehended inductively or, alternatively, apprehended holistically by completely "unrolling" the induction. From the first perspective, S is the merger of $\check{S}$ and $\hat{S}$, which is derived from $\hat{S}'$, all of which are analyzed inductively. By iteratively unrolling the decomposition of $\check{S}$ and $\hat{S}'$, we obtain a *derivation tree* $\mathcal{T}$ whose nodes represent every block in every sequence encountered in the recursive decomposition of S. Whereas S occupies the leaves of $\mathcal{T}$, derived sequences such as $\hat{S}'$ occupy levels higher in $\mathcal{T}$. Whereas S (and every sequence) is a static object, $\mathcal{T}$ can be thought of as a process for generating S whose history can be reasoned about explicitly. But how does $\mathcal{T}$ let us deduce something about the nestedness and non-nestedness of symbols in a common block?

Suppose we are interested in the nestedness of middle symbols a, b in block β, which corresponds to a leaf-node in $\mathcal{T}$. Imagine taking $\mathcal{T}$ and deleting every node whose block does not contain b, that is, *projecting* $\mathcal{T}$ onto the symbol b. What remains, $\mathcal{T}_{|b}$, is a tree rooted at the location in $\mathcal{T}$ where b is "born" and represents how occurrences of b have proliferated during the process that culminates in the construction of S. The block/node β occupies a location in $\mathcal{T}_{|b}$ and a location in $\mathcal{T}_{|a}$, whose node sets are only guaranteed to intersect at β. Some locations in $\mathcal{T}_{|a}$ and $\mathcal{T}_{|b}$ are intrinsically bad—these are called *feathers* in Section 4. (Whether a node is a feather in $\mathcal{T}_{|a}$ depends solely on the structure of $\mathcal{T}_{|a}$, not how it is embedded in $\mathcal{T}$ nor its relationship to a different $\mathcal{T}_{|b}$.) We show that if β is not a feather in $\mathcal{T}_{|a}$ and not a feather in $\mathcal{T}_{|b}$, then a and b are nested in β. In other words, the middle symbols in β are partitioned into two equivalence classes, depending on whether or not they appear at feathers in their respective derivation trees. We could not outlaw interleavedness in general, yet we manage to outlaw it within one equivalence class! The question is, what are the relative sizes of these two equivalence classes, and in particular, how many feathers can a $\mathcal{T}_{|a}$ have?

Our aim is to get stronger asymptotic bounds on λ_s for odd s, which means the number of feathers should be a negligible ($o(1)$ fraction) of the size of $\mathcal{T}_{|a}$. In the same way that the multiplicities $\mu_{s,i}$ are bounded inductively, as in (1), for example, we are able to bound the number of feathers in $\mathcal{T}_{|a}$ inductively, call it $\nu_{s,i}$, in terms of $\nu_{s,i-1}$ and $\nu_{s-1,i}$. However, now $\mu_{s,i}$ is bounded in terms of $\mu_{s,i-1}$ (the multiplicity of symbols in the contracted sequence $\hat{S}'$), $\mu_{s-1,i}$ (the multiplicity of symbols in $\hat{S}$ begat by first and last occurrences in $\hat{S}'$), $\mu_{s-3,i}$ (the multiplicity of middle occurrences in $\hat{S}$ begat by non-feathers in $\hat{S}'$), and both $\nu_{s,i-1}$ and $\mu_{s-2,i}$, which count the number of feathers in $\hat{S}'$ and the multiplicity of middle occurrences in $\hat{S}$ begat by feathers in $\hat{S}'$. This leads to a system of three interconnected recurrences: one for $\mu_{s,i}$ at odd s, one for $\mu_{s,i}$ at even s, and one for the feather count $\nu_{s,i}$. An elementary (though necessarily detailed) proof by induction gives solutions for $\mu_{s,i}$ and $\nu_{s,i}$ that ultimately lead to the upper bounds of Theorem 1.3.

3. BASIC UPPER BOUNDS

In Section 3.1, we review and expand on the notation introduced informally in Section 2. It will be used repeatedly throughout Sections 4–6.

3.1. Sequence Decomposition

Let S be a sequence over an $n = \|S\|$ letter alphabet consisting of $m = [\![S]\!]$ blocks. Suppose we partition S into $\hat{m}$ intervals of consecutive blocks $S_1 S_2 \cdots S_{\hat{m}}$, where $m_q = [\![S_q]\!]$ is the number of blocks in interval q. Let $\check{\Sigma}_q$ be the alphabet of symbols *local* to S_q (that do not appear in any S_p, $p \neq q$), and let $\hat{\Sigma} = \Sigma(S) \backslash \bigcup_q \check{\Sigma}_q$ be the alphabet of all other *global* symbols. The cardinalities of $\check{\Sigma}_q$ and $\hat{\Sigma}$ are $\check{n}_q$ and $\hat{n}$, thus $n = \hat{n} + \sum_{q=1}^{\hat{m}} \check{n}_q$. A global symbol in S_q is called *first*, *last*, or *middle* if it appears in no earlier interval, no later interval, or appears in both earlier and later intervals, respectively. Let $\acute{\Sigma}_q$, $\grave{\Sigma}_q$, $\bar{\Sigma}_q$, $\hat{\Sigma}_q$ be the subset of $\Sigma(S_q)$ consisting of, respectively, first, last, middle, and all global symbols, and let $\acute{n}_q, \grave{n}_q, \bar{n}_q$, and $\hat{n}_q$ be their cardinalities. Let $\check{S}_q, \hat{S}_q, \acute{S}_q, \grave{S}_q, \bar{S}_q$ be the projection of S_q onto $\check{\Sigma}_q$, $\hat{\Sigma}_q$, $\acute{\Sigma}_q$, $\grave{\Sigma}_q$, and $\bar{\Sigma}_q$. Note that $\hat{S}_1$ consists solely of first occurrences; if the last occurrence of a symbol appeared in $\hat{S}_1$ the symbol would be classified as local to S_1, not global. The same argument shows that $\hat{S}_{\hat{m}}$ consists solely of last occurrences. Let $\check{S}$, $\hat{S}$, $\acute{S}$, $\grave{S}$, and $\bar{S}$ be the subsequences of local, global, first, last, and middle occurrences, respectively, that is, $\check{S} = \check{S}_1 \cdots \check{S}_{\hat{m}}$, $\hat{S} = \hat{S}_1 \cdots \hat{S}_{\hat{m}}$, $\acute{S} = \acute{S}_1 \cdots \acute{S}_{\hat{m}-1}$, $\grave{S} = \grave{S}_2 \cdots \grave{S}_{\hat{m}}$, and $\bar{S} = \bar{S}_2 \cdots \bar{S}_{\hat{m}-1}$, the last of which would be empty if $\hat{m} = 2$. Let $\hat{S}' = \beta_1 \cdots \beta_{\hat{m}}$ be an $\hat{m}$-block sequence obtained from $\hat{S}$ by replacing each $\hat{S}_q$ with a single block β_q containing its alphabet $\hat{\Sigma}_q$, listed in order of *first* appearance in $\hat{S}_q$.

3.2. 2-Sparse vs. Blocked Sequences

Every analysis of Davenport-Schinzel sequences since Hart and Sharir [1986] uses Lemma 3.1(2) to reduce the problem of bounding 2-sparse DS sequences to bounding m-block DS sequences, that is, expressing $\lambda_s(n)$ in terms of $\lambda_s(n, m)$, where $m = O(n)$.

LEMMA 3.1. *Let $\gamma_s(n) : \mathbb{N} \to \mathbb{N}$ be a nondecreasing function such that $\lambda_s(n) \leq \gamma_s(n) \cdot n$.*

(1) *For $s \geq 1$, $\lambda_s(n, m) \leq m - 1 + \lambda_s(n)$.*
(2) *For $s \geq 3$, $\lambda_s(n) \leq \gamma_{s-2}(n) \cdot \lambda_s(n, 2n - 1)$. (This generalizes the proof of Hart and Sharir [1986] for $s = 3$.)*
(3) *For $s \geq 2$, $\lambda_s(n) \leq \gamma_{s-1}(n) \cdot \lambda_s(n, n)$.*
(4) *For $s \geq 3$, $\lambda_s(n) \leq \gamma_{s-2}(\gamma_s(n)) \cdot \lambda_s(n, 3n - 1)$.*

Parts 2 and 3 of Lemma 3.1 are due to Sharir [1987].

Part 4 of Lemma 3.1 improves on parts 2 and 3 when $\gamma_{s-2}(n) = \omega(1)$, that is, when $s \geq 5$. Our upper bounds do not need the power of part 4. Part 2 suffices when $s \leq 4$ or $s \geq 6$ and $s = 5$ is handled as a special case. Nonetheless, an analogue of part 4 can be useful when analyzing generalized DS sequences [Pettie 2015]. We include a proof of Lemma 3.1 in Appendix A.

3.3. Orders 1 and 2

In the interest of completeness, we shall reestablish the known bounds on order-1 and order-2 DS sequences, in both their 2-sparse and blocked forms.

LEMMA 3.2 (DAVENPORT AND SCHINZEL [1965a]). *The extremal functions for order-1 and order-2 DS sequences are*

$$\begin{aligned} \lambda_1(n) &= n, \\ \lambda_2(n) &= 2n - 1, \\ \lambda_1(n, m) &= n + m - 1, \\ \lambda_2(n, m) &= 2n + m - 2. && \text{for } m \geq 2 \end{aligned}$$

PROOF. Let S be a 2-sparse sequence with $n = \|S\|$. If $|S| > n$, then there are two copies of some symbol, say a. The as cannot be adjacent, due to 2-sparseness, so S must contain a subsequence aba, for some $b \neq a$. Such an S is not an order-1 DS sequence, hence $\lambda_1(n) \leq n$.

If S has order 2, then some symbol must appear exactly once. To see this, consider the closest pair of occurrences of some symbol, say a. If every symbol b appearing between this pair of as occurred twice in S then S would contain $baba$, $abab$, or $abba$. The first two are precluded since S has order 2 and the third violates the fact that the two as are the closest such pair. Thus, every symbol b between the two as occurs once. Remove one such b; if this causes the two a's to become adjacent, remove one of the a's. What remains is a 2-sparse sequence over an $(n-1)$-letter alphabet, so $\lambda_2(n) \le \lambda_2(n-1)+2$. Since $\lambda_2(1) = 1$ we have $\lambda_2(n) \le 2n-1$.

Lemma 3.1(1) and the bounds previously established imply $\lambda_1(n, m) \le n + m - 1$ and $\lambda_2(n, m) \le 2n + m - 2$. All these upper bounds are tight. The unique extremal order-1, 2-sparse DS sequence is $123 \cdots n$, which can be converted into an extremal m-block sequence $[123 \cdots n][n]^{m-1}$. Brackets mark block boundaries. There are exponentially many extremal DS sequences of order 2, each corresponding to an Euler tour around a rooted tree with vertex labels from $\{1, \ldots, n\}$. For example, $123 \cdots (n-1)n(n-1) \cdots 321$ and $1213141 \cdots 1(n-1)1n1$ are extremal 2-sparse, order-2 DS sequences. The first corresponds to an Euler tour around a path, the second an Euler tour around a star. The first sequence can be converted into an extremal m-block, order-2 DS sequence $[12 \cdots (n-1)n][n(n-1) \cdots 21][1]^{m-2}$, assuming that $m \ge 2$. When there is only one block we have $\lambda_s(n, 1) = n$, regardless of the order s. □

3.4. Nivasch's Recurrence

Nivasch's [2010] upper bounds (Niv) are a consequence of a recurrence for λ_s that is stronger than the one of Agarwal et al. [1989]. Here we present a streamlined version of Nivasch's recurrence.

RECURRENCE 3.3. *Let m, n, and $s \ge 3$ be the block count, alphabet size, and order parameters. For any $\hat{m} < m$, any block partition $\{m_q\}_{1 \le q \le \hat{m}}$, and the corresponding alphabet partition $\{\hat{n}\} \cup \{\check{n}_q\}_{1 \le q \le \hat{m}}$, where $m = \sum_q m_q$ and $n = \hat{n} + \sum_q \check{n}_q$, we have*

$$\lambda_s(n, m) \;\le\; \sum_{q=1}^{\hat{m}} \lambda_s(\check{n}_q, m_q) \;+\; 2 \cdot \lambda_{s-1}(\hat{n}, m) \;+\; \lambda_{s-2}(\lambda_s(\hat{n}, \hat{m}) - 2\hat{n}, m).$$

PROOF. We adopt the notation and definitions from Section 3.1, where S is an order-s DS sequence with $\|S\| = n$ and $[\![S]\!] = m$. We shall bound $|S|$ by considering its four constituent subsequences $\check{S}$, $\acute{S}$, $\grave{S}$, and $\bar{S}$.

Each $\check{S}_q$ is an order-s DS sequence, therefore the contribution of local symbols is $|\check{S}| \le \sum_{q=1}^{\hat{m}} \lambda_s(\check{n}_q, m_q)$. We claim each $\acute{S}_q$ is an order-$(s-1)$ DS sequence. By virtue of being categorized as first in $\hat{S}_q$, every symbol in $\acute{S}_q$ appears at least once after $\acute{S}_q$. Therefore an occurrence of an alternating sequence $\sigma_{s+1} = abab \cdots$ (length $s+1$), in $\acute{S}_q$ would imply an occurrence of σ_{s+2} in S, a contradiction. By symmetry, it also follows that $\grave{S}_q$ is an order-$(s-1)$ DS sequence, hence $|\acute{S}| = \sum_{q=1}^{\hat{m}-1} \lambda_{s-1}(\acute{n}_q, m_q)$ and $|\grave{S}| = \sum_{q=2}^{\hat{m}} \lambda_{s-1}(\grave{n}_q, m_q)$. Since λ_s is clearly superadditive[11], we can bound these sums by $\lambda_{s-1}(\hat{n}, m - m_{\hat{m}})$ and $\lambda_{s-1}(\hat{n}, m - m_1)$. (Note that $\sum_q \acute{n}_q = \hat{n}$ and $\sum_q \grave{n}_q = \hat{n}$, as each sum counts each global symbol exactly once.) The contribution of first and last symbols is therefore upper bounded by $2 \cdot \lambda_{s-1}(\hat{n}, m)$.

The same argument shows that $\bar{S}_q$ is an order-$(s-2)$ DS sequence. Symbols in $\bar{S}_q$ were categorized as middle, so an alternating subsequence $\sigma_s = baba \cdots$ (length s) in $\bar{S}_q$, together with an a preceding $\bar{S}_q$ and either an a or b following $\bar{S}_q$ (depending on whether s is even or odd), yields an instance of σ_{s+2} in S, a contradiction. Thus the

[11]It is straightforward to show that $\lambda_s(n', m') + \lambda_s(n'', m'') \le \lambda_s(n' + n'', m' + m'' - 1)$, for all n', n'', m', m''.

contribution of middle symbols is

$$
\begin{aligned}
|\bar{S}| &\le \sum_{q=2}^{\hat{m}-1} \lambda_{s-2}(\bar{n}_q, m_q) \\
&\le \lambda_{s-2}\left(\sum_{q=2}^{\hat{m}-1} \bar{n}_q, m - m_1 - m_{\hat{m}}\right) && \{\text{superadditivity of } \lambda_{s-2}\} \\
&\le \lambda_{s-2}(|\hat{S}'| - 2\hat{n}, m - m_1 - m_{\hat{m}}), && (3) \\
&\le \lambda_{s-2}(\lambda_s(\hat{n}, \hat{m}) - 2\hat{n}, m). && (4)
\end{aligned}
$$

Inequality (3) follows from the fact that $\sum_q \bar{n}_q$ counts the length of $\hat{S}'$, save the first and last occurrence of each global symbol, that is, $2\hat{n}$ occurrences in total. Since $\hat{S}'$ is a subsequence of S, it too is an order-s DS sequence, so $|\hat{S}'| \le \lambda_s(\hat{n}, \hat{m})$. Inequality (4) follows. □

Recurrence 3.3 offers us the freedom to choose the block partition $\{m_q\}_{1 \le q \le \hat{m}}$, but it does not suggest what the optimal partition might look like. One natural starting place [Agarwal et al. 1989; Hart and Sharir 1986; Nivasch 2010] is to always choose $\hat{m} = 2$, partitioning the sequence into two intervals each containing $m/2$ blocks. This choice leads to $O(n + m \log^{s-2} m)$ upper bounds on $\lambda_s(n, m)$, which is $O(n + m)$ if the alphabet/block density $n/m = \Omega(\log^{s-2} m)$. Call this Analysis (1). Given Analysis (1), we can conduct a stronger Analysis (2) by selecting $\hat{m} = m/\log^{s-2} m$, so each interval contains $\log^{s-2} m$ blocks. The $\lambda_s(\hat{n}, \hat{m})$ term is bounded via Analysis (1) (i.e., $\lambda_s(\hat{n}, \hat{m}) = O(\hat{n} + \hat{m} \log^{s-2} \hat{m}) = O(\hat{n} + m)$) and the remaining terms bounded inductively via analysis Analysis (2). This leads to bounds of the form $\lambda_s(n, m) = O(n + m\,\mathrm{poly}(\log^\star m))$. By iterating this process, Analysis (i) gives bounds of the form $O(n + m\,\mathrm{poly}(\log^{[i-1]} m))$.[12] We cannot conclude that $\lambda_s(n, m) = O(n + m)$ since the constant hidden by the asymptotic notation, call it $\mu_{s,i}$, increases with i and s.

This is merely meant to foreshadow the analysis of Recurrence 3.3 and subsequent Recurrences 5.1 and 5.2 (see Appendix B). We have made every attempt to segregate recurrences and structural arguments from their quantitative analyses, which are important but nonetheless rote. As a consequence, Ackermann's function, its various inverses, and quantities such as $\{\mu_{s,i}\}$ will be introduced as late as possible.

3.5. The Evolution of Recurrence 3.3

The statement of Recurrence 3.3 is simple and arguably cannot be made simpler. We feel it is worthwhile to recount how it has been assembled over the years [Agarwal et al. 1989; Hart and Sharir 1986; Klazar 1999; Nivasch 2010; Sharir 1987].

When s is fixed, the function $\lambda_s(n)$ depends only on one parameter, n, a situation that would not ordinarily lead to expressions involving α, which is most naturally expressed as a function of two independent parameters.[13] Hart and Sharir's [1986] insight recognizes an additional parameter m (the block count) and obtains bounds on $\lambda_s(n)$ via bounds on $\lambda_s(n, m)$ (see Lemma 3.1).

[12] The $[i-1]$ here being short for $i-1\ \star$s.

[13] In graph algorithms, these parameters typically correspond to nodes and edges [Chazelle 2000; Lengauer and Tarjan 1979; Tarjan 1979], in matrix problems [Klawe 1992; Klawe and Kleitman 1990] to rows and columns, and in data structures they may correspond to elements and queries [Gabow 1985; Tarjan 1975], query time, and preprocessing time [Pettie 2006], or input size and storage space [Alon and Schieber 1987; Chazelle and Rosenberg 1991; Yao 1982].

Implicit in Hart and Sharir's analysis is a classification of symbols into local and global, and of global occurrences into first, middle, and last.[14] Agarwal et al. [1989] make this local/global and first/middle/last classification explicit, and arrive at a recurrence very close to Recurrence 3.3.[15] However, they do not bound the contribution of global middle occurrences in the same way. Whereas $\bar{S}_q$ is an m_q-block sequence, it can be converted to a 2-sparse one by removing up to $m_q - 1$ repeated symbols at block boundaries. By Lemma 3.1(1,2),

$$|\bar{S}_q| < m_q + \lambda_{s-2}(\bar{n}_q) \;\leq\; m_q + \gamma_{s-2}(\bar{n}_q) \cdot \bar{n}_q \;\leq\; m_q + \gamma_{s-2}(n) \cdot \bar{n}_q.$$

In other words, when contracting $\bar{S}$ to form $\bar{S}'$, the shrinkage factor is at most $\gamma_{s-2}(n)$. A similar statement holds for first and last occurrences, where the shrinkage factor is at most $\gamma_{s-1}(n)$. This leads to a recurrence [Agarwal et al. 1989, p. 249] that forgets the role of m when analyzing global occurrences.

$$\lambda_s(n, m) \;\leq\; \sum_{q=1}^{\hat{m}} \lambda_s(\check{n}_q, m_q) \;+\; 2 \cdot \gamma_{s-1}(n) \cdot n \;+\; \gamma_{s-2}(n) \cdot \lambda_s(\hat{n}, \hat{m}) \;+\; O(m).$$

Nivasch's recurrence [2010, Recurrence 3.1] improves that of Agarwal et al. [1989] by not forgetting that $\bar{S}$ is an m-block sequence. In particular, $|\bar{S}| \leq \sum_q \lambda_{s-2}(\hat{n}_q, m_q)$, where $|\bar{S}'| < \sum_q \hat{n}_q \leq \lambda_s(\hat{n}, \hat{m})$. Recurrence 3.3 is substantively no different than that of Nivasch [2010], but it is more succinct, for two reasons. First, the superadditivity of λ_s lets us bound the number of middle occurrences with the single term $\lambda_{s-2}(\lambda_s(\hat{n}, \hat{m}) - 2\hat{n}, m)$.[16] Second, the function equivalent to $\lambda_s(n, m)$ from Agarwal et al. [1989] and Nivasch [2010] is the extremal function of order-s DS sequences that are both 2-sparse and have m blocks. This small change introduces $O(m)$ terms [Nivasch 2010, Recurrence 3.1; Agarwal et al. 1989, p. 249], since the derived sequences $\hat{S}$, $\hat{S}'$, and $\{\check{S}_q, \acute{S}_q, \bar{S}_q, \grave{S}_q\}_{1\leq q\leq \hat{m}}$ are not necessarily 2-sparse, and must be made 2-sparse by removing $O(m)$ symbols at block boundaries.

Recurrence 3.3 could be made yet more succinct by removing the $-2\hat{n}$ from the estimation of global middle occurrences. This would not affect the solution asymptotically, but keeping it is essential for obtaining bounds on $\lambda_3(n)$ tight to the leading constant.

4. DERIVATION TREES

A *derivation tree* $\mathcal{T}(S)$ for an m-block sequence S is a rooted, ordered tree whose nodes are identified with the blocks encountered in recursively decomposing S, as in Section 3.1 and Recurrence 3.3. Let $\mathcal{B}(u)$ be the block associated with node $u \in \mathcal{T}(S)$. The leaf level of $\mathcal{T}(S)$ coincides with S, that is, the pth leaf of $\mathcal{T}(S)$ holds the pth block of S. As we are sometimes indifferent to the order of symbols within a block, $\mathcal{B}(v)$ is often treated as a set. We assume without loss of generality that no symbol appears just once in S.[17] As usual, we adopt the sequence decomposition notation from Section 3.1.

Base Cases. Suppose $S = \beta_1\beta_2$ is a two-block sequence, where each block contains the whole alphabet $\Sigma(S)$. The tree $\mathcal{T}(S)$ consists of three nodes u, u_1, and u_2, where u

[14] This part of their analysis is ostensibly about nodes and path compressions, not blocks and symbols.

[15] Sharir [1987] splits global occurrences into two categories—first and non-first—which leads to a near-linear upper bound of $\lambda_s(n) < n \cdot \alpha(n)^{O(\alpha(n))^{s-3}}$.

[16] One might think it would be dangerous to bound middle occurrences with one aggregated term, since we "forget" that $\bar{S}$ is partitioned into $\hat{m} - 2$ order-$(s-2)$ DS sequences. Doing this does not affect the solution of $\lambda_s(n, m)$ asymptotically.

[17] This property is preserved throughout the recursive decomposition. If a symbol appears at least twice in S, it also appears at least twice in $\hat{S}'$ or some $\check{S}_q$, depending on whether it is global or local.

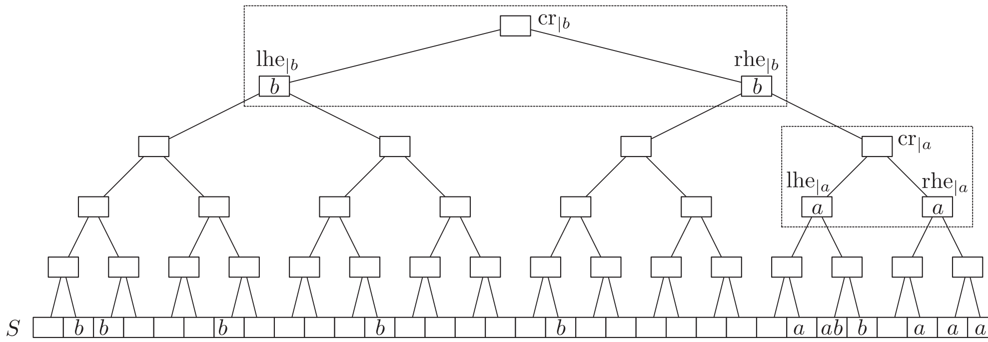

Fig. 1. A derivation tree $\mathcal{T}(S)$ for a 32-block sequence S. The tree is generated by always choosing $\hat{m} = 2$ and the block partition $m_1 = m_2 = m/2$, where $m > 2$ is the number of blocks in the given sequence. The frames isolate the base case derivation trees that assign the crown and heads for symbols $a, b \in \Sigma(S)$.

is the parent of u_1 and u_2, $\mathcal{B}(u_1) = \beta_1$, $\mathcal{B}(u_2) = \beta_2$, and $\mathcal{B}(u)$ does not exist. For every $a \in \Sigma(S)$, call u its *crown* and u_1 and u_2 its left and right *heads*, respectively. These nodes are denoted $\mathrm{cr}_{|a}$, $\mathrm{lhe}_{|a}$, and $\mathrm{rhe}_{|a}$.

When $S = \beta_1$ consists of a single block, $\mathcal{T}(S)$ consists of two nodes, u_1 and u, where $\mathcal{B}(u_1) = \beta_1$, $\mathcal{B}(u)$ does not exist, and u is the parent of u_1.

Inductive Case. If S contains $m > 2$ blocks, choose an $\hat{m} < m$ and an arbitrary block partition $\{m_q\}_{1\le q\le \hat{m}}$. Inductively construct derivation trees $\hat{\mathcal{T}} = \mathcal{T}(\hat{S}')$ and $\{\check{\mathcal{T}}_q\}_{1\le q\le \hat{m}}$, where $\check{\mathcal{T}}_q = \mathcal{T}(\check{S}_q)$, then identify the root of $\check{\mathcal{T}}_q$ (which has no block) with the qth leaf of $\hat{\mathcal{T}}$. Finally, place the blocks of S at the leaves of $\mathcal{T}$. This last step is necessary since only local symbols appear in the blocks of $\{\check{\mathcal{T}}_q\}$, whereas the leaves of $\mathcal{T}$ must be identified with the blocks of S. Note that nodes at or above the leaf level of $\hat{\mathcal{T}}$ carry only global symbols in their blocks and that internal nodes in $\{\check{\mathcal{T}}_q\}$ carry only local symbols in their blocks. Local and global symbols only mingle at the leaf level of $\mathcal{T}$.

The crown and heads of each symbol $a \in \Sigma(S)$ are inherited from $\hat{\mathcal{T}}$, if a is global, or some $\check{\mathcal{T}}_q$ if a is local to S_q. See Figure 1 for an illustration.

Remark 4.1. Trees defined recursively are typically built in a bottom-up or top-down fashion. Our algorithm for constructing $\mathcal{T}$ is somewhat unusual in that the trees defined by the two recursive invocations are joined at a level midway between the leaf level and root of $\mathcal{T}$. However, in the base case of our analysis, we happen to choose $\hat{m} = 2$ and $m_1 = m_2 = m/2$, as in Figure 1. In this special case, $\mathcal{T}$ is actually built in a top-down fashion.

4.1. Anatomy of the Tree

The *projection of* $\mathcal{T}$ *onto* $a \in \Sigma(S)$, denoted $\mathcal{T}_{|a}$, is the tree on the node set $\{\mathrm{cr}_{|a}\} \cup \{v \in \mathcal{T} \mid a \in \mathcal{B}(v)\}$ that inherits the ancestor/descendant relation from $\mathcal{T}$, that is, the parent of v in $\mathcal{T}_{|a}$, where $v \notin \{\mathrm{cr}_{|a}, \mathrm{lhe}_{|a}, \mathrm{rhe}_{|a}\}$, is v's nearest strict ancestor u for which $a \in \mathcal{B}(u)$. For example, in Figure 1, $\mathcal{T}_{|a}$ consists of $\mathrm{cr}_{|a}$, its children $\mathrm{lhe}_{|a}$, $\mathrm{rhe}_{|a}$, and five grandchildren at the leaf level of $\mathcal{T}$. As one can see, even though $\mathcal{T}$ is binary, $\mathcal{T}_{|a}$ is not necessarily binary.

Definition 4.2 (*Anatomy*).

—The leftmost and rightmost leaves of $\mathcal{T}_{|a}$ are *wingtips*, denoted $\mathrm{lwt}_{|a}$ and $\mathrm{rwt}_{|a}$.
—The left and right *wings* are those paths in $\mathcal{T}_{|a}$ extending from $\mathrm{lhe}_{|a}$ to $\mathrm{lwt}_{|a}$ and from $\mathrm{rhe}_{|a}$ to $\mathrm{rwt}_{|a}$.

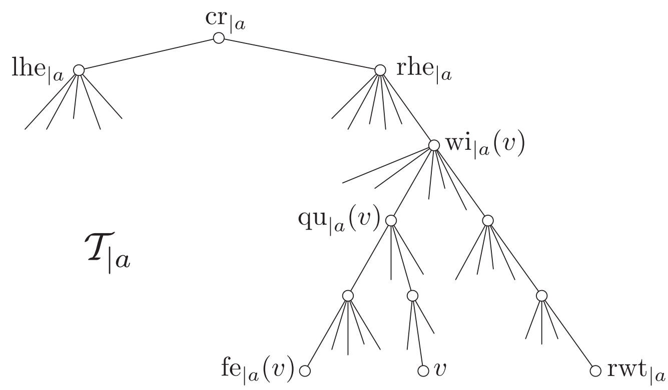

Fig. 2. In this example, v is a hawk leaf in $\mathcal{T}_{|a}$ since it is a descendant of $\mathrm{rhe}_{|a}$. Its wing node $\mathrm{wi}_{|a}(v)$, quill $\mathrm{qu}_{|a}(v)$, and feather $\mathrm{fe}_{|a}(v)$ are indicated.

—Descendants of $\mathrm{lhe}_{|a}$ and $\mathrm{rhe}_{|a}$ in $\mathcal{T}_{|a}$ are called *doves* and *hawks*, respectively.
—A child of a wing node that is not itself on the wing is called a *quill*.
—A leaf is called a *feather* if it is the rightmost descendant of a dove quill or leftmost descendant of a hawk quill.
—Suppose v is a node in $\mathcal{T}_{|a}$. Let $\mathrm{wi}_{|a}(v)$ be the nearest wing node ancestor of v, $\mathrm{qu}_{|a}(v)$ the quill ancestral to v, and $\mathrm{fe}_{|a}(v)$ the feather descending from $\mathrm{qu}_{|a}(v)$. See Figure 2 for an illustration.

Once $a \in \Sigma(S)$ is known or specified, we will use these terms (feather, wingtip, etc.) to refer to nodes in $\mathcal{T}_{|a}$ or to the occurrences of a within those blocks. For example, an occurrence of a in S would be a feather if it appears in a block $\mathcal{B}(v)$ in S, where v is a feather in $\mathcal{T}_{|a}$.

Note that the nodes $\mathrm{lhe}_{|a}, \mathrm{rhe}_{|a}, \mathrm{wi}_{|a}(v), \mathrm{qu}_{|a}(v)$, and $\mathrm{fe}_{|a}(v)$ are not necessarily all distinct. It may be that $\mathrm{wi}_{|a}(v)$ is equal to $\mathrm{lhe}_{|a}$ or $\mathrm{rhe}_{|a}$, and it may be that $v = \mathrm{qu}_{|a}(v) = \mathrm{fe}_{|a}(v)$, if v's parent in $\mathcal{T}_{|a}$ is $\mathrm{wi}_{|a}(v)$.

Lemma 4.3 identifies one property of $\mathcal{T}$ used in the proof of Lemma 4.4.

LEMMA 4.3. *Suppose that on a leaf-to-root path in $\mathcal{T}$, we encounter nodes u, v, x, and y (the last two possibly identical), where $u, x \in \mathcal{T}_{|a}$ and $v, y \in \mathcal{T}_{|b}$. It must be that $a \in \mathcal{B}(v)$ and therefore $v \in \mathcal{T}_{|a}$.*

PROOF. Consider the decomposition of $\mathcal{T}$ into a global derivation tree $\hat{\mathcal{T}}$ and local derivation trees $\{\check{\mathcal{T}}_q\}$. If v were an internal node in some $\check{\mathcal{T}}_q$ then b would be classified as local. This implies $y \in \check{\mathcal{T}}_q$ as well and the claim follows by induction on the construction of $\check{\mathcal{T}}_q$. If v were an internal node in $\hat{\mathcal{T}}$, then let u' be the leaf of $\hat{\mathcal{T}}$ ancestral to u. The nodes $u', v, x, y \in \hat{\mathcal{T}}$ also satisfy the criteria of the lemma; the claim follows by induction on the construction of $\hat{\mathcal{T}}$. Thus, we can assume v is a leaf of $\hat{\mathcal{T}}$ and u is a leaf of $\mathcal{T}$. See Figure 3. By construction, all global symbols in $\mathcal{B}(u)$ also appear in $\mathcal{B}(v)$. Since $x \in \hat{\mathcal{T}}$, the symbol a is classified as global and must appear in $\mathcal{B}(v)$. □

4.2. Nesting Habits

Suppose a block β in S contains two symbols a, b that are not wingtips, that is, they make neither their first nor last appearance in β. We call a and b *nested in* β if S contains either $ab\,\beta\,ba$ or $ba\,\beta\,ab$ and call them *interleaved in* β otherwise, that is, if

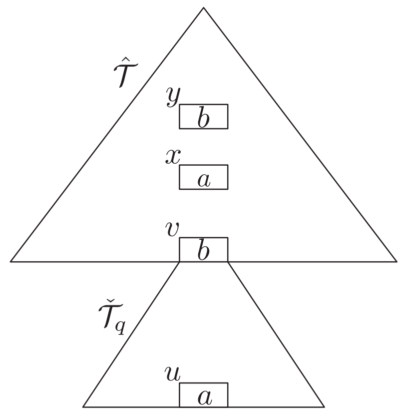

Fig. 3. The case where v is a leaf of $\hat{\mathcal{T}}$. Both x and y are necessarily in $\hat{\mathcal{T}}$, which implies that a and b are global, and further implies that u is a leaf of $\check{\mathcal{T}}_q$ since global symbols do not appear in the internal nodes of $\check{\mathcal{T}}_q$. All global symbols of u also appear in v.

the occurrences of a and b in S take the form $a^*b^*\ \beta\ a^*b^*$ or $b^*a^*\ \beta\ \ b^*a^*$. Lemma 4.4 is the critical structural lemma used in our analysis. It provides us with simple criteria for nestedness.

LEMMA 4.4. *Suppose that $v \in \mathcal{T}(S)$ is a leaf and a, b are symbols in a block $\mathcal{B}(v)$ of S. If the following two criteria are satisfied, then a and b are nested in $\mathcal{B}(v)$.*

(*1*) *v is not a wingtip in either $\mathcal{T}_{|a}$ or $\mathcal{T}_{|b}$.*
(*2*) *v is not a feather in either $\mathcal{T}_{|a}$ or $\mathcal{T}_{|b}$.*

PROOF. By symmetry, we can assert two additional criteria.

(3) $\mathrm{cr}_{|b}$ is equal to or strictly ancestral to $\mathrm{cr}_{|a}$.
(4) v is a dove in $\mathcal{T}_{|a}$.

According to Criteria (1, 2, 4), v is distinct from $\mathrm{lwt}_{|a}$ and $\mathrm{fe}_{|a}(v)$. They are all descendants of $\mathrm{wi}_{|a}(v)$ and appear in the order $\mathrm{lwt}_{|a}$, v, $\mathrm{fe}_{|a}(v)$.

We partition the sequence S outside of $\mathcal{B}(v)$ into the following four intervals.

—I_1: everything preceding the a in $\mathcal{B}(\mathrm{lwt}_{|a})$.
—I_2: everything from the end of I_1 to $\mathcal{B}(v)$.
—I_3: everything from $\mathcal{B}(v)$ to the a in $\mathcal{B}(\mathrm{fe}_{|a}(v))$.
—I_4: everything following I_3.

Since v is not a wingtip of $\mathcal{T}_{|b}$, there must be occurrences of b in S both before and after $\mathcal{B}(v)$. If, contrary to the claim, a and b are not nested in $\mathcal{B}(v)$, all other occurrences of b must appear exclusively in I_1 and I_3 or exclusively in I_2 and I_4. We show that both possibilities lead to contradictions. Figures 4 and 5 illuminate the proof.

Case 1: b does not appear in I_1 or I_3. Since v is not a wingtip in $\mathcal{T}_{|b}$ (Criterion (1)), the left wingtip $\mathrm{lwt}_{|b}$ of $\mathcal{T}_{|b}$ appears in interval I_2. Since $\mathrm{lwt}_{|b}$ and v are descendants of $\mathrm{wi}_{|a}(v)$, which is a strict descendant of $\mathrm{cr}_{|a}$, which, by Criterion (3), is a descendant of $\mathrm{cr}_{|b}$, it must also be that $\mathrm{lwt}_{|b}$ and v descend from the same child of $\mathrm{cr}_{|b}$, that is,

(5) v is a dove in $\mathcal{T}_{|b}$.

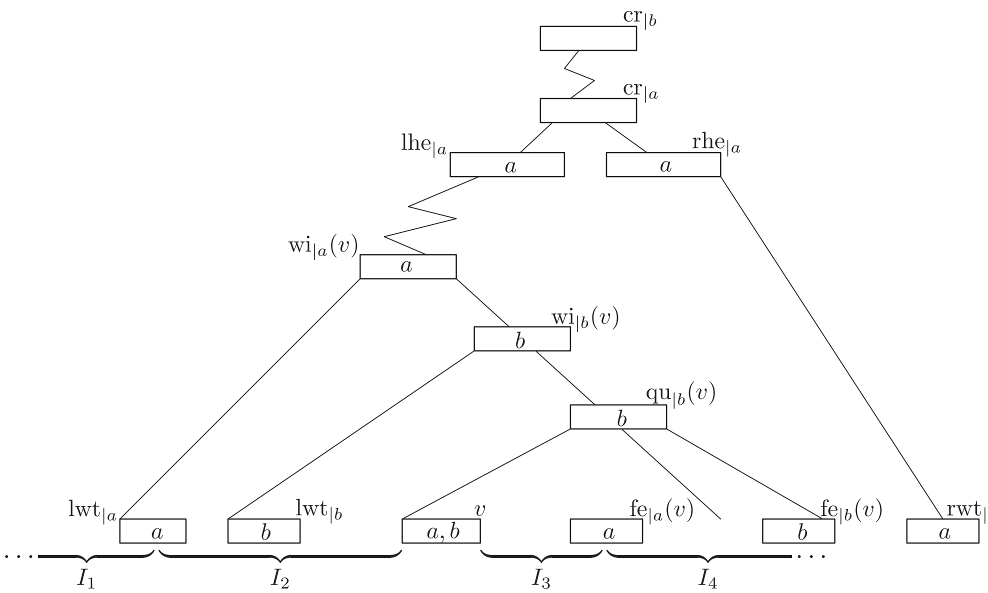

Fig. 4. Boxes represent nodes in $\mathcal{T}(S)$ and their associated blocks. The blocks at the leaf-level correspond to those in S. In Case 1, all occurrences of b outside of $\mathcal{B}(v)$ appear in intervals I_2 and I_4. Contrary to the depiction, it may be that $\mathrm{cr}_{|a}$ and $\mathrm{cr}_{|b}$ are identical, that $\mathrm{lwt}_{|a}$ and $\mathrm{lwt}_{|b}$ are identical, that $\mathrm{fe}_{|a}(v)$ and $\mathrm{fe}_{|b}(v)$ are identical, and that $\mathrm{wi}_{|b}(v)$ is not a descendant of $\mathrm{wi}_{|a}(v)$.

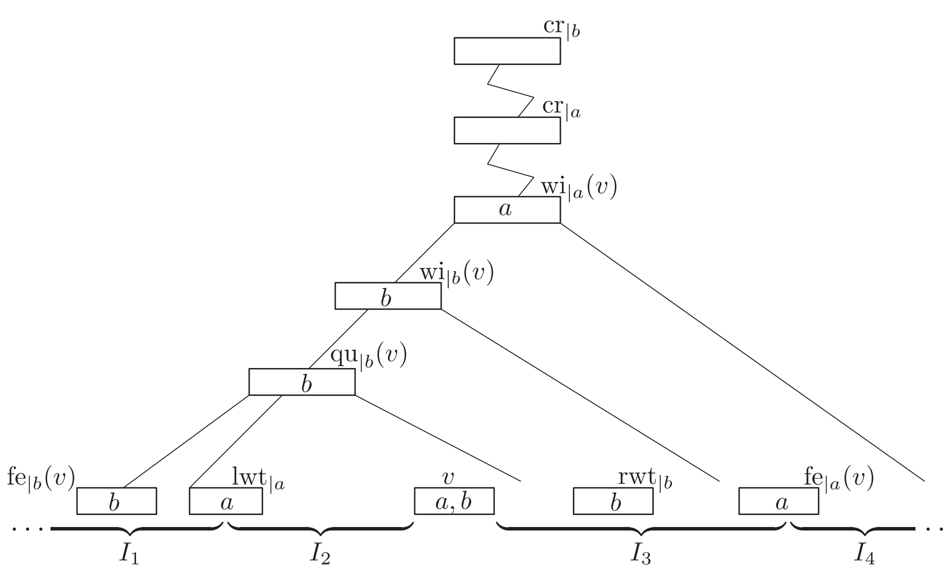

Fig. 5. In Case 2, all occurrences of b outside of $\mathcal{B}(v)$ appear in intervals I_1 and I_3.

We shall next argue the following.

(6) In $\mathcal{T}$, $\mathrm{qu}_{|b}(v)$ is a strict descendant of $\mathrm{wi}_{|a}(v)$ and a strict ancestor of $\mathrm{fe}_{|a}(v)$.
(7) $\mathrm{fe}_{|b}(v)$ lies in interval I_4.

The least common ancestor of v and $\mathrm{lwt}_{|b}$ in $\mathcal{T}_{|b}$ is by definition $\mathrm{wi}_{|b}(v)$. The quill $\mathrm{qu}_{|b}(v)$ is a child of $\mathrm{wi}_{|b}(v)$ not on a wing, hence $\mathrm{qu}_{|b}(v)$ cannot be ancestral to $\mathrm{lwt}_{|b}$ and therefore must be a strict descendant of $\mathrm{wi}_{|a}(v)$. Since v is a non-feather dove in $\mathcal{T}_{|b}$ (by Criterion (1) and Inference (5)), $\mathrm{fe}_{|b}(v)$ is the rightmost leaf descendant of $\mathrm{qu}_{|b}(v)$ and distinct from v. However, by supposition, I_3 contains no occurrences of b, so $\mathrm{fe}_{|b}(v)$ must lie in interval I_4. For $\mathrm{qu}_{|b}(v)$ to have descendants in both I_2 and I_4, it must be a strict ancestor of $\mathrm{fe}_{|a}(v)$ in $\mathcal{T}$. As we explain, a consequence of Inference (6) is as follows.

(8) $\mathrm{rwt}_{|a}$ lies to the right of $\mathrm{fe}_{|b}(v)$.

According to Criterion (4) and Inference (6), $\mathrm{qu}_{|b}(v)$ is a descendant of $\mathrm{wi}_{|a}(v)$, which is a descendant of $\mathrm{lhe}_{|a}$. Since $\mathrm{rwt}_{|a}$ is a descendant of $\mathrm{rhe}_{|a}$, the right sibling of $\mathrm{lhe}_{|a}$, $\mathrm{rwt}_{|a}$ must lie to the right of $\mathrm{fe}_{|b}(v)$.

Let us review the situation. Scanning the leaves from left to right, we see the blocks $\mathrm{lwt}_{|a}$, $\mathrm{lwt}_{|b}$, v, $\mathrm{fe}_{|a}(v)$, $\mathrm{fe}_{|b}(v)$, and $\mathrm{rwt}_{|a}(v)$. It may be that $\mathrm{lwt}_{|a}$ and $\mathrm{lwt}_{|b}$ are equal and it may be that $\mathrm{fe}_{|a}(v)$ and $\mathrm{fe}_{|b}(v)$ are equal. If either of these cases hold, then the a precedes the b in the given block. The blocks $\mathrm{lwt}_{|a}$, $\mathrm{lwt}_{|b}$, v, $\mathrm{fe}_{|b}(v)$, $\mathrm{rwt}_{|a}(v)$ certify that a and b are nested in $\mathcal{B}(v)$.

Case 2: b does not appear in I_2 or I_4. The right wingtip $\mathrm{rwt}_{|b}$ is distinct from v, by Criterion (1) and must therefore lie in I_3. Following the same reasoning from Case 1, we can deduce the following.

(9) v is a hawk in $\mathcal{T}_{|b}$.
(10) In $\mathcal{T}$, $\mathrm{qu}_{|b}(v)$ is a strict descendant of $\mathrm{wi}_{|a}(v)$ and a strict ancestor of $\mathrm{lwt}_{|a}$.

Inference (9) follows since v and $\mathrm{rwt}_{|b}$ must be descendants of the same head in $\mathcal{T}_{|b}$. This implies that $\mathrm{fe}_{|b}(v)$ is the leftmost leaf descendant of $\mathrm{qu}_{|b}(v)$. Since $\mathrm{fe}_{|b}(v)$ is distinct from v and interval I_2 is free of bs, it must be that $\mathrm{fe}_{|b}(v)$ lies in I_1 and therefore that $\mathrm{qu}_{|b}(v)$ is a strict descendant of $\mathrm{wi}_{|a}(v)$ and a strict ancestor of $\mathrm{lwt}_{|a}$. Inference (10) follows. See Figure 5.

It follows from Criterion (3) and Inference (10) that on a leaf-to-root path, one encounters distinct nodes $\mathrm{lwt}_{|a}$, $\mathrm{qu}_{|b}(v)$, $\mathrm{wi}_{|a}(v)$, and $\mathrm{cr}_{|b}$, in that order. Lemma 4.3 implies that $a \in \mathcal{B}(\mathrm{qu}_{|b}(v))$. We have deduced that $\mathrm{qu}_{|b}(v)$ is in $\mathcal{T}_{|a}$, is a strict descendant of $\mathrm{wi}_{|a}(v)$, and is ancestral to both $\mathrm{lwt}_{|a}$ and v. This contradicts the fact that $\mathrm{wi}_{|a}(v)$ is the "least" common ancestor of v and $\mathrm{lwt}_{|a}$ in $\mathcal{T}_{|a}$. $\square$

Note that Lemma 4.4 applies to any blocked sequence and an associated derivation tree. It has nothing to do with Davenport-Schinzel sequences as such.

5. A RECURRENCE FOR ODD ORDERS

Lemma 4.4 may be rephrased as follows. Every blocked sequence S is the union of four sequences: two comprising wingtips (first occurrences and last occurrences, each of length n), one comprising all feathers, and one comprising non-wingtip non-feathers. The last sequence is distinguished by the property that each pair of symbols in any block is nested with respect to S, which is a "good" thing if we are intent on giving strong upper bounds on odd-order sequences. The sequence comprising feathers is "bad" in this senses, therefore, we must obtain better-than-trivial upper bounds on its length if this strategy is to bear fruit.

Recall that the definition of a *feather* is not absolute: it is with respect to a derivation tree $\mathcal{T}$, that is, with respect to some strategy for choosing block partitions. To obtain good bounds on the number of feathers we only consider a specific way to choose block partitions.

Canonical Derivation Trees. Define the *canonical* derivation tree $\mathcal{T}^\star(S)$ of an m-block sequence as follows. Choose $\hat{m} = \lceil m/2 \rceil$ and $m_q = 2$ for all $q < \hat{m}$. The derivation trees

$\{\check{\mathcal{T}}_q\}$ of $\{\check{S}_q\}$ are necessarily two-leaf base case trees. Generate the canonical derivation tree $\hat{\mathcal{T}} = \mathcal{T}^\star(\hat{S}')$ inductively and, as usual, let $\mathcal{T}^\star(S)$ be the composition of $\hat{\mathcal{T}}$ and the local trees $\{\check{\mathcal{T}}_q\}$, placing the blocks of S at the leaves of $\mathcal{T}^\star(S)$.

The tree $\mathcal{T}^\star(S)$ can also be defined non-inductively as the unique tree satisfying the following three criteria. Structurally, it must be a fragment of a full binary tree with height $\lceil \log[\![S]\!] \rceil$. The crown $\mathrm{cr}_{|a}$ is the least common ancestor of all leaves whose blocks contain a. If $a \in \mathcal{B}(v)$ and v is a leaf, a also appears in the blocks of every ancestor strictly between v and $\mathrm{cr}_{|a}$. With this definition $\mathcal{T}^\star_{|a}$ is binary: all nodes have one or two children. The branching nodes on the wings of $\mathcal{T}^\star_{|a}$ are associated with exactly one quill and one feather. Bounding the number of feathers is therefore tantamount to bounding the number of branching wing nodes.

Recurrence 5.1. *Let S be an m-block, order-s DS sequence over an n-letter alphabet. Define $\Phi_s(n, m)$ to be the maximum number of feathers of one type (dove or hawk) in such a sequence, where* feather *is with respect to $\mathcal{T}^\star(S)$. When $m = 2$ or $s = 2$, we have*

$$\Phi_s(n, 2) = 0$$
$$\Phi_2(n, m) < m.$$

Let $\{m_q\}_{1 \le q \le \hat{m}}$ be any block partition in which $m_1 = \cdots = m_{\hat{m}-1}$ are equal powers of two and $m_{\hat{m}}$ may be smaller. Let $\{\hat{n}\} \cup \{\check{n}_q\}_{1 \le q \le \hat{m}}$ be the corresponding alphabet partition. Then,

$$\Phi_s(n, m) \le \sum_{q=1}^{\hat{m}} \Phi_s(\check{n}_q, m_q) + \Phi_s(\hat{n}, \hat{m}) + \Phi_{s-1}(\hat{n}, m) + \hat{n}.$$

Proof. Suppose we only wish to bound the number of dove feathers with respect to $\mathcal{T}^\star = \mathcal{T}^\star(S)$. If there are only two blocks, then all occurrences are wingtips and feathers are not wingtips. This gives the first equality. In the most extreme case, every non-wingtip is a dove feather, so $\Phi_s(n, m) \le \lambda_s(n, m) - 2n$. In particular, $\Phi_2(n, m) \le \lambda_2(n, m) - 2n < m$. Decompose S into $\hat{S}, \hat{S}', \acute{S}_q, \grave{S}_q, \bar{S}_q$ in the usual way with respect to the given block partition. Define $\hat{\mathcal{T}} = \mathcal{T}^\star(\hat{S}')$ to be the canonical derivation tree of the contracted global sequence $\hat{S}'$ and define $\acute{\mathcal{T}}_q = \mathcal{T}^\star(\acute{S}_q)$ to be the canonical derivation tree of the global *first occurrences* in S_q. Since we forced $m_1, \ldots, m_{\hat{m}-1}$ to be equal powers of two, $\hat{S}'$ occupies a level in $\mathcal{T}^\star$. Thus, both $\hat{\mathcal{T}}_{|a}$ and $(\acute{\mathcal{T}}_q)_{|a}$ (where $a \in \Sigma(\acute{S}_q)$) are contained as subtrees in $\mathcal{T}^\star_{|a}$.

The branching nodes on the left wing of $\mathcal{T}^\star_{|a}$, where $a \in \Sigma(\acute{S}_q)$, consist of (i) the branching nodes on the left wing of $\hat{\mathcal{T}}_{|a}$, (ii) the branching nodes on the left wing of $(\acute{\mathcal{T}}_q)_{|a}$, and (iii) the crown $\acute{\mathrm{cr}}_{|a}$ of $(\acute{\mathcal{T}}_q)_{|a}$, which is on the left wing of $\mathcal{T}^\star_{|a}$ but not $(\acute{\mathcal{T}}_q)_{|a}$. See Figure 6. Each branching node is identified with one feather in $\mathcal{T}^\star_{|a}$. The total number of branching nodes/feathers covered by (i), summed over all $a \in \Sigma(\hat{S})$, is at most $\Phi_s(\hat{n}, \hat{m})$. The total number covered by (ii), summed over all $q \le \hat{m}$ and all $a \in \Sigma(\acute{S}_q)$, is $\sum_q \Phi_{s-1}(\acute{n}_q, m_q) \le \Phi_{s-1}(\hat{n}, m)$. (Remember that $\acute{S}_q$ is an order-$(s-1)$ DS sequence.) The number covered by (iii) is clearly $\hat{n}$, which gives the last inequality. □

We now have all the elements in place to provide a recurrence for odd-order Davenport-Schinzel sequences.

Recurrence 5.2. *Let m, n, and s be the block count, alphabet size, and order parameters, where $s \ge 5$ is odd. For any block partition $\{m_q\}_{1 \le q \le \hat{m}}$ and the corresponding*

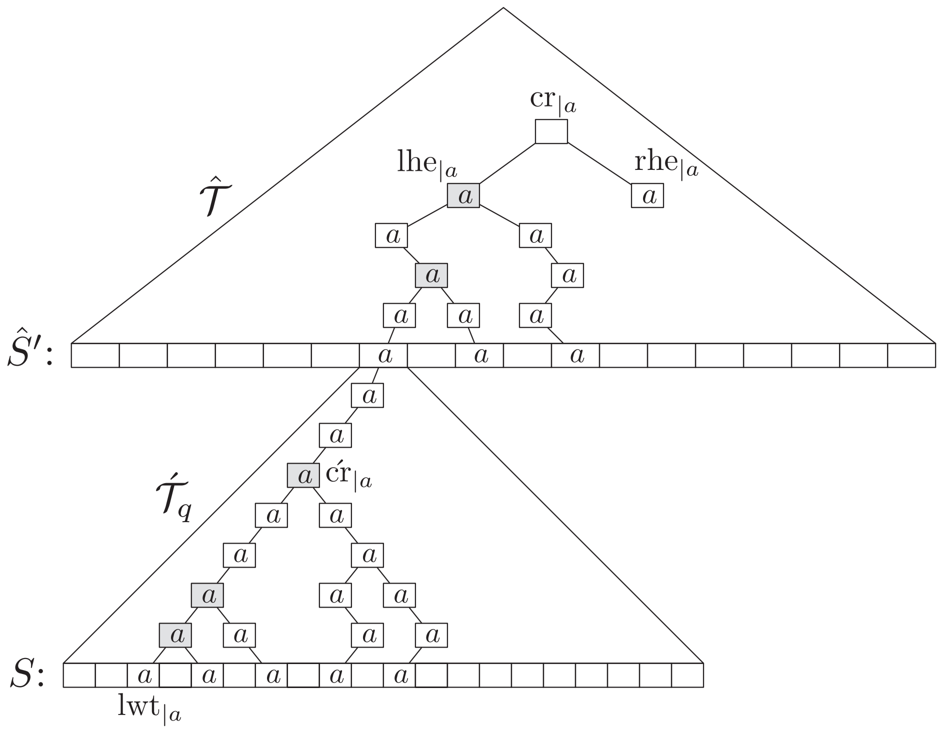

Fig. 6. Counting dove feathers in $\mathcal{T}^{\star}_{|a}$ is tantamount to counting branching nodes on the left wing of $\mathcal{T}^{\star}_{|a}$. Branching nodes are shaded.

alphabet partition $\{\hat{n}\} \cup \{\check{n}_q\}_{1\le q\le \hat{m}}$, *we have*

$$\lambda_s(n, m) \le \sum_{q=1}^{\hat{m}} \lambda_s(\check{n}_q, m_q) \;+\; 2\cdot\lambda_{s-1}(\hat{n}, m) \;+\; \lambda_{s-2}(2\cdot\Phi_s(\hat{n}, \hat{m}), m) \;+\; \lambda_{s-3}(\lambda_s(\hat{n}, \hat{m}), m).$$

PROOF. Define $\hat{\mathcal{T}}$ as in the proof of Recurrence 5.1. In Recurrence 3.3, we partitioned S into local and global symbols and partitioned the occurrences of global symbols into first, middle, and last. We now partition the middle occurrences one step further. Define $\tilde{S}'$ and $\breve{S}'$ to be the subsequences of $\hat{S}'$ consisting of feathers (according to $\hat{\mathcal{T}}$) and non-feather, non-wingtips, respectively. That is, $|\hat{S}'| = |\tilde{S}'| + |\breve{S}'| + 2\hat{n}$. In an analogous fashion, define $\tilde{S}$ and $\breve{S}$ to be the subsequences of $\hat{S}$ begat by occurrences in $\tilde{S}'$ and $\breve{S}'$. The sequences $\acute{S}$ and $\grave{S}$ consist of occurrences begat by dove and hawk wingtips in $\hat{\mathcal{T}}$. Thus, $|S| = \sum_q |\check{S}_q| + |\acute{S}| + |\grave{S}| + |\tilde{S}| + |\breve{S}|$.

The local sequences $\{\check{S}_q\}$ are order-s DS sequences. According to the standard argument, $\acute{S}$ and $\grave{S}$ are order-$(s-1)$ DS sequences and $\tilde{S} = \tilde{S}_1\cdots\tilde{S}_{\hat{m}}$ is obtained from $\tilde{S}'$ by substituting for its qth block an order-$(s-2)$ DS sequence $\tilde{S}_q$. From the superadditivity of λ_{s-2}, it follows that $|\tilde{S}| \le \lambda_{s-2}(|\tilde{S}'|, m) \le \lambda_{s-2}(2\cdot\Phi_s(\hat{n}, \hat{m}), m)$. (Recall that $\Phi_s(\hat{n}, \hat{m})$ is the number of feathers of one type, so $2\cdot\Phi_s(\hat{n}, \hat{m})$ bounds the total number of feathers.)

We claim that $\breve{S} = \breve{S}_1\cdots\breve{S}_{\hat{m}}$ is obtained from $\breve{S}'$ by substituting for its qth block an order-$(s-3)$ DS sequence $\breve{S}_q$, which, if true, would imply that $|\breve{S}| \le \lambda_{s-3}(|\breve{S}'|, m) < \lambda_{s-3}(\lambda_s(\hat{n}, \hat{m}), m)$. Suppose for the purpose of obtaining a contradiction that the qth block

β in $\breve{S}'$ contains $a, b \in \hat{\Sigma}$, and that $\breve{S}_q$ is not an order-$(s-3)$ DS sequence, that is, it contains an alternating subsequence $ab\cdots ab$ of length $s-1$. Note that $s-1$ is even. By definition β is a non-feather, non-wingtip in both $\hat{\mathcal{T}}_{|a}$ and $\hat{\mathcal{T}}_{|b}$. According to Lemma 4.4, a and b must be nested in β, which implies that S contains a subsequence of the form

$$a\cdots b\cdots \Big|\cdots \overbrace{a\cdots b\cdots a\cdots b}^{s-1}\cdots \Big|\cdots b\cdots a$$

$$\text{or}\qquad b\cdots a\cdots \Big|\cdots \overbrace{a\cdots b\cdots a\cdots b}^{s-1}\cdots \Big|\cdots a\cdots b$$

where the portion between bars is in S_q. In either case, S contains an alternating subsequence with length $s+2$, contradicting the fact that S is an order-s DS sequence. □

5.1. Analysis of the Recurrences

The dependencies between λ and Φ established by Recurrences 3.3, 5.1, and 5.2 are rather intricate. For even s, λ_s is a function of λ_s, λ_{s-1} and λ_{s-2}, and for odd s, λ_s is a function of $\Phi_s, \lambda_s, \lambda_{s-1}, \lambda_{s-2}$, and λ_{s-3}, while Φ_s is a function of Φ_s and Φ_{s-1}.

The proof of Lemma 5.3 is by induction over parameters: s, n, c, i, and j, where s is the order, n the alphabet size, $c \geq s-2$ a constant that determines how $\hat{m}$ and the block partition is chosen, $i \geq 1$ is an integer, and j is minimal such that the block count $m \leq a^c_{i,j}$. Some level of complexity is therefore unavoidable. Furthermore, when $s \geq 5$ is odd, λ_s is so sensitive to approximations of λ_{s-3} that we must treat $s \in \{1,2,3,4,5\}$ as distinct base cases and treat even and odd $s \geq 6$ as separate inductive cases. Given these constraints, we feel our analysis is reasonably simple.

LEMMA 5.3. *Let $s \geq 1$ be the order parameter, $c \geq s-2$ be a constant, and $i \geq 1$ be an arbitrary integer. Define j to be maximum such that $m \leq a^c_{i,j}$. The following upper bounds on λ_s and Φ_s hold for all $s \geq 1$.*

$$\begin{aligned}
\lambda_1(n,m) &= n+m-1 && s=1\\
\lambda_2(n,m) &= 2n+m-2 && s=2\\
\lambda_3(n,m) &\leq (2i+2)n+(3i-2)cj(m-1) && s=3\\
\lambda_s(n,m) &\leq \mu_{s,i}(n+(cj)^{s-2}(m-1)) && \textit{all } s\geq 4\\
\Phi_s(n,m) &\leq \nu_{s,i}(n+(cj)^{s-2}(m-1)) && \textit{all } s\geq 5.
\end{aligned}$$

The values $\{\mu_{s,i}, \nu_{s,i}\}$ are defined as follows, where $t = \lfloor\frac{s-2}{2}\rfloor$ and C is an absolute constant.

$$\mu_{s,i} = \begin{cases} 2^{\binom{i+t+C}{t}} - 6(i+2) & \textit{even } s\geq 4\\[1ex] 3i\cdot 2^{\binom{i+t+C}{t}} & \textit{odd } s\geq 5\end{cases}$$

$$\nu_{s,i} = \binom{i+s-2}{s-2} - 1 \qquad \textit{all } s\geq 5.$$

One may want to keep in mind that we will eventually substitute $\alpha(n,m)+O(1)$ for the parameter i, and that $\binom{i+t+C}{t} = i^t/t! + O(i^{t-1})$. Lemma 5.3 will, therefore, imply bounds on $\lambda_s(n,m)$ analogous to those claimed for $\lambda_s(n)$ in Theorem 1.3. The proof of Lemma 5.3 appears in Appendix B.

5.2. The Upper Bounds of Theorem 1.3

Fix $s \geq 3, n, m$ and let $c = s-2$. For $i \geq 1$, let j_i be minimum such that $m \leq a^c_{i,j_i}$. Lemma 5.3 implies that an order-s DS sequence has length at most $\mu_{s,i}(n + (cj_i)^{s-2}m)$. Choose ι to be minimum such that[18] $(cj_\iota)^{s-2} \leq \max\{\frac{n}{m}, (c\cdot 3)^{s-2}\}$. One can show that $\iota = \alpha(n,m) + O(1)$. By choice of ι, it follows that $(cj_\iota)^{s-2}m = O(m+n)$, so $\lambda_s(n,m) = O(\mu_{s,\iota}(n+m))$. According to Lemma 5.3's definition of $\mu_{s,\iota}$, we have

$$\lambda_3(n,m) = O((n+m)\alpha(n,m))$$
$$\lambda_4(n,m) = O\big((n+m)2^{\alpha(n,m)}\big)$$
$$\lambda_5(n,m) = O\big((n+m)\alpha(n,m)2^{\alpha(n,m)}\big)$$
$$\lambda_s(n,m) = (n+m)\cdot 2^{\alpha^t(n,m)/t! + O(\alpha^{t-1}(n,m))} \quad \text{both even and odd } s \geq 6, \text{ where } t = \lceil \tfrac{s-2}{2} \rceil.$$

The bound on $\lambda_5(n,m)$ follows since $\mu_{5,\iota} = O(\iota 2^\iota)$. When $s \geq 6$ and $t \geq 2$, $\mu_{s,\iota} = O(\iota 2^{\binom{\iota+t+C}{t}}) = 2^{\iota^t/t! + O(\iota^{t-1})}$.

Theorem 1.3 stated bounds on $\lambda_s(n)$ rather than $\lambda_s(n,m)$. If it were known that extremal order-s DS sequences consisted of $m = O(n)$ blocks, we could simply substitute $\alpha(n)$ for $\alpha(n,m)$ in the preceding bounds, but this is not known to be true. According to Lemma 3.1(2,4), if γ_s is such that $\lambda_s(n) \leq \gamma_s(n)\cdot n$, then $\lambda_s(n) \leq \gamma_{s-2}(n)\cdot\lambda_s(n, 2n-1)$ and $\lambda_s(n) \leq \gamma_{s-2}(\gamma_s(n))\cdot\lambda_s(n, 3n-1)$. Applying Lemma 3.1 when $s \in \{3,4\}$ has no asymptotic affect, since $\gamma_1 = 1$ and $\gamma_2 = 2$. It has no *perceptible* effect when $s \geq 6$, since $\gamma_{s-2}(n)$ or $\gamma_{s-2}(\gamma_s(n))$ is dwarfed by the lower order terms in the exponent. However, for $s \in \{3,5\}$, these reductions only show that $\lambda_3(n) = O(n\alpha(n))$ and that $\lambda_5(n) = O(n\alpha(\alpha(n))\alpha(n)2^{\alpha(n)})$, which are weaker than the bounds claimed in Theorem 1.3.

In Section 6, we prove the remaining upper bounds of Theorem 1.3: $\lambda_3(n) = 2n\alpha(n) + O(n)$ and $\lambda_5(n) = \Theta(n\alpha(n)2^{\alpha(n)})$. The bound on λ_3 is a tiny improvement over Klazar's bound [1999], though it is within $O(n)$ of the construction of Nivasch [2010] and is therefore optimal in the Ackermann-invariant sense. See Remark 1.1. Section 7 gives a matching lower bound on order-5 DS sequences.

6. SHARP BOUNDS AT ORDERS 3 AND 5

6.1. Order-3 DS Sequences

Let S be an order-3 DS sequence over an n-letter alphabet. According to Lemma 3.1, $|S| \leq \lambda_3(n) \leq \lambda_3(n,m)$, where $m = 2n-1$. Letting ι be minimum such that $m \leq a_{\iota,3}$, Lemma 5.3 implies that $\lambda_3(n,m) < (2\iota+2)n + (3\iota-2)m < (8\iota-2)n$. It is straightforward to show that $\iota \leq \alpha(n) + O(1)$. The problem is clearly that there are too many blocks. Were there fewer than $(2n-1)/\iota$ blocks, Lemma 5.3 would give a bound of $(2\iota+2)n + O(\iota m/\iota) = 2n\alpha(n) + O(n)$. We can invoke Recurrence 3.3 to divide S into a global $\hat{S}$ and local $\check{S} = \check{S}_1\cdots\check{S}_{\hat{m}}$, where $\hat{m} = m/\iota \leq (2n-1)/\iota$, that is, each $\check{S}_q$ is an ι-block sequence. Using Lemma 5.3, we will bound $\hat{S}$ with $i = \iota$ and each of the $\{\check{S}_q\}_q$

[18]We want $(cj_\iota)^{s-2}m$ not to be the dominant term, so $(cj_\iota)^{s-2}$ should be less than $\lceil n/m \rceil$. On the other hand, the first and second columns of Ackermann's function ($a_{i,1}$ and $a_{i,2}$) do not exhibit sufficient growth, so j_ι must also be at least 3.

with $i = 1$.

$$\begin{aligned}
|S| &\le \lambda_3(n) \le \lambda_3(n,m) && \{\text{where } m = 2n-1\}\\
&\le \textstyle\sum_{q=1}^{\hat{m}} \lambda_3(\check{n}_q, \iota) \;+\; 2\cdot\lambda_2(\hat{n},m) \;+\; \lambda_1(\lambda_3(\hat{n},\hat{m}) - 2\hat{n}, m) && \{\text{Recurrence 3.3}\}\\
&< \textstyle\sum_{q=1}^{\hat{m}} [4\check{n}_q \;+\; \min\{\iota\lceil\log\iota\rceil,\;\; (\iota-1) + (2\check{n}_q - 1)\lceil\log(2\check{n}_q - 1)\rceil\}] && (*)\\
&\quad +\; [4\hat{n} + 2m] \;+\; [2\iota\hat{n} \;+\; (3\iota-2)\hat{m} \;+\; m] && \{\text{Lemmas 3.2, 5.3}\}\\
&< \sum_{q=1}^{\hat{m}} [4\check{n}_q + \iota + 2\check{n}_q\lceil\log\iota\rceil] + [4\hat{n} + 2m] + [2\iota\hat{n} + 4m]. && (**)
\end{aligned}$$

The bound on local symbols in Line (*) follows from Lemma 5.3 and Hart and Sharir's [1986] observation that $\lambda_3(n) \le \lambda_3(n, 2n-1)$. When $i = 1$ and $j = \lceil\log\iota\rceil$, Lemma 5.3 gives us a bound of $\lambda_3(\check{n}_q, \iota) \le 4\check{n}_q + \iota\lceil\log\iota\rceil$. Alternatively, we could make $\check{S}_q$ 2-sparse by removing up to $\iota - 1$ duplicated symbols at block boundaries, then partition the remaining sequence into $2\check{n}_q - 1$ blocks, hence $\lambda_3(\check{n}_q, \iota) \le \iota - 1 + \lambda_3(\check{n}_q, 2\check{n}_q - 1) \le \iota - 1 + 4\check{n}_q + (2\check{n}_q - 1)\lceil\log(2\check{n}_q - 1)\rceil$. If $\iota < 2\check{n}_q$ we apply the first method, otherwise we apply the second method. The minimum of the two is therefore less than $4\check{n}_q + \iota + 2\check{n}_q\lceil\log\iota\rceil$, which justifies Line (**). Continuing with the inequalities,

$$\begin{aligned}
&< [m + (n-\hat{n})(4 + 2\lceil\log\iota\rceil)] + (2\iota+4)\hat{n} + 6m && \{\textstyle\sum_q \iota = m\},\\
&< (2\iota+4)n + 7m && \{\text{maximized when } \hat{n} = n\},\\
&\le 2n\alpha(n) + O(n) && \{\iota = \alpha(n) + O(1)\}.
\end{aligned}$$

This matches the lower bound of Nivasch [2010] on $\lambda_3(n)$ to within $O(n)$.

6.2. Order-5 DS Sequences

Lemma 5.3 states that for any i, $\lambda_5(n,m) < \mu_{5,i}(n + (3j_i)^3 m)$, where j_i is minimum such that $m \le a^3_{i,j_i}$. Choose $\iota \ge 1$ to be minimum such that $(3j_\iota)^3 \le \max\{\frac{n}{m}, (3\cdot 3)^3\}$. One can show that $\iota = \alpha(n,m) + O(1)$, implying that $\lambda_5(n,m) = O((n+m)\mu_{5,\iota}) = O((n+m)\alpha(n,m)2^{\alpha(n,m)})$, matching the construction from Section 7. According to Lemma 3.1(2,4), $\lambda_5(n) = O(\alpha(\alpha(n))) \cdot \lambda_5(n, 3n-1)$. In this section, we present a more efficient reduction from 2-sparse, order-5 DS sequences to blocked order-5 sequences, thereby removing the extra $\alpha(\alpha(n))$ factor.

THEOREM 6.1. *$\lambda_5(n) = O(n\alpha(n)2^{\alpha(n)})$ and $\lambda_5(n,m) = m + O(n\alpha(n,m)2^{\alpha(n,m)})$.*

PROOF. The second bound is asymptotically the same as $O((n+m)\alpha(n,m)2^{\alpha(n,m)})$ if $m = O(n)$. If not, we remove up to $m-1$ repeated symbols at block boundaries, yielding a 2-sparse, order-5 DS sequence. Our remaining task is therefore to prove that $\lambda_5(n) = O(n\alpha(n)2^{\alpha(n)})$.

Let S be a 2-sparse, order-5 DS sequence with $\|S\| = n$. Greedily partition S into maximal order-3 DS sequences $S_1 S_2 \cdots S_m$. According to Sharir's [1987] argument, $m < 2n$. See the proof of Lemma 3.1(2) in Appendix A. As usual, let $\check{S}, \hat{S} \prec S$ be the subsequences of local and global symbols, and let $\hat{S}'$ be derived by contracting each interval to a single block. The number of global symbols is $\hat{n} = \|\hat{S}\|$. In contrast to the situations we considered earlier, $\check{S}$ and $\hat{S}$ are neither 2-sparse nor partitioned into blocks.

Define $\hat{\mathcal{T}} = \mathcal{T}^\star(\hat{S}')$ to be the canonical derivation tree for $\hat{S}'$. Let $\tilde{S}', \check{S}' \prec \hat{S}'$ be the subsequences of feathers and non-feather, non-wingtips, respectively. Let $\acute{S}, \grave{S}, \tilde{S}, \check{S} \prec S$ be the subsequences of S begat symbols in $\hat{S}'$ categorized as dove wingtips, hawk wingtips, feathers, and non-feather, non-wingtips. Define $\acute{S}^*$, $\grave{S}^*$, $\tilde{S}^*$, and $\breve{S}^*$, to be their maximal length 2-sparse subsequences, that is, what remains after replacing runs $aaa\cdots a$ with a single a. Define $\check{S}^*$ to be the maximal length 2-sparse subsequence of $\check{S}$.

As we argue next, Lemma 3.1(2) and the arguments from Recurrence 5.1 imply that

$$\begin{aligned}|\acute{S}^*| + |\grave{S}^*| &\le 2\cdot\lambda_4(\hat{n}) \le 4\cdot\lambda_4(\hat{n},2\hat{n})\\ |\tilde{S}^*| &\le \lambda_3(\varphi) \le \lambda_3(\varphi,2\varphi) && \text{where } \varphi = 2\cdot\Phi_5(\hat{n},2n)\\ |\breve{S}^*| &\le \lambda_2(\lambda_5(\hat{n},2n)) \le 2\cdot\lambda_5(\hat{n},2n)\\ |\check{S}^*| &\le \lambda_3(n-\hat{n}) \le \lambda_3(n-\hat{n},2(n-\hat{n})).\end{aligned}$$

The sequences $\acute{S}^*$ and $\grave{S}^*$ are 2-sparse, order-4 DS sequences; hence their contribution is $2\cdot\lambda_4(\hat{n})$, which is less than $4\cdot\lambda_4(\hat{n},2\hat{n})$, by Lemma 3.1(2). The sequence $\tilde{S}^*$ is obtained by substituting for each block in $\tilde{S}'$ a 2-sparse, order-3 DS sequence. Since $|\tilde{S}^*| \le \varphi \stackrel{\text{def}}{=} 2\cdot\Phi_5(\hat{n},2n)$, by the superadditivity of λ_3, we have $|\tilde{S}^*| \le \lambda_3(\varphi)$, which is at most $\lambda_3(\varphi,2\varphi)$, by Lemma 3.1(2). Finally, $\breve{S}^*$ is obtained by substituting for each block in $\breve{S}'$ a 2-sparse, order-2 DS sequence. Since $|\breve{S}'| < \lambda_5(\hat{n},2n)$, we have $|\breve{S}^*| \le \lambda_2(\lambda_5(\hat{n},2n)) < 2\cdot\lambda_5(\hat{n},2n)$. We can also conclude that $|\check{S}^*| \le \lambda_3(n-\hat{n}) \le \lambda_3(n-\hat{n},2(n-\hat{n}))$.

In bounding these various sequences, the second argument of λ_s and Φ_s is never more than φ. Choose ι to be minimal such that $\varphi \le a^3_{\iota,3}$, so $j=3$ will be constant whenever we invoke Lemmas 5.3 with $s \le 5$ and $c=3$. It is straightforward to show that $\iota = \alpha(n)+O(1)$.

Observe that S can be constructed by shuffling its five non-2-sparse constituent subsequences $\check{S}$, $\acute{S}$, $\grave{S}$, $\tilde{S}$, $\breve{S}$ in some fashion that restores 2-sparseness. In other words, there is a 1-1 map between positions in S and positions in its five constituents, and a surjective map ψ from positions in S to positions in its 2-sparse constituents $\check{S}^*$, $\acute{S}^*$, $\grave{S}^*$, $\tilde{S}^*$, $\breve{S}^*$. Partition S into intervals $T_1T_2\cdots T_{\lceil |S|/h\rceil}$, each with length $h = \lceil\frac{\lambda_5(5)+1}{2}\rceil = O(1)$. The image of ψ on two consecutive intervals T_{p-1} and T_p (where $p < \lceil |S|/h\rceil$) cannot be identical, for otherwise $T_{p-1}T_p$ would be a 2-sparse, order-5 DS sequence with length $2h > \lambda_5(5)$ over a 5-letter alphabet, a contradiction. In other words, T_p must introduce a symbol from one of the five constituents subsequences that was not seen in T_{p-1}. It follows that

$$\begin{aligned}|S| &= |\check{S}| + |\acute{S}| + |\grave{S}| + |\tilde{S}| + |\breve{S}|\\ &\le h\cdot(|\check{S}^*| + |\acute{S}^*| + |\grave{S}^*| + |\tilde{S}^*| + |\breve{S}^*|)\\ &= h\cdot n\cdot O(\mu_{3,\iota} + 2\mu_{4,\iota} + 2\mu_{3,\iota}\nu_{5,\iota} + 2\mu_{5,\iota})\\ &= O(n\iota 2^\iota) \qquad \{\text{Since } \mu_{3,\iota} = O(\iota),\ \mu_{4,\iota} = O(2^\iota),\ \nu_{5,\iota} = O(\iota^3), \text{ and } \mu_{5,\iota} = O(\iota 2^\iota).\}\\ &= O(n\alpha(n)2^{\alpha(n)}). \quad \square\end{aligned}$$

7. LOWER BOUNDS ON ORDER-5 DS SEQUENCES

We have established every bound claimed in Theorem 1.3 except for the lower bound on order-5 DS sequences. In this section, we give a construction that yields bounds of $\lambda_5(n,m) = \Omega(n\alpha(n,m)2^{\alpha(n,m)})$ and $\lambda_5(n) = \Omega(n\alpha(n)2^{\alpha(n)})$. This is the first construction that is asymptotically longer than the order-4 DS sequences of Agarwal et al. [1989] having length $\Theta(n2^{\alpha(n)})$. Our construction is based on generalized forms of sequence *composition* and *shuffling* used by Agarwal et al. [1989], Nivasch [2010], and Pettie [2011b].

Recall from Section 1.1 that $\|S\| = |\Sigma(S)|$ is the alphabet size of S and, if S is partitioned into blocks, $[\![S]\!]$ is its block count.

7.1. Composition and Shuffling

In its generic form, a sequence S is assumed to be over the alphabet $\{1, \ldots, \|S\|\}$, that is, any totally ordered set with size $\|S\|$. To *substitute* S for a block $\beta = [a_1 \ldots a_{|\beta|}]$ means to replace β with a copy $S(\beta)$ under the alphabet mapping $k \mapsto a_k$, where $|\beta| \leq \|S\|$. If $|\beta|$ is strictly smaller than $\|S\|$, any occurrences of the $\|S\| - |\beta|$ unused symbols of $\Sigma(S)$ do not appear in $S(\beta)$. We always assume that S is in *canonical form*: the symbols are ordered according to the position of their first appearance in S.

Composition. If S_{mid} is a sequence in canonical form with $\|S_{\text{mid}}\| = j$ and S_{top} a sequence partitioned into blocks with length at most j, $S_{\text{sub}} = S_{\text{top}} \circ S_{\text{mid}}$ is obtained by substituting for each block β in S_{top} a copy $S_{\text{mid}}(\beta)$. Clearly, $[\![S_{\text{sub}}]\!] = [\![S_{\text{top}}]\!] \cdot [\![S_{\text{mid}}]\!]$. If S_{mid} and S_{top} contain μ and μ' occurrences of each symbol, respectively, then S_{sub} contains $\mu\mu'$ occurrences of each symbol. Composition preserves canonical form, that is, if S_{mid} and S_{top} are in canonical form, so is S_{sub}.

Shuffling. If S_{bot} is a j'-block sequence and S_{sub} is partitioned into blocks of length at most j', we can form the *shuffle* $S_{\text{sh}} = S_{\text{sub}} \diamond S_{\text{bot}}$ as follows. First create a sequence S^*_{bot} consisting of the concatenation of $[\![S_{\text{sub}}]\!]$ copies of S_{bot}, each copy being over an alphabet disjoint from the other copies and disjoint from that of S_{sub}. By design the length of S_{sub} is at most the number of blocks in S^*_{bot}, and precisely the same if all blocks in S_{sub} have their maximum length j'. The sequence $S_{\text{sub}} \diamond S_{\text{bot}}$ is obtained by shuffling the at most j' symbols of the lth block of S_{sub} into the j' blocks of the lth copy of S_{bot} in S^*_{bot}. Specifically, the kth symbol of the lth block is inserted at the end of the kth block of the lth copy of S_{bot}. If there is no kth symbol, then nothing is inserted into the kth block.

Three-Fold Composition. Our construction of order-5 DS sequences uses a generalized form of composition that treats symbols in β differently based on context. Suppose S_{top} is partitioned into blocks with length at most j and $S^{\text{f}}_{\text{mid}}$, $S^{\text{m}}_{\text{mid}}$, and $S^{\text{l}}_{\text{mid}}$ are sequences with alphabet size $\|S^{\text{f}}_{\text{mid}}\| = \|S^{\text{m}}_{\text{mid}}\| = \|S^{\text{l}}_{\text{mid}}\| = j$. The three-fold composition $S_{\text{top}} \circ \langle S^{\text{f}}_{\text{mid}}, S^{\text{m}}_{\text{mid}}, S^{\text{l}}_{\text{mid}}\rangle$ is formed as follows. For each block β in S_{top}, categorize its symbols as *first* if they occur in no earlier block, *last* if they occur in no later block, and *middle* otherwise. Let β^{f}, β^{m}, and β^{l} be the subsequences of β consisting of first, middle, and last symbols. These three sequences do not necessarily occur contiguously in β, but each is nonetheless a subsequence of β. Substitute for β the concatenation of $S^{\text{f}}_{\text{mid}}(\beta^{\text{f}})$, $S^{\text{m}}_{\text{mid}}(\beta^{\text{m}})$, and $S^{\text{l}}_{\text{mid}}(\beta^{\text{l}})$. This substitution preserves canonical form. Note that if S_{top}, $S^{\text{f}}_{\text{mid}}$, $S^{\text{m}}_{\text{mid}}$, and $S^{\text{l}}_{\text{mid}}$ contain $\mu \geq 2$, μ^{f}, μ^{m}, and μ^{l} occurrences of each symbol then $S_{\text{top}} \circ \langle S^{\text{f}}_{\text{mid}}, S^{\text{m}}_{\text{mid}}, S^{\text{l}}_{\text{mid}}\rangle$ contains $\mu^{\text{f}} + \mu^{\text{l}} + (\mu - 2)\mu^{\text{m}}$ occurrences of each symbol. Figure 7 gives a schematic of the generation of the sequence $(S_{\text{top}} \circ \langle S^{\text{f}}_{\text{mid}}, S^{\text{m}}_{\text{mid}}, S^{\text{l}}_{\text{mid}}\rangle) \diamond S_{\text{bot}}$.

7.2. Sequences of Orders 4 and 5

The sequences $S_4(i, j)$ and $S_5(i, j)$ are defined inductively. As we will prove, $S_4(i, j)$ is an order-4 DS sequence partitioned into blocks of length precisely j in which each symbol appears 2^i times, whereas $S_5(i, j)$ is an order-5 DS sequence partitioned into blocks of length at most j in which each symbol appears $(2i - 3)2^i + 4$ times. Let $B_s(i, j) = [\![S_s(i, j)]\!]$ and $N_s(i, j) = \|S_s(i, j)\|$ be, respectively, the number of blocks in $S_s(i, j)$ and the alphabet size of $S_s(i, j)$. By definition, $|S_4(i, j)| = 2^i \cdot N_4(i, j) = j \cdot B_4(i, j)$ and $|S_5(i, j)| = ((2i - 3)2^i + 4) \cdot N_5(i, j) \leq j \cdot B_5(i, j)$. The construction of S_4 is the same as Nivasch's [2010] and similar to that of Agarwal et al. [1989].

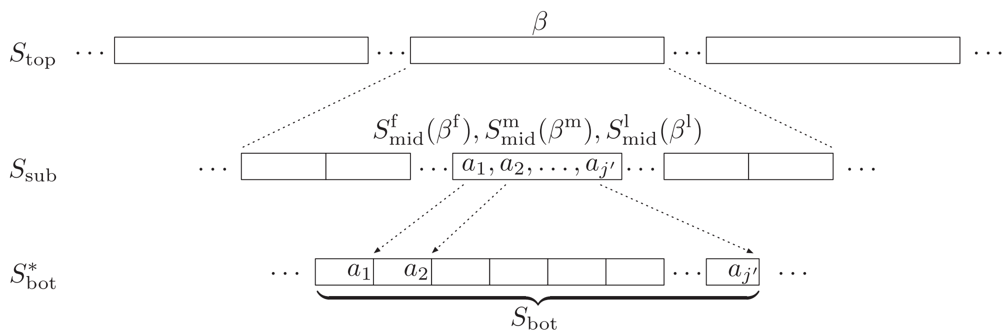

Fig. 7. Three-fold composition followed by shuffling. Each block β in $S_{\rm top}$ is replaced with the concatenation of $S^{\rm f}_{\rm mid}(\beta^{\rm f})$, $S^{\rm m}_{\rm mid}(\beta^{\rm m})$, and $S^{\rm l}_{\rm mid}(\beta^{\rm l})$, and each block of that sequence is shuffled with a single copy of $S_{\rm bot}$ in $S^*_{\rm bot}$. In general, blocks in $S^{\rm f}_{\rm mid}(\beta^{\rm f})$, $S^{\rm m}_{\rm mid}(\beta^{\rm m})$, and $S^{\rm l}_{\rm mid}(\beta^{\rm l})$ will not attain their maximum length j'.

The base cases for our sequences are given here, where square brackets indicate blocks:

$$\begin{aligned}
S_2(j) &= [12\cdots(j-1)j]\ \ [j(j-1)\cdots 21] && \text{two blocks with length } j,\\
S_4(1,j) &= S_5(1,j) = S_2(j) &&\\
S_4(i,1) &= [1]^{2^i} && 2^i \text{ identical blocks},\\
S_5(i,1) &= [1]^{(2i-3)2^i+4} && (2i-3)2^i+4 \text{ identical blocks}.
\end{aligned}$$

Observe that these base cases satisfy the property that symbols appear precisely 2^i times in $S_4(i,\cdot)$ and $(2i-3)2^i+4$ times in $S_5(i,\cdot)$. Define $S_4(i,j)$ as

$$S_4(i,j) = (S_4(i-1,y)\circ S_2(y))\diamond S_4(i,j-1), \qquad \text{where } y = B_4(i,j-1),$$

and $S_5(i,j)$ as

$$\begin{aligned}
&S_5(i,j) = \left(S_{\rm top}\circ\left\langle S^{\rm f}_{\rm mid}, S^{\rm m}_{\rm mid}, S^{\rm l}_{\rm mid}\right\rangle\right)\diamond S_{\rm bot},\\
&\text{where } \ S_{\rm bot} = S_5(i,j-1), && z = B_5(i,j-1)\\
&\qquad S^{\rm f}_{\rm mid} = S^{\rm l}_{\rm mid} = S_4(i,z),\\
&\qquad S^{\rm m}_{\rm mid} = S_2(N_4(i,z)),\\
&\text{and } \ S_{\rm top} = S_5(i-1, N_4(i,z)).
\end{aligned}$$

By definition, $S^{\rm f}_{\rm mid}$ and $S^{\rm l}_{\rm mid}$ are partitioned into blocks with length z. In the three-fold composition operation, we also interpret $S_2(N_4(i,z))$ as a sequence of blocks of length precisely z.[19] We argue by induction that symbols appear with the correct multiplicity in S_4 and S_5. In the case of S_4, each symbol appears 2^{i-1} times in $S_4(i-1,y)$ (by the inductive hypothesis), twice in $S_2(y)$, and therefore 2^i times in $S_4(i-1,y)\circ S_2(y)$. Symbols in copies of $S_4(i,j-1)$ already appear 2^i times by the inductive hypothesis. In $S_5(i-1,N_4(i,z))$, each symbol appears $(2i-5)2^{i-1}+4$ times. The three-fold composition operation increases the multiplicity of such symbols to $2((2i-5)2^{i-1}+2)+2(2^i) = (2i-3)2^i+4$, where the first term accounts for the blowup in middle occurrences and the second term for the blowup in first and last occurrences. It follows that B and N

[19]This requires that $N_4(i,z)$ is a multiple of z. By induction, $B_4(i,z)$ is a multiple of 2^i. Since $2^i\cdot N_4(i,z) = z\cdot B_4(i,z)$, $N_4(i,z)$ must be a multiple of z.

are defined inductively as follows.

$$\begin{aligned}
B_4(1,j) &= B_5(1,j) = B_2(j) = 2 \\
B_4(i,1) &= 2^i \\
B_5(i,1) &= (2i-3)2^i + 4 \\
B_4(i,j) &= B_4(i-1,y)\cdot 2\cdot y && \text{where } y = B_4(i,j-1) \\
B_5(i,j) &= B_5(i-1,N_4(i,z))\cdot(2+2^{-i+1})\cdot B_4(i,z)\cdot z && \text{where } z = B_5(i,j-1) \\
N_4(1,j) &= N_5(1,j) = N_2(j) = j \\
N_4(i,1) &= N_5(i,1) = 1 \\
N_4(i,j) &= N_4(i-1,y) \;+\; B_4(i-1,y)\cdot 2\cdot N_4(i,j-1) \\
N_5(i,j) &= N_5(i-1,N_4(i,z)) \;+\; B_5(i-1,N_4(i,z))\cdot(2+2^{-i+1})\cdot B_4(i,z)\cdot N_5(i,j-1).
\end{aligned}$$

The $2+2^{-i+1}$ factor in the definition of $B_5(i,j)$ and $N_5(i,j)$ comes from the fact that in the shuffling step, $S_2(N_4(i,z))$ is interpreted as having $|S_2(N_4(i,z))|/z$ blocks of length z, where

$$\frac{|S_2(N_4(i,z))|}{z} = \frac{2\cdot N_4(i,z)}{z} = \frac{2\cdot z\cdot B_4(i,z)}{z\cdot 2^i} = 2^{-i+1}B_4(i,z).$$

LEMMA 7.1. *For $s\in\{4,5\}$, $S_s(i,j)$ is an order-s Davenport-Schinzel sequence.*

PROOF. We use brackets to indicate block boundaries in forbidden patterns, for example, $[ba]ba$ is a pattern where the first ba appears in one block and the last ba appears outside that block. One can easily show by induction that $ba[ba] \not\prec S_s(i,j)$ and $[ba]ab \not\prec S_s(i,j)$ for all $s\in\{4,5\}, i>1, j\ge 1$. The base cases are trivial. When a is shuffled into the indicated block in a copy of $S_s(i,j-1)$, all b's appear in that copy and all other a's are shuffled into different copies, hence $[ba]$ cannot be preceded by ba or followed by ab. This also implies that two symbols cannot both appear in two blocks of $S_s(i,j)$, for all $i>1$. It follows that the patterns $ababab$ (and $abababa$) cannot be introduced into S_4 (and S_5) by the shuffling operation but must come from the composition (and three-fold composition) operation. Suppose $ba\prec\beta$ for some block β in $S_4(i-1,y)$. It follows that composing β with $S_2(y)$ (a $baba$-free sequence) does not introduce an $ababab$ pattern. (Substituting $\underline{bab}$ for $ba\prec\beta$ and projecting onto $\{a,b\}$ yields sequences of the form $a^*b^*\underline{bab}b^*a^*$.)

Turning to S_5, suppose $ab\prec\beta$ for some block β in S_{top}. If a and b are both middle symbols in β then, by the same argument, composing β with $S^{\text{m}}_{\text{mid}} = S_2(N_4(i,z))$ does not introduce an $ababab$ pattern much less an $abababa$ pattern. If both a,b are first then composing β with an order-4 DS sequence $S^{\text{f}}_{\text{mid}} = S_4(i,z)$ and projecting onto $\{a,b\}$ yields patterns of the form $\underline{a^*b^*a^*b^*a^*}a^*b^*$, where the underlined portion originated from β. The case when a and b are last is symmetric. The cases when a and b are of different types (first-middle, first-last, last-middle) are handled similarly. $\square$

We have shown that $\lambda_4(N_4(i,j),B_4(i,j)) \ge 2^iN_4(i,j)$ and $\lambda_5(N_5(i,j),B_5(i,j)) \ge ((2i-3)2^i+4)N_5(i,j)$. Since any blocked sequence can be turned into a 2-sparse sequence by removing duplicates at block boundaries this also implies that $\lambda_4(N_4(i,j)) \ge 2^iN_4(i,j) - B_4(i,j) > (1-1/j)2^iN_4(i,j)$. Remember that all blocks in $S_4(i,j)$ have length exactly j. There is no such guarantee for S_5, however. It is conceivable that it consists largely of long runs of identical symbols (each in a block of length 1), nearly all of which would be removed when converting it to a 2-sparse sequence. That is, statements of the form $\lambda_5(N_5(i,j)) \ge ((2i-3)2^i+4)N_5(i,j) - B_5(i,j)$ become trivial if the $B_5(i,j)$ term dominates. Lemma 7.2 shows that for j sufficiently large, this does not occur and therefore removing duplicates at block boundaries does not affect the length of $S_5(i, j)$ asymptotically.

LEMMA 7.2. *$N_5(i, j) \geq j \cdot B_5(i, j)/\xi(i)$, where $\xi(i) = 3^i 2^{\binom{i+1}{2}}$.*

PROOF. When $i = 1$ we have $N_5(1, j) = j \geq j \cdot B_5(1, j)/\xi(1) = 2j/6$. When $j = 1$, we have $N_5(i, 1) = 1 \geq B_5(i, 1)/\xi(i) = \left((2i-3)2^i + 4\right)/3^i 2^{\binom{i+1}{2}}$. Now assume inductively that the claim holds for all $(i', j') < (i, j)$ lexicographically. By the definition of N_5, we have

$$N_5(i, j) = N_5(i-1, N_4(i, z)) + B_5(i-1, N_4(i, z)) \cdot (2 + 2^{-i+1}) \cdot B_4(i, z) \cdot N_5(i, j-1).$$

Applying the inductive hypothesis to the last factor and using the definition of $z = B_5(i, j-1)$, the previous line is bounded as

$$\geq N_5(i-1, N_4(i, z)) + \frac{1}{\xi(i)} B_5(i-1, N_4(i, z)) \cdot (2 + 2^{-i+1}) \cdot B_4(i, z) \cdot (j-1) \cdot z,$$

which, by the definition of B_5, is exactly

$$= N_5(i-1, N_4(i, z)) + \frac{j-1}{\xi(i)} B_5(i, j).$$

Applying the inductive hypothesis once more, the previous line is bounded by

$$\geq \frac{1}{\xi(i-1)} N_4(i, z) \cdot B_5(i-1, N_4(i, z)) + \frac{j-1}{\xi(i)} B_5(i, j),$$

and since $N_4(i, z) = \frac{z}{2^i} B_4(i, z)$ and $2 + 2^{-i+1} \leq 3$, the previous line is

$$\begin{aligned} &\geq \frac{1}{\xi(i-1) \cdot 2^i} \cdot z \cdot B_4(i, z) \cdot B_5(i-1, N_4(i, z)) + \frac{j-1}{\xi(i)} B_5(i, j) \\ &\geq \frac{1}{\xi(i-1) \cdot 2^i \cdot 3} \cdot (2 + 2^{-i+1}) \cdot z \cdot B_4(i, z) \cdot B_5(i-1, N_4(i, z)) + \frac{j-1}{\xi(i)} B_5(i, j). \end{aligned}$$

Finally, by the definition of B_5 and ξ, the previous line is equal to

$$= \frac{1}{\xi(i)} B_5(i, j) + \frac{j-1}{\xi(i)} B_5(i, j) = \frac{j}{\xi(i)} B_5(i, j),$$

which concludes the induction. □

THEOREM 7.3. *For any n and m, $\lambda_5(n, m) = \Omega(n\alpha(n, m)2^{\alpha(n,m)})$ and $\lambda_5(n) = \Omega(n\alpha(n)2^{\alpha(n)})$.*

PROOF. Consider the sequence $S_5 = S_5(i, j)$, where $j \geq \xi(i)$, and let S_5' be obtained by removing duplicates at block boundaries. It follows that S_5' is 2-sparse and, from Lemma 7.2, that $|S_5'| \geq ((2i-3)2^i + 3)N_5(i, j)$. It is straightforward to prove that $i = \alpha(N_5(i, j), B_5(i, j)) + O(1)$ and that $i = \alpha(N_5(i, j)) + O(1)$ when $j = \xi(i)$. □

8. DISCUSSION AND OPEN PROBLEMS

Davenport-Schinzel sequences have been applied almost exclusively to problems in combinatorial and computational geometry, with only a smattering of applications in other areas (e.g., see [Alstrup et al. 1997; Pettie 2008, 2010; Di Salvo and Proietti 2007]). One explanation for this, which is undoubtedly true, is that there is a natural fit between geometric objects and their characterizations in terms of forbidden substructures.[20] An equally compelling explanation, in our opinion, is that DS sequences are simply underpublicized and that the broader algorithms community is not used to analyzing algorithms and data structures with forbidden substructure arguments. We are optimistic that with increased awareness of DS sequences and their generalizations (e.g., forbidden 0-1 matrices) the *forbidden substructure method* [Pettie 2010] will become a standard tool in every algorithms researcher's toolbox.

Our bounds on Davenport-Schinzel sequences are sharp for every order s, leaving little room for improvement.[21] However, there are many open problems on the geometric realizability of DS sequences and on various generalizations of DS sequences. The most significant realizability result is due to Wiernik and Sharir [1988], who prove that the lower envelope of n line segments (i.e., n linear functions, each defined over a different interval) has complexity $\Theta(\lambda_3(n)) = \Theta(n\alpha(n))$. It is an open question whether this result can be generalized to degree-s polynomials or polynomial segments. In particular, it may be that the lower envelope of any set of n degree-s polynomials has complexity $O(n)$, where s only influences the leading constant.

There are several challenging open problems in the realm of generalized Davenport-Schinzel sequences, the foremost being that of characterizing the set of all linear forbidden subsequences: those σ for which $\mathrm{Ex}(\sigma, n) = O(n)$ [Klazar 2002; Pettie 2011b]. Linear forbidden subsequences and minimally nonlinear ones were exhibited by Adamec et al. [1992], Klazar and Valtr [1994], and Pettie [2011a, 2011b, 2011c, 2011d]. It is also an open problem to characterize minimally nonlinear forbidden 0-1 matrices [Füredi and Hajnal 1992]. Though far from being solved, there has been significant progress on this problem in the last decade [Fulek 2009; Geneson 2009; Marcus and Tardos 2004; Pettie 2011a, 2011b, Tardos 2005].

APPENDIXES

A. PROOF OF LEMMA 3.1

Recall the four parts of Lemma 3.1.

Restatement of Lemma 3.1. *Let $\gamma_s(n) : \mathbb{N} \to \mathbb{N}$ be a nondecreasing function such that $\lambda_s(n) \le \gamma_s(n) \cdot n$.*

(*1*) *For $s \ge 1$, $\lambda_s(n, m) \le m - 1 + \lambda_s(n)$.*
(*2*) *For $s \ge 3$, $\lambda_s(n) \le \gamma_{s-2}(n) \cdot \lambda_s(n, 2n-1)$. (This generalizes the proof of Hart and Sharir [1986] for $s = 3$.)*
(*3*) *For $s \ge 2$, $\lambda_s(n) \le \gamma_{s-1}(n) \cdot \lambda_s(n, n)$.*
(*4*) *For $s \ge 3$, $\lambda_s(n) \le \gamma_{s-2}(\gamma_s(n)) \cdot \lambda_s(n, 3n-1)$.*

Proof. Removing at most $m - 1$ repeated symbols at block boundaries makes any sequence 2-sparse, which implies part (1).

For parts (2) and (3), consider the following method for greedily partitioning a 2-sparse, order-s DS sequence S with $\|S\| = n$. Write S as $S_1S_2 \cdots S_m$, where S_1 is the longest order-$(s-2)$ prefix of S, S_2 is the longest order-$(s-2)$ prefix of the remainder of the sequence, and so on. Each S_q contains the first or last occurrence of some symbol, which implies $m \le 2n - 1$, since S_1 must contain the first occurrence of at least two symbols. To see this, consider the symbol b which caused the termination of S_q, that is, S_q has order $s - 2$ but S_qb contains an alternating subsequence $\sigma_s = aba \cdots ab$ or $ba \cdots ab$ with length s; whether it starts with a depends on the parity of s. If S_q contains

[20] For example, in general, two lines do not share two points, three spheres do not share three points, degree-d polynomials do not have $d + 1$ zeros, and so on.
[21] That is, they cannot be expressed more tightly using a generic inverse-Ackermann function $\alpha(n)$. See Remark 1.1.

neither the first nor last occurrence of both a and b, S would contain an alternating subsequence σ_{s+2} of length $s+2$, a contradiction. Obtain S' from S replacing each S_q with a block containing exactly one occurrence of each symbol in $\Sigma(S_q)$. Thus,

$$\begin{aligned}
|S| = \sum_{q=1}^{m} |S_q| &\leq \sum_{q=1}^{m} \gamma_{s-2}(\|S_q\|) \cdot \|S_q\| && \{S_q \text{ has order } s-2, \text{defn. of } \gamma_{s-2}\} \\
&\leq \gamma_{s-2}(n) \cdot \sum_{q=1}^{m} \|S_q\| && \{\gamma_{s-2} \text{ is nondecreasing}\} \\
&= \gamma_{s-2}(n) \cdot |S'| \leq \gamma_{s-2}(n) \cdot \lambda_s(n,m) && \{S' \prec S \text{ has order } s\},
\end{aligned}$$

which proves part (2). Part (3) is proved in the same way except that we partition S into order-$(s-1)$ DS sequences. In this case, each S_q must contain the last occurrence of some symbol, so $m \leq n$. We turn now to part (4).

Partition S into order-$(s-2)$ sequences $S_1 S_2 \cdots S_m$ as follows. After $S_1, \ldots, S_{q-1}$ have been selected, let S_q be the longest prefix of the remaining sequence that (i) has order $s-2$ and (ii) has length at most $\gamma_s(n)$. The number of such sequences that are terminated due to (i) is at most $2n-1$, by the same argument from part (2). The number terminated due to (ii) is at most n, since $|S| \leq \gamma_s(n) \cdot n$, so $m \leq 3n-1$. Obtain an m-block sequence S' in the usual way, by replacing each S_q with a block containing its alphabet. Thus,

$$\begin{aligned}
|S| = \sum_{q=1}^{m} |S_q| &\leq \sum_{q=1}^{m} \gamma_{s-2}(\|S_q\|) \cdot \|S_q\| && \{S_q \text{ has order } s-2, \text{defn. of } \gamma_{s-2}\} \\
&\leq \gamma_{s-2}(\gamma_s(n)) \cdot \sum_{q=1}^{m} \|S_q\| && \{\gamma_s \text{ is nondecreasing}, \|S_q\| \leq |S_q| \leq \gamma_s(n)\} \\
&= \gamma_{s-2}(\gamma_s(n)) \cdot |S'| \leq \gamma_{s-2}(\gamma_s(n)) \cdot \lambda_s(n,m) && \{S' \prec S \text{ has order } s\}. \quad \square
\end{aligned}$$

Note that while part (4) is stronger than part (2), it requires an upper bound on $\gamma_s(n)$ to be applied, which is obtained by invoking part (2). In the end, it does not matter precisely what $\gamma_s(n)$ is. Once $\gamma_s(n)$ is known to be some primitive recursive function of $\alpha(n)$, it follows that $\gamma_{s-2}(\gamma_s(n)) = \gamma_{s-2}(\alpha(n)) + O(1)$.

B. PROOF OF LEMMA 5.3

Recall our definition of Ackermann's function: $a_{1,j} = 2^j$, $a_{i,1} = 2$, and $a_{i,j} = w \cdot a_{i-1,w}$, where $w = a_{i,j-1}$. Our task in this section is to prove the omnibus Lemma 5.3 in several stages.

Restatement of Lemma 5.3. Let $s \geq 1$ be the order parameter, $c \geq s-2$ be a constant, and $i \geq 1$ be an arbitrary integer. Define j to be maximum such that $m \leq a_{i,j}^c$. The following upper bounds on λ_s and Φ_s hold for all $s \geq 1$.

$$\begin{aligned}
&\lambda_1(n,m) = n + m - 1 && s = 1 \\
&\lambda_2(n,m) = 2n + m - 2 && s = 2 \\
&\lambda_3(n,m) \leq (2i+2)n + (3i-2)cj(m-1) && s = 3 \\
&\lambda_s(n,m) \leq \mu_{s,i}(n + (cj)^{s-2}(m-1)) && \textit{all } s \geq 4 \\
&\Phi_s(n,m) \leq \nu_{s,i}(n + (cj)^{s-2}(m-1)) && \textit{all } s \geq 5.
\end{aligned}$$

The values $\{\mu_{s,i}, \nu_{s,i}\}$ *are defined as follows, where* $t = \lfloor \frac{s-2}{2} \rfloor$ *and* C *is an absolute constant.*

$$\mu_{s,i} = \begin{cases} 2^{\binom{i+t+C}{t}} - 6(i+2) & \textit{even } s \geq 4 \\ 3i \cdot 2^{\binom{i+t+C}{t}} & \textit{odd } s \geq 5 \end{cases}$$

$$\nu_{s,i} = \binom{i+s-2}{s-2} - 1 \qquad \textit{all } s \geq 5.$$

Overview. The proof is by induction on (s, i, j) with respect to any fixed $c \geq s-2$. In Section B.1, we confirm that Lemma 5.3 holds when $i = 1$. In Section B.2, we discuss the role that Ackermann's function plays in selecting block partitions for Recurrences 3.3, 5.1, and 5.2. In Section B.3, we confirm Lemma 5.3's bounds on λ_s at $s = 3$. In Section B.4, we identify sufficient lower bounds on the elements of $\{\mu_{s,i}, \nu_{s,i}\}_{s\geq 2, i\geq 1}$, then, in Section B.5, prove that the particular ensemble $\{\mu_{s,i}, \nu_{s,i}\}_{s\geq 2, i\geq 1}$ proposed in Lemma 5.3 does, in fact, satisfy these lower bounds.

B.1. Base Cases

Call a block partition $\{m_q\}_{1\leq q\leq \hat{m}}$ *uniform with width* w if $m_1 = \cdots = m_{\hat{m}-1} = w$ and $m_{\hat{m}}$ may be smaller.

LEMMA B.1. *Let* $n, m,$ *and* $s \geq 2$ *be the alphabet size, block count, and order parameters. Given* $i \geq 1$, *let* j' *be minimum such that* $m \leq a_{i,j'}$. *Whether* $i = 1$ *and* $j' \geq 1$ *or* $j' = 1$ *and* $i > 1$, *we have*

$$\lambda_s(n, m) \leq 2^{s-1}n + j'^{s-2}(m-1)$$
$$\Phi_s(n, m) \leq (s-2)n + j'(m-1).$$

PROOF. The two claims are true when $s = 2$, by Lemma 3.2 and Recurrence 5.1. At $s \geq 3$, $j' = 1$, the claim is trivial, since there are only $a_{i,1} = 2$ blocks: $\lambda_s(n, 2) = 2n$ and $\Phi_s(n, 2) = 0$.

In the general case, we have $s \geq 3$ and $j' > 1$. Let S be an order-s, m-block sequence over an n-letter alphabet, where $m \leq a_{1,j'} = 2^{j'}$. Let $S = S_1S_2$ be the partition of S using a uniform block partition with width $a_{1,j'-1} = 2^{j'-1}$, so $[\![S_1]\!] = a_{1,j'-1}$ and $[\![S_2]\!] = m - a_{1,j'-1} \leq a_{1,j'-1}$. Note that $\hat{S}' = \beta_1\beta_2$ consists of two blocks, where each β_q is some permutation of the global alphabet $\hat{\Sigma}$.

Consider the first claim. Since there are no middle occurrences in $\hat{S}'$ or S, we can apply a simplified version of Recurrence 3.3.

$$\begin{aligned} &\lambda_s(n, m) \\ &\quad\leq \sum_{q=1,2} \lambda_s(\check{n}_q, [\![S_q]\!]) + \lambda_{s-1}(\hat{n}, [\![S_1]\!]) + \lambda_{s-1}(\hat{n}, [\![S_2]\!]) && \{\text{local, first, and last}\} \\ &\quad\leq 2^{s-1}(n-\hat{n}) + (j'-1)^{s-2}(m-2) + 2(2^{s-2}\hat{n}) + (j'-1)^{s-3}(m-2) && \{\text{ind. hyp.}\} \\ &\quad< 2^{s-1}n + j'^{s-2}(m-1). \end{aligned}$$

The last inequality follows from the fact that when $s \geq 3$, $(j'-1)^{s-2} + (j'-1)^{s-3} \leq j'^{s-2}$.

For the second, claim we invoke Recurrence 5.1.

$$\begin{aligned}
&\Phi_s(n, m)\\
&\le \sum_{q=1,2} \Phi_s(\check{n}_q, [\![S_q]\!]) \; + \; \Phi_s(\hat{n}, 2) + \Phi_{s-1}(\hat{n}, m) + \hat{n}\\
&\le (s-2)(n-\hat{n}) + (j'-1)^{s-2}(m-2) + (s-3)\hat{n} + (j'-1)^{s-3}(m-1) + \hat{n}\\
&\le (s-2)n + j'^{s-2}(m-1).
\end{aligned}$$

This concludes the induction. □

If we introduce the c parameter and define j to be minimum such that $m \le a^c_{i,j}$, Lemma B.1 implies that $\lambda_s(n, m) \le 2^{s-1}n + (cj)^{s-2}(m-1)$, since $j' \le cj$. Note that by definition of Ackermann's function, $a^c_{1,j} = (2^j)^c = a_{1,cj}$ and $a^c_{i,1} = 2^c = a_{1,c}$.

Lemma B.1 implies the claims of Lemma 5.3 when $i = 1$, for a sufficiently large constant C. For any s, $\nu_{s,i} = \binom{i+s-2}{s-2} - 1 = s-2$. When $s = 3$, $\mu_{3,i} = 2i + 2 = 4 = 2^{s-1}$, and $3i - 2 = 1$. When $s \ge 4$ is even, $\mu_{s,i} = 2^{\binom{i+t+C}{t}} - 6(i+2) \ge 2^{2t+1} = 2^{s-1}$. When $s \ge 5$ is odd, $\mu_{s,i} = 3 \cdot 2^{\binom{i+t+C}{t}} \ge 2^{2t+2} = 2^{s-1}$. These bounds also imply that Lemma 5.3 holds at $j = 1$ and $i > 1$, since $a^c_{i,1} = a^c_{1,1}$ and both $\mu_{s,i}$ and $\nu_{s,i}$ are increasing in i.

B.2. Block Partitions and Inductive Hypotheses

When analyzing order-s DS sequences, we express the block count m and partition size $\hat{m}$ in terms of constant powers of Ackermann's function $\{a^c_{i,j}\}$, where the constant $c \ge s-2$ is fixed. Recall that once i is selected, j is minimal such that $m \le a^c_{i,j}$. The base cases $i = 1$ and $j = 1$ have been handled so we can assume both are at least 2. Let $w = a_{i,j-1}$.

We always choose a uniform block partition $\{m_q\}_{1\le q\le \hat{m}}$ with width w^c, that is, $m_q = w^c$ for all $q < \hat{m} = \lceil m/w^c \rceil$, and the leftover $m_{\hat{m}}$ may be smaller. When invoking the inductive hypothesis on the $\hat{m}$-block sequence $\hat{S}'$, we use parameter $i-1$. In all other invocations of the inductive hypothesis, we use parameter i. When applied to any m_q-block sequences, the "j" parameter is decremented, since $m_q \le w^c = a^c_{i,j-1}$. When applied to an $\hat{m}$-block sequence, the "j" parameter is w, since

$$\hat{m} = \left\lceil \frac{m}{w^c} \right\rceil \le \left(\frac{a_{i,j}}{w}\right)^c = a^c_{i-1,w}.$$

Furthermore, in such an invocation, the dependence on $\hat{m}$ will always be at most linear in m, since $(cw)^{s-2}(\hat{m}-1) \le (cw)^{s-2}(\lceil \frac{m}{w^c} \rceil - 1) \le c^{s-2}(m-1)$. This is the reason we require the lower bound $c \ge s-2$.

If one is more familiar with the slowly growing row-inverses of Ackermann's function, it may be helpful to remember that $cj = \log m - O(1)$ when $i = 1$ and that $j = \log^{[i-1]}(m) - O(1)$ when $i > 1$, the effect of the c parameter being negligible since $a_{i,j}$ and $a^c_{i,j}$ are essentially identical relative to any sufficiently slowly growing function.[22] Thus, the bounds of Lemma 5.3 could be rephrased as $\lambda_s(n, m) \le \mu_{s,i}(n + O(m(\log^{[i-1]}(m))^{s-2}))$. Since $\mu_{s,i}$ is increasing in i, the best bounds are obtained by choosing i to be minimal such that $\log^{[i-1]}(m) = n/m + O(1)$.

[22] Recall that $\log^{[i-1]}(m)$ is short for $\log^{\star\cdots\star}(m)$ with $i-1$ $\star$s.

B.3. Bounds at Order 3

LEMMA B.2. *Let n and m be the alphabet size and block count of an order-3 DS sequence S. For any $i, c \geq 1$, define j to be minimum such that $m \leq a_{i,j}^c$. Then λ_3 is bounded by*

$$\lambda_3(n, m) \leq (2i+2)n + (3i-2)cj(m-1).$$

PROOF. The base cases $i = 1$ and $j = 1$ have been handled already. Let $i, j > 1$ and $w = a_{i,j-1}$. We invoke Recurrence 3.3 with the uniform block partition $\{m_q\}_{1 \leq q \leq \hat{m}}$, where $\hat{m} = \lceil m/w^c \rceil$ (see Section B.2).

$$\begin{aligned}
\lambda_3(n, m) &\leq \textstyle\sum_{q=1}^{\hat{m}} \lambda_3(\check{n}_q, m_q) \;+\; 2 \cdot \lambda_2(\hat{n}, m) \;+\; \lambda_1(\lambda_3(\hat{n}, \hat{m}) - 2\hat{n}, m) \\
&\leq (2i+2)(n-\hat{n}) \;+\; (3i-2)c(j-1)(m-\hat{m}) && \{\text{local symbols}\} \\
&\quad +\; 4\hat{n} \;+\; 2(m-1) && \{\text{global first and last}\} \\
&\quad +\; (2i-2)\hat{n} \;+\; (3(i-1)-2)cw(\hat{m}-1) \;+\; (m-1) && \{\text{global middle}\} \\
&\leq (2i+2)n \;+\; (3i-2)cj(m-1) \\
&\quad +\; [-(2i+2)+4+(2i-2)]\hat{n} \;+\; [-c(3i-2)+(3i-5)+3](m-1) \\
&\leq (2i+2)n \;+\; (3i-2)cj(m-1).
\end{aligned}$$

The last inequality holds since $c \geq s-2 = 1$. □

At $s = 2$ and $s = 3$, the terms involving n and m have different leading constants, namely, 2 and 1 when $s = 2$ and $2i+2$ and $3i-2$ when $s = 3$. To provide some uniformity in the following analyses, we will use the inequalities $\lambda_2(n, m) \leq \mu_{2,i}(n+m-1)$ and $\lambda_3(n, m) \leq \mu_{3,i}(n + (cj)(m-1))$ when invoking the inductive hypothesis at $i \geq 2$ and $s \in \{2, 3\}$. By definition, $\mu_{2,i} = 2$ and $\mu_{3,i} = 3i$. Note that when $i \geq 2$, $\mu_{3,i} = 3i \geq \max\{2i+2, 3i-2\}$.

B.4. Lower Bounds on $\mu_{s,i}$ and $\nu_{s,i}$

Call an ensemble of values $\{\mu_{s',i'}, \nu_{s',i'}\}_{(s',i') \leq (s,i)}$ *happy* if $\lambda_{s'}(n, m) \leq \mu_{s',i'}(n+(cj)^{s'-2}(m-1))$ and $\Phi_{s'}(n, m) \leq \nu_{s',i'}(n + (cj)^{s'-2}(m-1))$, where c and j are defined as usual. (In the subscript, '$\leq$' represents lexicographic ordering on tuples.) In Lemma B.3, we determine lower bounds on $\mu_{s,i}$ and $\nu_{s,i}$ in a happy ensemble. In Section B.5, we prove that the *specific* ensemble proposed in Lemma 5.3 is, in fact, happy.

LEMMA B.3. *Let $s \geq 4$ and $i \geq 2$. Define n, m, c, and j as usual. If $\{\mu_{s',i'}, \nu_{s',i'}\}_{(s',i') \leq (s,i-1)}$ is happy, then $\{\mu_{s',i'}, \nu_{s',i'}\}_{(s',i') \leq (s,i)}$ is as well, so long as*

$$\begin{aligned}
\mu_{s,i} &\geq 2\mu_{s-1,i} \;+\; \mu_{s-2,i}\mu_{s,i-1} && \textit{even } s \\
\mu_{s,i} &\geq 2\mu_{s-1,i} \;+\; 2\mu_{s-2,i}\nu_{s,i-1} \;+\; \mu_{s-3,i}\mu_{s,i-1} && \textit{odd } s \\
\nu_{s,i} &\geq \nu_{s,i-1} \;+\; \nu_{s-1,i} \;+\; 1 && \textit{all } s.
\end{aligned}$$

PROOF. When $s \geq 4$ is even, Recurrence 3.3 implies that

$$\begin{aligned}
\lambda_s(n, m) &\leq \sum_{q=1}^{\hat{m}} \lambda_s(\check{n}_q, m_q) \;+\; 2 \cdot \lambda_{s-1}(\hat{n}, m) \;+\; \lambda_{s-2}(\lambda_s(\hat{n}, \hat{m}), m) \\
&\leq \mu_{s,i}((n-\hat{n}) \;+\; (c(j-1))^{s-2}(m-\hat{m})) && \{\text{happiness of the ensemble}\}
\end{aligned}$$

$$
\begin{aligned}
&\quad + \; 2\mu_{s-1,i}(\hat{n} \; + \; (cj)^{s-3}(m-1)) \\
&\quad + \; \mu_{s-2,i}(\mu_{s,i-1}(\hat{n} \; + \; (cw)^{s-2}(\hat{m}-1)) \; + \; (cj)^{s-4}(m-1)) \\
&\le \mu_{s,i}(n \; + \; (cj)^{s-2}(m-1)) \\
&\quad + \; [-\mu_{s,i} + 2\mu_{s-1,i} + \mu_{s-2,i}\mu_{s,i-1}] \cdot \hat{n} && (5) \\
&\quad + \; [-\mu_{s,i}c^{s-2}j^{s-3} + 2\mu_{s-1,i}(cj)^{s-3} \\
&\quad + \mu_{s-2,i}\mu_{s,i-1}c^{s-2} + \mu_{s-2,i}(cj)^{s-4}] \cdot (m-1) && (6) \\
&\le \mu_{s,i}(n \; + \; (cj)^{s-2}(m-1)). && (7)
\end{aligned}
$$

Inequality (7) will be satisfied whenever (5) and (6) are nonpositive, that is, when

$$
\mu_{s,i} \;\ge\; 2\mu_{s-1,i} \;+\; \mu_{s-2,i}\mu_{s,i-1} \tag{8}
$$

$$
\mu_{s,i} \;\ge\; \frac{2\mu_{s-1,i}}{s-2} \;+\; \frac{\mu_{s-2,i}\mu_{s,i-1}}{2^{s-3}} \;+\; \frac{\mu_{s-2,i}}{2(s-2)^2}. \tag{9}
$$

Inequality (9) was obtained by dividing (6) through by $c^{s-2}j^{s-3}$ and noting that $c \ge s-2 \ge 2$ and $j \ge 2$. Note that Inequality (9) is weaker than Inequality (8), since $\mu_{s,i} > \mu_{s-1,i} > \mu_{s-2,i}$, so it suffices to consider only the former.

When $s \ge 5$ is odd, Recurrence 5.2 implies that

$$
\begin{aligned}
\lambda_s(n,m) &\le \sum_{q=1}^{\hat{m}} \lambda_s(\check{n}_q, m_q) + 2 \cdot \lambda_{s-1}(\hat{n}, m) + \lambda_{s-2}(2 \cdot \Phi_s(\hat{n}, \hat{m}), m) + \lambda_{s-3}(\lambda_s(\hat{n}, \hat{m}), m) \\
&\le \mu_{s,i}((n-\hat{n}) \; + \; (c(j-1))^{s-2}(m-\hat{m})) \qquad \{\text{happiness of the ensemble}\} \\
&\quad + \; 2\mu_{s-1,i}(\hat{n} \; + \; (cj)^{s-3}(m-1)) \\
&\quad + \; \mu_{s-2,i}(2 \cdot \nu_{s,i-1}(\hat{n} \; + \; (cw)^{s-2}(\hat{m}-1)) \; + \; (cj)^{s-4}(m-1)) \\
&\quad + \; \mu_{s-3,i}(\mu_{s,i-1}(\hat{n} \; + \; (cw)^{s-2}(\hat{m}-1)) \; + \; (cj)^{s-5}(m-1)) \\
&\le \mu_{s,i}(n \; + \; (cj)^{s-2}(m-1)) \\
&\quad + \; [-\mu_{s,i} \; + \; 2\mu_{s-1,i} \; + \; 2\mu_{s-2,i}\nu_{s,i-1} \; + \; \mu_{s-3,i}\mu_{s,i-1}] \cdot \hat{n} && (10) \\
&\quad + \; \big[-\mu_{s,i}c^{s-2}j^{s-3} \; + \; 2\mu_{s-1,i}(cj)^{s-3} \; + \; 2\mu_{s-2,i}\nu_{s,i-1}c^{s-2} \\
&\quad + \; \mu_{s-2,i}(cj)^{s-4} \; + \; \mu_{s-3,i}\mu_{s,i-1}c^{s-2} \; + \; \mu_{s-3,i}(cj)^{s-5}\big] \cdot (m-1) && (11) \\
&\le \mu_{s,i}(n \; + \; (cj)^{s-2}(m-1)). && (12)
\end{aligned}
$$

Inequality (12) will be satisfied whenever (10) and (11) are nonpositive, that is, when

$$
\mu_{s,i} \;\ge\; 2\mu_{s-1,i} \;+\; 2\mu_{s-2,i}\nu_{s,i-1} \;+\; \mu_{s-3,i}\mu_{s,i-1} \tag{13}
$$

$$
\mu_{s,i} \;\ge\; \frac{2\mu_{s-1,i}}{s-2} \;+\; \frac{\mu_{s-2,i}\nu_{s,i-1}}{2^{s-2}} \;+\; \frac{\mu_{s-2,i}}{2(s-2)^2} \;+\; \frac{\mu_{s-3,i}\mu_{s,i-1}}{2^{s-3}} \;+\; \frac{\mu_{s-3,i}}{4(s-2)^3}. \tag{14}
$$

The denominators of Inequality (14) follow by dividing (11) through by $c^{s-2}j^{s-3}$ and noting that $c \ge s-2 \ge 3$ and $j \ge 2$. Inequality (14) is weaker than Inequality (13), since $\mu_{s-1,i} > \mu_{s-2,i} > \mu_{s-3,i}$, so it suffices to consider only Inequality (13).

Using similar calculations, one derives from Recurrence 5.1 the claimed lower bound on $\nu_{s,i}$.

$$
\nu_{s,i} \;\ge\; \nu_{s,i-1} \;+\; \nu_{s-1,i} \;+\; 1. \quad \square \tag{15}
$$

B.5. The Happiness of the Ensemble

From this point, on we argue the happiness of the *specific* ensemble $\{\mu_{s,i}, \nu_{s,i}\}$ stated in Lemma 5.3. We can say an individual value $\mu_{s,i}$ or $\nu_{s,i}$ is happy if it satisfies the appropriate lower bound inequality, either (9), (14), or (15).

LEMMA B.4. *The ensemble $\{\mu_{s,i}, \nu_{s,i}\}$ defined in Lemma 5.3 is happy.*

PROOF. The $\{\nu_{s,i}\}$ are happy since

$$\nu_{s,i-1} + \nu_{s-1,i} + 1 = \left(\binom{i-1+s-2}{s-2} - 1\right) + \left(\binom{i+s-3}{s-3} - 1\right) + 1 = \binom{i+s-2}{s-2} - 1 = \nu_{s,i}.$$

When $s = 4$ and $t = \lfloor\frac{s-2}{2}\rfloor = 1$, the expression for $\mu_{4,i}$ simplifies to $2^{i+1+C} - 6(i+2)$. The happiness of $\mu_{4,i}$ follows easily, as seen here.

$$\begin{aligned}
&2\mu_{3,i} + \mu_{2,i}\mu_{4,i-1}\\
&\quad = 2(3i) + 2\cdot(2^{i+C} - 6(i+1))\\
&\quad = 2^{i+1+C} + 6i - 12(i+1)\\
&\quad = 2^{i+1+C} - 6(i+2) = \mu_{4,i}.
\end{aligned}$$

When $s = 5$ and $t = \lfloor\frac{s-2}{2}\rfloor = 1$, the expression for $\mu_{5,i}$ simplifies to $3i\cdot 2^{i+1+C}$, which lets us quickly certify the happiness of $\mu_{5,i}$.

$$\begin{aligned}
&2\mu_{4,i} + 2\mu_{3,i}\nu_{5,i-1} + \mu_{2,i}\mu_{5,i-1}\\
&\quad < 2(2^{i+1+C}) + 2\cdot 3i\cdot\binom{i-1+3}{3} + 2\cdot 3(i-1)\cdot 2^{i+C}\\
&\quad < 3i\cdot 2^{i+1+C} = \mu_{5,i}.
\end{aligned}$$

The last inequality follows from the fact that $6i\cdot\binom{i-1+3}{3} < 2^{i+1+C}$ when C is sufficiently large. We now turn to the happiness of $\mu_{s,i}$ for even $s \geq 6$. Note that when we invoke the definition of $\mu_{s-1,i}$ and $\mu_{s-2,i}$, their "t" parameter is $t - 1 = \lfloor\frac{(s-1)-2}{2}\rfloor = \lfloor\frac{(s-2)-2}{2}\rfloor$.

$$\begin{aligned}
&2\mu_{s-1,i} + \mu_{s-2,i}\mu_{s,i-1}\\
&\quad < 2\cdot\left[3i\cdot 2^{\binom{i+t-1+C}{t-1}}\right] + \left[2^{\binom{i+t-1+C}{t-1}} - 6(i+2)\right]\cdot\left[2^{\binom{i-1+t+C}{t}} - 6(i+1)\right]\\
&\quad < 2^{\binom{i+t+C}{t}} - 6(i+2) = \mu_{s,i}.
\end{aligned}$$

In other words, $\mu_{s,i}$ satisfies Inequality (8) when $s \geq 6$ is even. It also satisfies Inequality (13) at odd $s \geq 7$, which can be seen as follows. Note that the t parameter for $s-1$ is t, whereas it is $t-1$ for $s-2$ and $s-3$.

$$\begin{aligned}
&2\mu_{s-1,i} + 2\mu_{s-2,i}\nu_{s,i-1} + \mu_{s-3,i}\mu_{s,i-1}\\
&\quad < \left[2\cdot 2^{\binom{i+t+C}{t}}\right] + \left[2\cdot 3i\cdot 2^{\binom{i+t-1+C}{t-1}}\cdot\binom{i-1+s-2}{s-2}\right] + \left[2^{\binom{i+t-1+C}{t-1}}\cdot 3(i-1)\cdot 2^{\binom{i-1+t+C}{t}}\right]\\
&\quad = \left[2\cdot 2^{\binom{i+t+C}{t}}\right] + \left[2\cdot 3i\cdot 2^{\binom{i+t-1+C}{t-1}}\cdot\binom{i-1+s-2}{s-2}\right] + \left[3(i-1)\cdot 2^{\binom{i+t+C}{t}}\right]\\
&\quad < 3i\cdot 2^{\binom{i+t+C}{t}} = \mu_{s,i}.
\end{aligned}$$

The last inequality follows from that fact that the middle term, $6i\cdot 2^{\binom{i+t-1+C}{t-1}}\cdot\binom{i-1+s-2}{s-2}$, is less than $2^{\binom{i+t+C}{t}}$ when C is sufficiently large.

We have shown that $\{\mu_{s,i}\}$ and $\{\nu_{s,i}\}$ are happy over the full range of parameters. This concludes the proof of Lemma 5.3. $\square$

ACKNOWLEDGMENT

We thank the two anonymous reviewers for their careful reading of the article, and for offering us thoughtful suggestions for improvements.

REFERENCES

M. A. Abam, M. de Berg, and J. Gudmundsson. 2010. A simple and efficient kinetic spanner. *Comput. Geom.* 43, 3 (2010), 251–256.

R. Adamec, M. Klazar, and P. Valtr. 1992. Generalized Davenport-Schinzel sequences with linear upper bound. *Discrete Math.* 108, 1–3 (1992), 219–229.

P. Agarwal and M. Sharir. 2002. The number of congruent simplices in a point set. *Discrete Computat. Geom.* 28, 2 (2002), 123–150.

P. Agarwal, M. Sharir, and P. Shor. 1989. Sharp upper and lower bounds on the length of general Davenport-Schinzel sequences. *J. Combin. Theory Ser. A* 52 (1989), 228–274.

P. K. Agarwal, H. Kaplan, and M. Sharir. 2008. Kinetic and dynamic data structures for closest pair and all nearest neighbors. *ACM Trans. Algorithms* 5, 1 (2008).

G. Albers, L. J. Guibas, J. S. B. Mitchell, and T. Roos. 1998. Voronoi diagrams of moving points. *Int. J. Comput. Geom. Appl.* 8, 3 (1998), 365–380.

A. Alon and B. Schieber. 1987. *Optimal Preprocessing for Answering On-line Product Queries*. Technical Report TR-71/87. Institute of Computer Science, Tel Aviv University.

N. Alon, H. Kaplan, G. Nivasch, M. Sharir, and S. Smorodinsky. 2008. Weak ϵ-nets and interval chains. *J. ACM* 55, 6 (2008).

S. Alstrup, P. W. Lauridsen, P. Sommerlund, and M. Thorup. 1997. Finding cores of limited length. In *Proceedings of the 5th International Workshop on Algorithms and Data Structures (WADS)*. 45–54.

H. Aonuma, H. Imai, K. Imai, and T. Tokuyama. 1990. Maximin location of convex objects in a polygon and related dynamic Voronoi diagrams. In *Proceedings of the 6th Annual Symposium on Computational Geometry (SoCG)*. 225–234.

B. Aronov, K. Buchin, M. Buchin, B. M. P. Jansen, T. de Jong, M. J. van Kreveld, M. Löffler, J. Luo, R. I. Silveira, and B. Speckmann. 2011. Connect the dot: Computing feed-links for network extension. *J. Spatial Inf. Sci.* 3, 1 (2011), 3–31.

B. Aronov and M. de Berg. 2012. Unions of fat convex polytopes have short skeletons. *Discrete Computat. Geom.* 48, 1 (2012), 53–64.

B. Aronov and D. Drusvyatskiy. 2011. Complexity of a single face in an arrangement of s-intersecting curves. *CoRR* abs/1108.4336 (2011).

M. J. Atallah. 1985. Some dynamic computational geometry problems. *Comput. Math. Appl.* 11, 12 (1985), 1171–1181.

J. Augustine, B. Putnam, and S. Roy. 2010. Largest empty circle centered on a query line. *J. Discrete Algorithms* 8, 2 (2010), 143–153.

F. Aurenhammer and B. Jüttler. 2012. On computing the convex hull of (piecewise) curved objects. *Math. Comput. Sci.* 6, 3 (2012), 261–266.

S. W. Bae and Y. Okamoto. 2012. Querying two boundary points for shortest paths in a polygonal domain. *Comput. Geom.* 45, 7 (2012), 284–293.

A. Baltsan and M. Sharir. 1988. On the shortest paths between two convex polyhedra. *J. ACM* 35, 2 (1988), 267–287.

M. W. Bern, D. Eppstein, P. E. Plassman, and F. F. Yao. 1991. Horizon theorems for lines and polygons. In *Discrete and Computational Geometry: Papers from the DIMACS Special Year*, Jacob E. Goodman, Richard Pollack, and William Steiger (Eds.). Number 6 in DIMACS Ser. Discrete Math. and Theoretical Computer Science. Amer. Math. Soc., 45–66.

R. P. Berretty, K. Goldberg, M. H. Overmars, and A. F. van der Stappen. 2001. Trap design for vibratory bowl feeders. *Int. J. Robot. Res.* 20, 11 (2001), 891–908.

R. E. Burkard and H. Dollani. 2001. Robust location problems with pos/neg weights on a tree. *Networks* 38, 2 (2001), 102–113.

B. Chazelle. 2000. A minimum spanning tree algorithm with inverse-Ackermann type complexity. *J. ACM* 47, 6 (2000), 1028–1047.

B. Chazelle and B. Rosenberg. 1991. The complexity of computing partial sums off-line. *Int. J. Comput. Geom. Appl.* 1, 1 (1991), 33–45.

J. Cibulka and J. Kynčl. 2012. Tight bounds on the maximum size of a set of permutations with bounded VC-dimension. *J. Combin. Theory Ser. A* 119, 7 (2012), 1461–1478.

R. Cole and M. Sharir. 1989. Visibility problems for polyhedral terrains. *J. Symbo. Computat.* 7, 1 (1989), 11–30.

H. Davenport. 1970/1971. A combinatorial problem connected with differential equations. II. *Acta Arith.* 17 (1970/1971), 363–372.

H. Davenport and A. Schinzel. 1965a. A combinatorial problem connected with differential equations. *Am. J. Math.* 87 (1965), 684–694.

H. Davenport and A. Schinzel. 1965b. A note on sequences and subsequences. *Elemente Der Mathematik* 20 (1965), 63–64.

M. de Berg. 2010. Better bounds on the union complexity of locally fat objects. In *Proceedings of the 26th Annual Symposium on Computational Geometry (SoCG)*. 39–47.

A. Di Salvo and G. Proietti. 2007. Swapping a failing edge of a shortest paths tree by minimizing the average stretch factor. *Theor. Comput. Sci.* 383, 1 (2007), 23–33.

A. Dumitrescu, M. Sharir, and C. D. Tóth. 2009. Extremal problems on triangle areas in two and three dimensions. *J. Combin. Theory Ser. A* 116, 7 (2009), 1177–1198.

H. Edelsbrunner, J. Pach, J. T. Schwartz, and M. Sharir. 1987. On the lower envelope of bivariate functions and its applications. In *Proceedings of the 28th IEEE Symposium on Foundations of Computer Science (FOCS)*. 27–37.

A. Efrat. 2005. The complexity of the union of (α, β)-covered objects. *SIAM J. Comput.* 34, 4 (2005), 775–787.

E. Ezra and W. Mulzer. 2012. Convex hull of points lying on lines in $o(n \log n)$ time after preprocessing. In *Proceedings of the 28th Annual Symposium on Computational Geometry (SoCG)*. 11–20.

J. Fox, J. Pach, and A. Suk. 2013. The number of edges in k-quasi-planar graphs. *SIAM J. Discrete Math.* 27, 1 (2013), 550–561.

R. Fulek. 2009. Linear bound on extremal functions of some forbidden patterns in 0-1 matrices. *Discrete Math.* 309, 6 (2009), 1736–1739.

Z. Füredi and P. Hajnal. 1992. Davenport-Schinzel theory of matrices. *Discrete Math.* 103, 3 (1992), 233–251.

H. N. Gabow. 1985. A scaling algorithm for weighted matching on general graphs. In *Proceedings of the 26th IEEE Symposium on Foundations of Computer Science (FOCS)*. 90–100.

J. T. Geneson. 2009. Extremal functions of forbidden double permutation matrices. *J. Combin. Theory Ser. A* 116, 7 (2009), 1235–1244.

L. J. Guibas, J. S. B. Mitchell, and T. Roos. 1992. Voronoi diagrams of moving points in the plane. In *Proceedings of the 17th International Workshop on Graph-Theoretic Concepts in Computer Science (WG'91)*. Lecture Notes in Computer Science, Vol. 570. Springer, Berlin, 113–125.

S. Har-Peled. 2000. Taking a walk in a planar arrangement. *SIAM J. Comput.* 30, 4 (2000), 1341–1367.

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

D. P. Huttenlocher and K. Kedem. 1990. Computing the minimum Hausdorff distance for point sets under translation. In *Proceedings of the 6th Annual Symposium on Computational Geometry (SoCG)*. 340–349.

D. P. Huttenlocher, K. Kedem, and J. M. Kleinberg. 1992. Voronoi diagrams of rigidly moving sets of points. *Info. Proc. Lett.* 43, 4 (1992), 217–223.

D. P. Huttenlocher, K. Kedem, and M. Sharir. 1993. The upper envelope of Voronoi surfaces and its applications. *Discrete Computat. Geom.* 9 (1993), 267–291.

O. Ilushin, G. Elber, D. Halperin, R. Wein, and M.-S. Kim. 2005. Precise global collision detection in multi-axis NC-machining. *Comput.-Aid. Des.* 37, 9 (2005), 909–920.

J. W. Jaromczyk and M. Kowaluk. 1988. Skewed projections with an application to line stabbing in $\mathbb{R}^3$. In *Proceedings of the 4th Annual Symposium on Computational Geometry (SoCG)*. 362–370.

H. Kaplan, N. Rubin, and M. Sharir. 2011. A kinetic triangulation scheme for moving points in the plane. *Computat. Geom.* 44, 4 (2011), 191–205.

M. M. Klawe. 1992. Superlinear bounds for matrix searching problems. *J. Algor.* 13, 1 (1992), 55–78.

M. M. Klawe and D. J. Kleitman. 1990. An almost linear time algorithm for generalized matrix searching. *SIAM J. Discr. Math.* 3, 1 (1990), 81–97.

M. Klazar. 1992. A general upper bound in extremal theory of sequences. *Comment. Math. Univ. Carolin.* 33, 4 (1992), 737–746.

M. Klazar. 1999. On the maximum lengths of Davenport-Schinzel sequences. In *Proceedings of the Contemporary Trends in Discrete Mathematics*. American Mathematical Society. 169–178.

M. Klazar. 2002. Generalized Davenport-Schinzel sequences: Results, problems, and applications. *Integers* 2 (2002), A11.

M. Klazar and P. Valtr. 1994. Generalized Davenport-Schinzel sequences. *Combinatorica* 14, 4 (1994), 463–476.

V. Koltun. 2004. Almost tight upper bounds for vertical decompositions in four dimensions. *J. ACM* 51, 5 (2004), 699–730.

V. Koltun and M. Sharir. 2003. 3-dimensional Euclidean Voronoi diagrams of lines with a fixed number of orientations. *SIAM J. Comput.* 32, 3 (2003), 616–642.

P. Komjáth. 1988. A simplified construction of nonlinear Davenport-Schinzel sequences. *J. Combin. Theory Ser. A* 49, 2 (1988), 262–267.

T. Lengauer and R. E. Tarjan. 1979. A fast algorithm for finding dominators in a flowgraph. *ACM Trans. Program. Lang. Syst.* 1, 1 (1979), 121–141.

D. Leven and M. Sharir. 1987. On the number of critical free contacts of a convex polygonal object moving in two-dimensional polygonal space. *Discrete Computat. Geom.* 2, 1 (1987), 255–270.

M. C. López-de-los Mozos, J. Puerto, and A. M. Rodríguez-Chía. 2013. Robust mean absolute deviation problems on networks with linear vertex weights. *Networks* 61, 1 (2013), 76–85.

A. Marcus and G. Tardos. 2004. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Combin. Theory Ser. A* 107, 1 (2004), 153–160.

J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl. 1994. Fat triangles determine linearly many holes. *SIAM J. Comput.* 23, 1 (1994), 154–169.

E. Moet, C. Knauer, and M. van Kreveld. 2008. Visibility maps of segments and triangles in 3D. *Computat. Geom.* 39, 3 (2008), 163–177.

G. Nivasch. 2010. Improved bounds and new techniques for Davenport-Schinzel sequences and their generalizations. *J. ACM* 57, 3, Article 17.

S. Pettie. 2006. An inverse-Ackermann type lower bound for online minimum spanning tree verification. *Combinatorica* 26, 2 (2006), 207–230.

S. Pettie. 2008. Splay trees, Davenport-Schinzel sequences, and the deque conjecture. In *Proceedings of the 19th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1115–1124.

S. Pettie. 2010. Applications of forbidden 0-1 matrices to search tree- and path compression-based data structures. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1457–1467.

S. Pettie. 2011a. Degrees of nonlinearity in forbidden 0-1 matrix problems. *Discrete Math.* 311 (2011), 2396–2410.

S. Pettie. 2011b. Generalized Davenport-Schinzel sequences and their 0-1 matrix counterparts. *J. Combin. Theory Ser. A* 118, 6 (2011), 1863–1895.

S. Pettie. 2011c. On the structure and composition of forbidden sequences, with geometric applications. In *Proceedings of the 27th Annual Symposium on Computational Geometry (SoCG)*. 370–379.

S. Pettie. 2011d. Origins of nonlinearity in Davenport-Schinzel sequences. *SIAM J. Discrete Math.* 25, 1 (2011), 211–233.

S. Pettie. 2015. Sharp bounds on formation-free sequences. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 592–604. arXiv:1401.5709.

M. Sharir. 1987. Almost linear upper bounds on the length of general Davenport-Schinzel sequences. *Combinatorica* 7, 1 (1987), 131–143.

M. Sharir. 1988. Improved lower bounds on the length of Davenport-Schinzel sequences. *Combinatorica* 8, 1 (1988), 117–124.

M. Sharir and P. Agarwal. 1995. *Davenport-Schinzel Sequences and their Geometric Applications*. Cambridge University Press.

M. Sharir, R. Cole, K. Kedem, D. Leven, R. Pollack, and S. Sifrony. 1986. Geometric applications of Davenport-Schinzel sequences. In *Proceedings of the 27th IEEE Symposium on Foundations of Computer Science (FOCS)*. 77–86.

M. Sharir and S. Toledo. 1994. Extremal polygon containment problems. *Comput. Geom.* 4 (1994), 99–118.

A. Suk. 2012. k-Quasi-planar graphs. In *Proceedings of the 19th International Symposium on Graph Drawing (GD'11)*. Lecture Notes in Computer Science, Vol. 7034. Springer, Berlin, 266–277.

E. Szemerédi. 1973/74. On a problem of Davenport and Schinzel. *Acta Arith.* 25 (1973/74), 213–224.

B. Tagansky. 1996. A new technique for analyzing substructures in arrangements of piecewise linear surfaces. *Discrete Computat. Geom.* 16, 4 (1996), 455–479.

G. Tardos. 2005. On 0-1 matrices and small excluded submatrices. *J. Combin. Theory Ser. A* 111, 2 (2005), 266–288.

R. E. Tarjan. 1975. Efficiency of a good but not linear set merging algorithm. *J. ACM* 22, 2 (1975), 215–225.

R. E. Tarjan. 1979. Applications of path compression on balanced trees. *J. ACM* 26, 4 (1979), 690–715.

P. Valtr. 1997. Graph drawings with no k pairwise crossing edges. In *Proceedings of the 5th International Symposium on Graph Drawing (GD)*. Lecture Notes in Computer Science, Vol. 1353. Springer, Berlin, 205–218.

M. A. Wahid, M. Kaykobad, and M. Hasan. 2010. Kinetisation of view of 3D point set. In *Proceedings of the 13th International Conference on Computer and Information Technology (ICCIT)*. 325–330.

A. Wiernik and M. Sharir. 1988. Planar realizations of nonlinear Davenport-Schinzel sequences by segments. *Discrete Comput. Geom.* 3, 1 (1988), 15–47.

A. C. Yao. 1982. Space-time tradeoff for answering range queries. In *Proceedings of the 14th ACM Symposium on Theory of Computing (STOC)*. 128–136.

Received November 2012; revised August 2014; accepted August 2015