On the Structure and Composition of Forbidden Sequences, with Geometric Applications*

Seth Pettie
Department of EECS
University of Michigan
Ann Arbor, MI
pettie@umich.edu

ABSTRACT

Forbidden substructure theorems have proved to be among of the most versatile tools in bounding the complexity of geometric objects and the running time of geometric algorithms. To apply them one typically *transcribes* an algorithm execution or geometric object as a sequence over some alphabet or a 0-1 matrix, proves that this object avoids some subsequence or submatrix σ, then uses an off the shelf bound on the maximum size of such a σ-free object. As a historical trend, expanding our library of forbidden substructure theorems has led to better bounds and simpler analyses of the complexity of geometric objects.

We establish new and tight bounds on the maximum length of *generalized* Davenport-Schinzel sequences, which are those whose subsequences are not isomorphic to some fixed sequence σ. (The standard Davenport-Schinzel sequences restrict σ to be of the form $abab\cdots$.)

1. We prove that N-shaped forbidden subsequences (of the form $abc\cdots xyzyx\cdots cbabc\cdots xyz$) have a linear extremal function. Our proof dramatically improves an earlier one of Klazar and Valtr in the leading constants and overall simplicity. This result tightens the (astronomical) leading constants in Valtr's $O(n\log n)$ bound on geometric graphs without $k = O(1)$ mutually crossing edges.

2. We prove tight $\Theta(n\alpha(n))$ bounds on sequences avoiding both $ababab$ and all M-shaped sequences of the form $ab\cdots yzzy\cdots baab\cdots yzzy\cdots ba$. A consequence of this result is that the complexity of the union of n δ-fat triangles is $O(n\log^* n\alpha(n))$, which improves, slightly, a recent bound of Ezra, Aronov, and Sharir. Here α is the inverse-Ackermann function.

3. We give a complete characterization of 3-letter linear and nonlinear forbidden subsequences without repetitions. Specifically, a repetition-free forbidden subsequence is nonlinear ($\Omega(n\alpha(n))$) if and only if contains $ababa$, $abcacbc$, or its reversal; all others are linear.

Many of our results are obtained by reinterpreting (forbidden) sequences as (forbidden) 0-1 matrices, which can alternatively be thought of as point sets with integer coordinates. By considering a dual sequence/matrix representation we can then apply techniques from both domains in tandem. For example, some of our results use a new composition operation on 0-1 matrices called *grafting*, which has no exact counterpart in the domain of sequences.

*This work is supported by NSF CAREER grant no. CCF-0746673 and a grant from the US-Israel Binational Science Foundation.

SCG'11, June 13–15, 2011, Paris, France.

Categories and Subject Descriptors

F.2.2 [**Nonnumerical Algorithms and Problems**]: Geometrical problems and computations

General Terms

Theory

Keywords

Davenport-Schinzel sequences, geometric graphs, fat triangles

1. INTRODUCTION

A *generalized* Davenport-Schinzel sequence over an n-letter alphabet is one whose subsequences are not isomorphic to some fixed *forbidden subsequence* σ. Let $\mathrm{Ex}(\sigma, n)$ be the extremal function for σ, i.e., the maximum length of such a σ-free sequence. When is $\mathrm{Ex}(\sigma, n)$ linear or nonlinear? and what characteristics of σ let us determine its asymptotic growth? These questions have been answered with startling precision [21] when σ is an alternating sequence $ababab\cdots$ with length $t+2$, also called the *order-t* Davenport-Schinzel sequence. These forbidden subsequences have found numerous geometric applications [2, 33, 5], largely because they relate to the complexity of the lower envelope of curves without $t+1$ pairwise crossings, e.g., degree-t polynomials. What can be said about forbidden subsequences σ with more evolved structure? It is known [15, 21] that $\mathrm{Ex}(\sigma, n)$ is bounded from above by $n \cdot 2^{\mathrm{poly}(\alpha(n))}$, where the polynomial depends on σ. However, tight asymptotic bounds on $\mathrm{Ex}(\sigma, n)$ are only known when σ fits into a couple well-structured classes. Before discussing prior work on generalized Davenport-Schinzel sequences and our contributions we briefly review some standard notation for sequences and 0-1 matrices.

1.1 Definitions and Notation

The length of a sequence is denoted $|\sigma|$. If $\sigma = (\sigma_i)_{0 \le i < |\sigma|}$ is a sequence let $\Sigma(\sigma) = \{\sigma_i\}_i$ be its *alphabet* and $\|\sigma\| = |\Sigma(\sigma)|$ be the alphabet size. Two equal length sequences σ, σ' are *isomorphic*, written $\sigma \sim \sigma'$, if there is a bijection $f : \Sigma(\sigma) \to \Sigma(\sigma')$ for which $f(\sigma_i) = \sigma'_i$. We say σ is a *subsequence* of σ', written $\sigma \mathrel{\bar{\prec}} \sigma'$, if there is a strictly increasing function $f : [|\sigma|] \to [|\sigma'|]$ for which $\sigma_i = \sigma'_{f(i)}$, for $0 \le i < |\sigma|$. Here $[k] = \{0, 1, \ldots, k-1\}$. We write $\sigma \prec \sigma'$ if σ is isomorphic to a subsequence of σ', that is, $\sigma \sim \sigma'' \mathrel{\bar{\prec}} \sigma'$ for some σ''. The phrase σ *appears in* (or *occurs in*) σ' means either $\sigma \prec \sigma'$ or $\sigma \mathrel{\bar{\prec}} \sigma'$, which one should be clear from context. A sequence σ' (or class of sequences) is σ*-free* if $\sigma \not\prec \sigma'$. A sequence σ is k*-sparse* if $\sigma_i = \sigma_j$ implies $|i-j| \ge k$. 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. Absent any knowledge of σ, the predicate $[\![\sigma]\!] = m$ asserts that there is some way to partition σ into at most m blocks. Let $\mathrm{dbl}(\sigma)$ be obtained from σ by doubling each letter save the first and last, e.g., $\mathrm{dbl}(abab) = abbaab$. There are two variants for the extremal function of σ-free sequences, one that specifies the number of blocks (without a sparseness criterion) and another that demands that the sequence be $\|\sigma\|$-sparse. Either criterion ensures that the function is well defined and finite.

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

We say a sequence σ is *linear* or *nonlinear* depending on whether $\mathrm{Ex}(\sigma, n)$ is linear or nonlinear in n. It is *minimally* nonlinear if no strict subsequence of σ is nonlinear.

The terminology for forbidden subsequences can be translated to forbidden 0-1 matrices. Let $S \in \{0,1\}^{n \times m}$ and $P \in \{0,1\}^{k \times l}$ be two matrices. We say P is *contained in* S, or $P \prec S$, if there are two strictly increasing functions $f : [k] \to [n]$ and $g : [l] \to [m]$ such that $P(i,j) = 1$ implies $S(f(i), g(j)) = 1$, i.e., a 0 in P matches either 0 or 1. The two functions f, g define a *submatrix* of S. If P is not contained in S then S if P*-free*. Let $P^{\obar}$, $P^{\ominus}$, $P^{\oslash}$, $P^{\obslash}$, $P^{\ogreaterthan}$, $P^{\odot}$, $P^{\olessthan}$ denote the horizontal, vertical, and diagonal reflections of P, and the right rotations by one, two, and three quarters, respectively. Let $|S|$ be the number of 1s in S, also called its *weight*. Define $\mathrm{Ex}(P, n, m) = \max\{|S| \mid S \in \{0,1\}^{n \times m} \text{ and } P \not\prec S\}$. Following a common convention, we write 0-1 matrices using bullets for 1s and blanks for 0s.

There is now a large body of work devoted to forbidden 0-1 matrices; see [4, 6, 7, 9, 10, 11, 12, 14, 18, 20, 22, 23, 25, 26, 27, 28, 30, 32]. Some of our results require precise bounds on U_1-free and U_2-free matrices. These and other matrices referenced in the paper are defined in Figure 1.

Theorem 1.1. *(Füredi-Hajnal [11] and Tardos [30])*

$$\mathrm{Ex}(U_1, n, m) < 2n + 2m$$
and
$$\mathrm{Ex}(U_2, n, m) = \mathrm{Ex}(\hat{V}_2, n, m) < 5n + m.$$

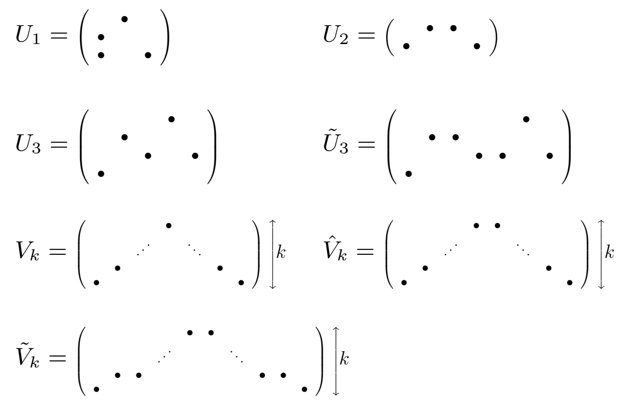

Figure 1: Several 0-1 matrices. By convention 1s and 0s are represented by bullets and blanks.

1.2 Linearity and Nonlinearity in Generalized Davenport-Schinzel Sequences

When $\sigma \in \{a, b\}^*$ is over a 2-letter alphabet we now understand the extremal function $\mathrm{Ex}(\sigma, n)$ *almost* perfectly. It is known that $\mathrm{Ex}(ababa, n) = \Theta(n\alpha(n))$ [13, 21] and that $\mathrm{Ex}(\mathrm{dbl}(abab), n) = O(n)$ [1, 17, 16], which implies that $ababa$ is the *only* minimally nonlinear sequence over two letters. Agarwal, Sharir, and Shor [3] proved tight bounds on $\mathrm{Ex}(ababab, n) = \Theta(n2^{\alpha(n)})$, and, very recently, Nivasch [21] proved essentially tight bounds on even-order Davenport-Schinzel sequences and nearly tight bounds on odd-order sequences:

$$\mathrm{Ex}(ababa, n) = \Theta(n\alpha(n)) \tag{1}$$
$$\mathrm{Ex}(ababab, n) = \Theta(n \cdot 2^{\alpha(n)}) \tag{2}$$
$$\mathrm{Ex}((ab)^{t+2}, n) = n \cdot 2^{(1 \pm o(1))\alpha(n)^t / t!} \tag{3}$$
$$\mathrm{Ex}((ab)^{t+2}a, n) \le n \cdot 2^{(1+o(1))\alpha(n)^t \log \alpha(n)/t!} \tag{4}$$

Pettie [27] showed that doubling an alternating sequence generally does not influence it's extremal function. In particular, the upper bounds in (2,3,4) continue to hold for $\mathrm{dbl}(ababab)$, $\mathrm{dbl}((ab)^{t+2})$, and $\mathrm{dbl}((ab)^{t+2}a)$. However, for $\mathrm{dbl}(ababa)$ the best known upper bound on $\mathrm{Ex}(\mathrm{dbl}(ababa), n)$ is $O(n\alpha^2(n))$.

What can we say about forbidden subsequences over larger alphabets? Klazar and Valtr [17] showed that N-shaped sequences of the form $\mathrm{dbl}(a_1 \cdots a_{k-1} a_k a_{k-1} \cdots a_2 a_1 a_2 \cdots a_k)$ are linear and that embedding one linear sequence in another results in a linear sequence. Specifically, if $\mathbf{u} = \mathbf{u}_1 aa \mathbf{u}_2$ and $\mathbf{v}$ are both linear and $\Sigma(\mathbf{u}) \cap \Sigma(\mathbf{v}) = \emptyset$ then $\mathbf{u}_1 a \mathbf{v} a \mathbf{u}_2$ is also linear. Using results on forbidden double permutation matrices [18, 12], Pettie [26] showed that $\sigma = \pi_1 \mathrm{dbl}(\pi_2)$ is linear for any two permutations π_1, π_2 of $\Sigma(\sigma)$, e.g., $\sigma = abcdeaddcceeb$ is such a sequence. Prior to the present work, all sequences known to be linear could be derived (via embeddings) from N-shaped sequences and double-permutation sequences.

Nivasch's nonlinear bounds on standard Davenport-Schinzel sequences are actually corollaries of a more general forbidden substructure theorem. Let $\mathrm{Perm}(r, s)$ be the set of all sequences of the form $\pi_1 \cdots \pi_s$, where each π_i is a permutation over r letters, e.g., $[abcd][acbd][dcab] \in \mathrm{Perm}(4, 3)$. A sequence is Perm$(r, s)$-free if it avoids *every* $\sigma \in$ Perm(r, s). It is shown [21, 15] that Ex(Perm$(r, s), n, m)$ is:

$$\begin{array}{ll} = O(n\alpha(n,m)) & \text{for } s = 4 \\ = O(n2^{\alpha(n,m)}) & \text{for } s = 5 \\ \leq n \cdot 2^{(1+o(1))\alpha^t(n,m)/t!} & s \geq 7 \text{ odd}, t = \frac{s-3}{2} \\ \leq n \cdot 2^{(1+o(1))\alpha^t(n,m)\log\alpha(n)/t!} & s \geq 6 \text{ even}, t = \frac{s-4}{2} \end{array} \quad (5)$$

This theorem gives a general upper bound on Ex(σ, n) since any σ-free sequence is necessarily Perm$(\|\sigma\|, |\sigma| - \|\sigma\| + 1)$-free; see [21, Lemma 1.4]. Note that the upper bounds in (1–4) are special cases of (5).

1.3 New Results

One implication of [1, 17] is that any sequence $\sigma \in \{a, b, c\}^*$ avoiding *ababa*, *abcacbc*, *abcbcac* or their reversals must be linear.[1] In Section 2 we show that *abcbcac* is linear and that Ex$(abcacbc, n) = O(n\alpha(n))$. Together with a matching lower bound on *abcacbc*-free sequences [27] it follows that a repetition-free sequence over three letters is linear *if and only if* it avoids *ababa*, *abcacbc* or its reversal. Thus, we have a nearly perfect understanding of linear forbidden subsequences over both two- and three-letter alphabets.

In Section 3 we show that the maximum length of a sequence avoiding *ababab* and any M-shaped sequence (of the form $ab \cdots yzzy \cdots baab \cdots yzzy \cdots ba$) has length $O(n\alpha(n))$, which, as a special case, implies that $\{ababab, \text{dbl}(ababa)\}$-free sequences have length $O(n\alpha(n))$. As an application of this theorem we show that the complexity of the union of n δ-fat triangles is $O(n \log^* n\alpha(n))$, which slightly improves a bound of $O(n \log^* n2^{\alpha(n)})$ by Ezra, Aronov, and Sharir [8]. Specifically, we prove that the complexity of the union of n nearly-isosceles right triangles, all of which intersect the x-axis, have complexity $O(n\alpha(n))$. In the conclusion of [8] Ezra et al. claim that the $2^{\alpha(n)}$ factor can be eliminated altogether using a different approach, that is, they do not prove an $O(n)$ bound on the type of arrangements we consider. It is an open problem whether $O(n\alpha(n))$ is tight for nearly-isosceles right triangles on the x-axis.

In Section 4 we introduce a new composition operation on forbidden 0-1 matrices called *grafting* and prove a lemma on the extremal functions of 0-1 matrices formed by multiple grafting operations. Among other corollaries, the grafting lemma implies that all N-shaped forbidden subsequences are linear. The leading constants in our linear bound are significantly smaller than [17]. They are both exponential in the alphabet size of the N-shaped sequence, with base of 5 in our case and about 1440 in [17]. These bounds immediately yield tighter $O(n \log n)$ bounds on the maximum number of edges in a geometric graph (and graphs whose edges are x-monotone curves) with no k mutually crossing edges; see Valtr [31].

2. FORBIDDEN SEQUENCES OVER THREE LETTERS

We obtain a nearly complete characterization of linear forbidden sequences over three letters. Theorem 2.1 is a consequence of prior work [13, 17, 1, 27] and Theorems 2.3, 2.4, and 2.5.

Theorem 2.1. *Let $\sigma \in \{a, b, c\}^*$ be a sequence on three letters.*

1. *The sequences ababa and abcacbc are minimally nonlinear and the only 2-sparse minimally nonlinear sequences over three letters.*

2. Ex(σ, n) *is $\Omega(n\alpha(n))$ if σ contains ababa or abcacbc and is $\Theta(n\alpha(n))$ if $\sigma \in \{ababa, abcacbc\}$.*

3. *If σ avoids ababa, abcacbc, and the three sequences obtained from a$\underline{b}\underline{c}bc\underline{a}$c by doubling one of the underlined symbols, then* Ex$(\sigma, n) = O(n)$.

The $\Omega(n\alpha(n))$ lower bounds can be found in [13, 27]. We prove that Ex$(abcacbc, n) = O(n\alpha(n))$, that *abcbcac* is linear, and, in fact, that *abcbbccac* is linear as well. Many of our proofs transform sequences into 0-1 matrices, usually in *canonical* form.

Definition 2.2. **(Canonical Form)** *Let $S = s_1 \cdots s_m$ be an m-block sequence over an n-symbol alphabet. The* canonical matrix *of S, denoted $A = A(S)$, is an $n \times m$ 0-1 matrix obtained by ordering $\Sigma(S)$ according to the* first appearance *in S, then letting $A(i, j) = 1$ if and only if the ith symbol appears in s_j.*

Theorem 2.3. Ex$(abcbcac, n) < 42n$ *and* Ex$(U_3, n, m) < 7n + 5m$.

Proof. Let S be an *abcbcac*-free sequence with length Ex$(abcbcac, n)$. Greedily partition $S = s_1 s_2 \cdots s_m$ into maximal *bcbcac*-free sequences (s_i), i.e., s_1 is the longest *bcbcac*-free prefix of S, s_2 is the longest *bcbcac*-free prefix of the remaining sequence, and so on. Since each s_i contains the first occurrence of some symbol, namely the 'a' in *bcbcac*, $m < n$. Let $S' = \Sigma(s_1)\Sigma(s_2) \cdots \Sigma(s_m)$ (i.e., replace each s_i by its alphabet $\Sigma(s_i)$, listed according to its order in s_i) and let $A = A(S')$ be the $n \times m$ canonical matrix for S'. Since $s_i \leq$ Ex$(bcbcac, \|s_i\|) \leq 3.5\|s_i\|$, $|S| \leq 3.5|S'|$.[2] If A contains U_3 this implies that S contains an ordered subsequence isomorphic to 42313, and, since A is canonical, that S contains $1232313 \sim abcbcac$. We will show that $|A| \leq$ Ex$(U_3, n, m) < 7n + 5m$, and therefore that Ex$(abcbcac, n) \leq 3.5 \cdot$ Ex$(U_3, n, n) = 42n$.

The remainder of the proof is structured as follows. Given A, we construct a set $\mathscr{Q}$ of overlapping boxes (contiguous submatrices) then convert $\mathscr{Q}$ into a set $\mathscr{R}$ of disjoint boxes with several properties: (i) after removing $3n$ 1s, no row or column has a non-zero intersection with more than one box in $\mathscr{R}$, (ii) each matrix in $\mathscr{R}$ is U_1-free, and (iii) the number of 1s not contained in any box is less than $2n + 3m$. By Theorem 1.1 the total number of 1s is $7n + 5m$.

To construct the set $\mathscr{Q}$ we examine each 1 in increasing order by column then increasing order by row. Let (i, j)

[1] A case analysis shows that these are the minimal sequences not contained in dbl$(abcbabc)$ and dbl$(abcbca)$, which are linear [1, 17].

[2] To see this, observe that any 3-sparse *bcbcac*-free sequence is also *bcbcc*-free as well; its 3-sparseness guarantees that there must be some a distinct from b and c located between the last two cs. We remove the last occurrence of each symbol in the sequence, then remove up to $n/2$ repetitions to restore 2-sparseness. The 3-sparseness of the original sequence guarantees that $n/2$ suffices. Thus, the length of the original sequence is at most $3n/2 +$ Ex$(bcbc, n) < 3.5n$.

be the current 1 and let $\mathscr{Q}$ be the set of boxes obtained so far. If (i,j) is the first 1 in its column, skip to the next 1. If (i,j) already lies in a box in $\mathscr{Q}$ then skip to the next 1. Otherwise let $(i',j') \in A$ be the 1 in A maximizing i' such that $j' < j$ and $i' > i$; if there is no such 1 then skip to the next 1. Include in $\mathscr{Q}$ the box $(i,i') \times (j,\infty)$. (Here $(x,y) = \{x+1,\ldots,y-1\}, [x,y) = \{x,\ldots,y-1\}$, etc.) Let $\mathscr{Q} = \{Q_1, Q_2, \ldots\}$ be the set of boxes in the order they were included in $\mathscr{Q}$. Let the set of boxes $\mathscr{R} = \{R_1, R_2, \ldots\}$ be such that $R_k = Q_k \backslash \bigcup_{l>k} Q_l$. Clearly boxes in $\mathscr{R}$ are disjoint. See Figure 2(A,B) for an example.

Before moving on we note that the matrix of 1s outside $\mathscr{R}$ is L-free, where $L = \begin{pmatrix} & \bullet & \\ & \bullet & \bullet \\ \bullet & & \end{pmatrix}$, and therefore has weight less than $2n+3m$. If there were such an L outside $\mathscr{R}$, the 1 in the third column would have been placed in a box when the second 1 in the second column was examined.

Let $(i_k, j_k), (i'_k, j'_k)$ be the 1s in A defining the dimensions of Q_k and R_k, i.e., R_k is of the form $(i_k, i'_k) \times (j_k, *)$. Let $f(j)$ be the row of the first 1 in column j.

Let $\hat{A}$ be derived from A be removing all 1s not contained in $\mathscr{R}$ and removing the first two 1s and last 1 in each row. We claim that no row in $\hat{A}$ has a nonzero intersection with more than one box. Suppose, to the contrary, that (i,j) and (i,j') are 1s in boxes R_q and R_r, where $j < j'$ and $q < r$. Figure 2(C) gives an example with (i,j) and (i,j') underlined. If $j < j_r$ (not depicted in Figure 2(C)) then the points $(i'_q, j'_q), (i_q, j_q), (i,j), (f(j_r), j_r), (i,j')$ form an instance of U_3. If $j = j_r$ (as in Figure 2(C)) then let $(i,j'') \in A$ be the first 1 in row i intersecting a box, say R_p. Then the 1s at positions $(i'_p, j'_p), (i_p, j_p), (i,j''), (f(j_r), j_r), (i,j')$ form an instance of U_3. Observe that R_p, R_q, and R_r may all have the same upper boundary (contrary to the depiction in Figure 2(C)), requiring us to use the point $(f(j_r), j_r)$ rather than (i_r, j_r) since it may be that $i_p = i_q = i_r$. We claim, further, that no column in $\hat{A}$ has a nonzero intersection with more than one box. Again, suppose to the contrary that (i,j) appears in box R_q and (i',j) in R_p, where $i' < i$ and $p < q$; see Figure 2(D). In A, (i,j) must appear between 1s at (i,j') and (i,j''), where $j' < j < j''$. The point (i,j'') might appear outside R_q but (i,j') will be in R_q, for if the two 1s in A preceding (i,j) lie in another box, they would create an instance of U_3, as in Figure 2(C). Thus, the 1s at positions $(i'_q, j'_q), (i_q, j_q), (i,j'), (i',j), (i,j'')$ form an instance of U_3. Finally, each box is clearly U_1-free since $\hat{A}$ omits the first two 1s in each row. Together with (i'_p, j'_p) these 1s form an instance of U_3. See Figure 2(E).

The row- and column-disjointness properties of $\hat{A}$ and the U_1-freeness of each box imply that $|\hat{A}| \le \mathrm{Ex}(U_1, n, m) < 2n + 2m$. Thus, the number of 1s in A contained in $\mathscr{R}$ is less than $5n + 2m$ and $|A| < 7n + 5m$. $\Box$

We are unable to show that dbl(*abcbcac*) is linear, though doubling the second *bc* does not affect the linearity of *abcbcac*. The proof of Theorem 2.4 appears in the appendix.

THEOREM 2.4. *$\mathrm{Ex}(abcbbccac, n) < 198n$ and $\mathrm{Ex}(\tilde{U}_3, n, m) < 11n + 7m$.*

It is possible to get an $O(n\alpha(n))$ upper bound on *abcacbc*-free sequences via a forbidden 0-1 matrix argument [27], but the proof is complex. We demonstrate that this bound follows directly from Equation 5.

THEOREM 2.5. *$\mathrm{Ex}(abcacbc, n) = O(n\alpha(n))$.*

PROOF. Consider the ways in which *abcacbc* might appear in some $\gamma = \pi_1\pi_2\pi_3\pi_4 \in \mathrm{Perm}(3,4)$, where, without loss of generality, $\pi_1 = abc$. It must be that (i) c precedes a in π_2, that (ii) b precedes c in π_3, and that (iii) c precedes b in π_4. If (i) fails to hold then $[abc][ac][b][c] \prec \gamma$; if (ii) fails to hold then $[abc][a][cb][c] \prec \gamma$; and if (iii) fails to hold then $[abc][a][c][bc] \prec \gamma$. If π_2 is either *cba* or *bca* then $[ac][ba][bc][b] \sim abcacbc \prec \gamma$; thus π_2 must be *cab*. In the same way we can deduce that for $\pi_3 = bca$ and $\pi_4 = acb$, $\gamma = [abc][cab][bca][acb]$ avoids *abcacbc*, so we have apparently failed to obtain a contradiction. However, we can show there is no $\gamma' = \pi'_1\pi'_2\pi'_3\pi'_4 \in \mathrm{Perm}(4,4)$ avoiding *abcacbc*. For each of the $\binom{4}{3}$ distinct $\mathrm{Perm}(3,4)$ sequences contained in γ', the first two blocks must be isomorphic to $[abc][cab]$. Thus, if $\pi'_1 = abcd$, π'_2 must contain both *cab* and *dbc*, an impossibility. $\Box$

3. M-SHAPED SEQUENCES AND FAT TRIANGLES

The M-shaped sequence μ_k generalized '*ababa*' to an alphabet of size k, where

$$\mu_k = a_1 \cdots a_{k-1} a_k^2 a_{k-1} \cdots a_2 a_1^2 a_2 \cdots a_{k-1} a_k^2 a_{k-1} \cdots a_2 a_1.$$

It is an open question whether there is a universal upper bound on μ_k-free sequences, i.e., $O(n 2^{\alpha^t(n)})$ for some t independent of k. In Theorem 3.1 we prove tight $\Theta(n\alpha(n))$ bounds on sequences that are both *ababab*-free and μ_k-free and in Theorem 3.2 we explain how this implies tighter bounds on the complexity of the union of δ-fat triangles.

THEOREM 3.1. *For any $k \ge 2$, $\mathrm{Ex}(\{\mu_k, ababab\}, n, m) = \Theta(n\alpha(n,m) + m)$. Observe that in the special case of $k = 2$, $\mathrm{Ex}(\{\mathrm{dbl}(ababa), ababab\}, n, m) = \Theta(n\alpha(n,m) + m)$.*

PROOF. We show that the claim follows from Equation 5. Suppose there is a $\gamma \in \mathrm{Perm}(k,4)$ avoiding both μ_k and *ababab*. Restricting our attention to two arbitrary symbols $a, b \in \Sigma(\gamma)$, each permutation in γ orders them as $[ab]$ or $[ba]$. If any type is repeated thrice then clearly $ababab \prec \gamma$. The remaining possibilities for γ are $[ab][ab][ba][ba]$, $[ab][ba][ba][ab]$, and $[ab][ba][ab][ba]$. The first two contain *ababab*. Since a and b are arbitrary symbols, *every* pair must be arranged in the third pattern, i.e., γ can only be $[a_1 \cdots a_k]$ $[a_k \cdots a_1]$ $[a_1 \cdots a_k]$ $[a_k \cdots a_1]$, which is μ_k. Thus, any $\{ababab, \mu_k\}$-free sequence is necessarily $\mathrm{Perm}(k,4)$-free as well. $\Box$

THEOREM 3.2. *(cf. Ezra, Aronov, and Sharir [8]) The complexity of the union of n δ-fat triangles in the plane is $O(n \log^* n \alpha(n) + n\delta^{-1} \log^2(\delta^{-1}))$.*

At one point in Ezra et al.'s [8] analysis they require a bound on the complexity of the union of n triangles with several restrictions: (i) each is a right, axis aligned triangle, (ii) each triangle intersects the x-axis, and (iii) each is nearly isosceles; the hypotenuse is between, say, 134 and 136 degrees. Matoušek et al. [19] showed that the portion of the boundary lying above the x-axis has at most $O(n)$ points, and in general, that it is sufficient to replace (ii) with (ii$'$): each triangle's uppermost point is on the x-axis. They proved that the complexity of such a union is $O(n2^{\alpha(n)})$ (the source of the $2^{\alpha(n)}$ factor in Ezra et al. [8]) using bounds on order-4 Davenport-Schinzel sequences. We prove an $O(n\alpha(n))$ bound by following Matoušek et al.'s

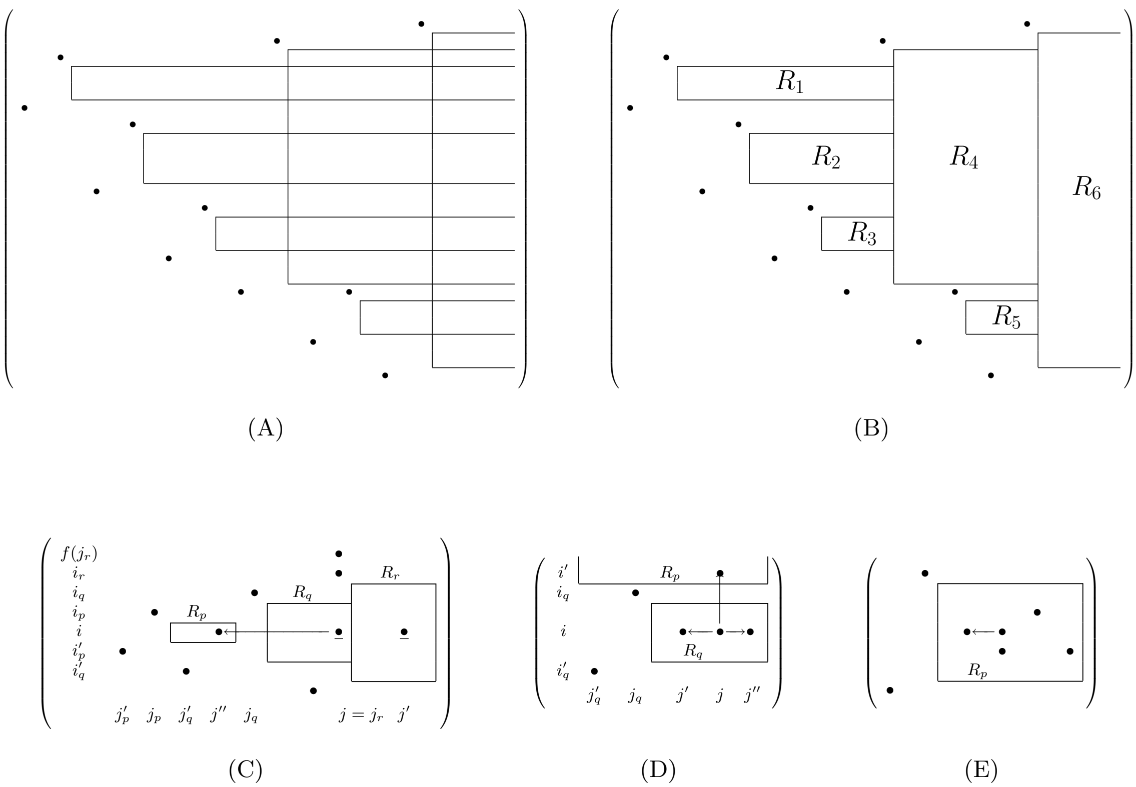

Figure 2: (A) The set $\mathscr{Q}$ of overlapping boxes. The two 1s defining the dimensions of each box are indicated. (B) A partition into non-overlapping boxes $\mathscr{R}$. (C) No row in $\hat{A}$ has a 1 in two distinct $\mathscr{R}$-boxes. (D) No column in $\hat{A}$ has a 1 in two distinct $\mathscr{R}$-boxes. (E) Every $\mathscr{R}$-box is U_1-free. Arrows indicate 1s that can be inferred to exist.

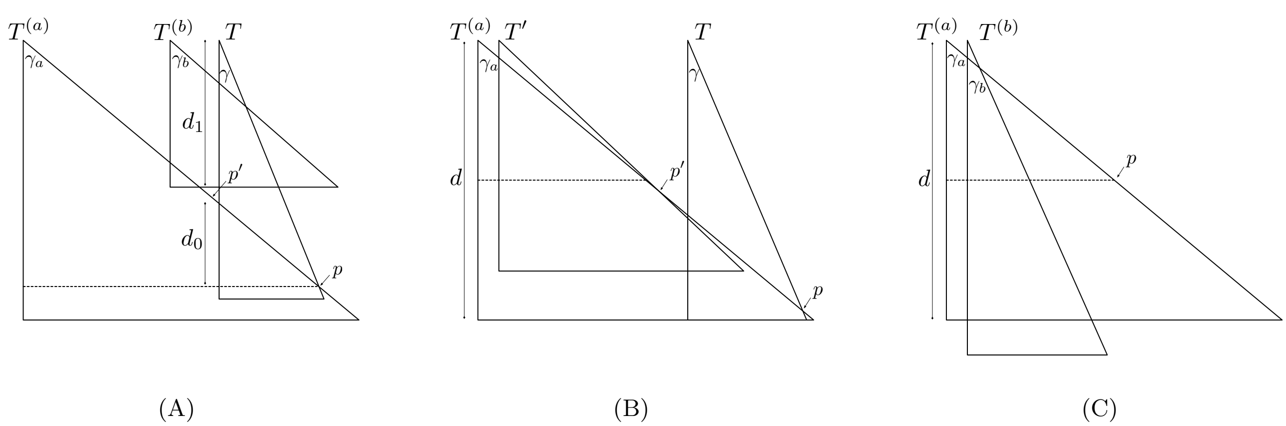

Figure 3: The points p' and p on the hypotenuse of $T^{(a)}$ intersect the hypotenuses of T' and T, respectively; T' does not appear in the figure.

proof but concluding with a different forbidden substructure argument.

Matoušek et al. [19] show that there are $O(n)$ intersections on the boundary involving the horizontal and/or vertical edges of some triangle. It suffices to bound the number of boundary points where two hypotenuses meet. Let $T^{(1)}, T^{(2)}, \ldots, T^{(n)}$ be the triangles sorted in increasing order by the x-coordinate of their uppermost points. If a point p on the boundary meets the hypotenuses of $T^{(i)}$ and $T^{(j)}$, with $i < j$, it is *labeled* i, i.e., the edge on the boundary just below p belongs to $T^{(i)}$'s hypotenuse. Let $S \in \{1, \ldots, n\}^*$ be the concatenation of the labels of all boundary points involving two hypotenuses, where the points are ordered from left to right.

LEMMA 3.3. *For any $a < b$, if the horizontal edge of $T^{(a)}$ lies below that of $T^{(b)}$ then $baab \bar{\not\prec} S$.*

PROOF. For $baab$ to appear in S there must be two triangles T', T whose hypotenuses intersect $T^{(a)}$'s on the boundary, at, say points p' and p, which lie under the horizontal edge of $T^{(b)}$. See Figure 3(A). Without loss of generality T' precedes T. It follows that p' must precede the vertical edge of T (otherwise p' would not be on the boundary) and therefore that $T^{(b)}$ precedes T. Let γ_a, γ_b, and γ be the angles of $T^{(a)}, T^{(b)}$, and T opposite their horizontal edges. Let d_0 be the vertical distance between the two intersections of T with $T^{(a)}$'s hypotenuse and let d_1 be the height of $T^{(b)}$. Thus $d_1 \tan \gamma_b > d_0 \tan \gamma_a > (d_0 + d_1) \tan \gamma$. The first inequality follows since $baab \bar{\prec} S$ and the second because $d_0 + d_1$ is less than the vertical distance from p to the top of T. This is clearly impossible since γ_a, γ_b, and γ are close to 45 degrees. □

LEMMA 3.4. *A boundary point labeled a is called* lower *if it appears in the lower half of $T^{(a)}$. No two lower points have the same label.*

PROOF. Suppose $T^{(a)}$ has two lower points, p' and p, lying on the hypotenuses of T' and T, respectively, where p' precedes the vertical edge of T. See Figure 3(B). Let d be the height of $T^{(a)}$ and let γ_a, γ be the angles of $T^{(a)}$ and T opposite their horizontal edges. It must be that $(d/2) \tan \gamma_a + d \tan \gamma$ is strictly less than the length of the horizontal edge of $T^{(a)}$, namely $d \tan \gamma_a$, which is impossible since γ_a and γ are close to 45 degrees. □

LEMMA 3.5. *Let S' be the result of removing the last occurrence of each symbol in S. For any $a < b$, if the horizontal edge of $T^{(a)}$ lies above that of $T^{(b)}$ then $baba \bar{\not\prec} S'$.*

PROOF. It follows from Lemma 3.4 that S' represents only non-lower boundary points. Let d be the height of $T^{(a)}$ and let p be the point on $T^{(a)}$'s hypotenuse whose vertical distance to the top of $T^{(a)}$ is $d/2$. Let γ_a and γ_b be defined as usual. See Figure 3(C). If $baba$ appears in S' then some portion of $T^{(b)}$'s hypotenuse lies below the horizontal edge of $T^{(a)}$ and to the left of p. Thus $(d/2) \tan \gamma_a > d \tan \gamma_b$, which is impossible since γ_a and γ_b are close to 45 degrees. □

PROOF. (Theorem 3.2) Let S' be defined as in Lemma 3.5. It follows from Lemmas 3.3 and 3.5 that S' is both $abaaba$-free and $ababab$-free. (These are the unique minimal sequences that contain $abba, baab, abab$, and $baba$.) It also follows that S' is composed of less than $2n$ blocks since different occurrences of a symbol must be separated by the beginning or end of some triangle. Thus, S' is at most $\mathrm{Ex}(\{\mathrm{dbl}(ababa), ababab\}, n, 2n)$, which is $O(n\alpha(n))$ by Theorem 3.1 and Equation 5. Note that $abaaba$-freeness alone might be sufficient to get an $O(n\alpha(n))$ bound; however, the best bound [27] on $\mathrm{Ex}(\mathrm{dbl}(ababa), n)$ is only $O(n\alpha^2(n))$. □

4. N-SHAPED SEQUENCES AND THE GRAFTING LEMMA

We introduce a composition operation on 0-1 matrices that preserves the linearity of the constituent matrices. This operation resembles or subsumes Klazar and Valtr's composition operations for sequences [17] and Keszegh's two composition operations on 0-1 matrices [14]. The linearity of N-shaped sequences follows almost immediately.

DEFINITION 4.1. **(Grafting)** *Let P_0 be a $k \times l$ matrix with 1s in the southwest and southeast corners, i.e., $P_0(k, 1) = P_0(k, l) = 1$, and let P_1 be a $k' \times l'$ matrix with adjacent 1s in the top row, i.e., $P_1(1, r) = P_1(1, r+1) = 1$ for some $1 \le r < l'$. Let Q be the $(k+k'-1) \times (l+l'-1)$ matrix obtained by identifying $P_1(1, r)$ with the 1 in the southwest corner of a copy of P_0, i.e., the submatrix of Q on the positions $[1, k] \times [r, r+l)$ is P_0, the submatrix of Q on the positions $[k, k+k') \times ([1, r] \cup [r+l, l+l'))$ is P_1 and Q is zero elsewhere.*

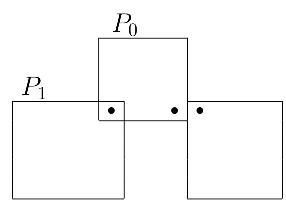

DEFINITION 4.2. **(Legality)** *A matrix R is* descending *if $R(i, j) = R(i', j') = 1$ and $i < i'$ imply $j < j'$ and* ascending *if they imply $j' < j$. A matrix S is* non-descending *(or* non-ascending*) if no interval of S's rows is descending (or ascending). A matrix is* legal *if it is either non-descending or non-ascending.*

LEMMA 4.3. **(The Grafting Lemma)** *Let P_0, P_1 be matrices satisfying the preconditions of Definition 4.1 where, in addition, P_0 is legal. Let Q be obtained by grafting P_0 onto P_1. If $\mathrm{Ex}(P_0, n, m) \le c_0 n + c_0' m$ and $\mathrm{Ex}(P_1, n, m) \le c_1 n + c_1' m$ then $\mathrm{Ex}(Q, n, m) \le cn + c'm$ where $c = (c_0+1)(c_1+1) - 1$ and $c' = (c_0 + 1)(c_1' + c_0')$.*

PROOF. Let A be an $n \times m$ Q-free matrix with weight $|A| = \mathrm{Ex}(Q, n, m)$. Choose a parameter $g > c_0$. Partition the 1s in each row of A into consecutive groups of g 1s, leaving up to $(g-1)n$ 1s ungrouped, $g - 1$ per row. We form a matrix $\tilde{A}$ from A by assigning each group to a distinct row in the following way. Let h_i be the number of groups in row i and $h_{<i} = \sum_{i'<i} h_{i'}$ be the number of groups in rows preceding i. If P_0 is non-descending then group j of row i in A is assigned to row $h_{<i} + j$ of $\tilde{A}$, whereas if P_0 is non-ascending this group is assigned to row $h_{<i} + h_i + 1 - j$ of $\tilde{A}$. It follows that $\tilde{A}$ is a $\lfloor \mathrm{Ex}(Q, n, m)/g \rfloor \times m$ matrix. Note that since P_0 is legal, a set of rows in $\tilde{A}$ containing an occurrence of P_0 must correspond to *distinct* rows in A, i.e., occurrences of P_0 in $\tilde{A}$ map injectively to occurrences in A.

Call a 1 in $\tilde{A}$ *good* if it is in the southwest corner of an instance of P_0 and call a row in $\tilde{A}$ bad if it contains no good 1s. The submatrix of bad rows is by definition P_0-free and contains exactly bg 1s, so $bg \le \mathrm{Ex}(P_0, b, m) \le c_0 b + c_0' m$, which implies $b < c_0' m/(g - c_0)$. We form an $n \times m$ matrix $\hat{A}$ that contains exactly one 1 from each group in A that corresponds to a good 1 in $\tilde{A}$. We claim $\hat{A}$ is P_1-free. Any two 1s at positions $P_1(1, r)$ and $P_1(1, r+1)$ in an occurrence of P_1 are, by definition, in different groups and, since $P_1(1, r)$ corresponds to a good 1 in $\tilde{A}$, it can be identified as the southwest corner of an instance of P_0 in A, all of whose columns strictly precede the 1 at the position corresponding to $P_1(1, r+1)$. Putting everything together we have:

$$\begin{aligned}
|A| &\le g|\hat{A}| + (g-1)n + gc_0' m/(g - c_0) \\
&\le g \cdot \mathrm{Ex}(P_1, n, m) + (g-1)n + gc_0' m/(g - c_0) \\
&\le g(c_1 n + c_1' m) + (g-1)n + gc_0' m/(g - c_0) \\
&= (g(c_1 + 1) - 1) \cdot n + g(c_1' + c_0'/(g - c_0)) \cdot m \\
&= ((c_0 + 1)(c_1 + 1) - 1) \cdot n + (c_0 + 1)(c_1' + c_0') \cdot m \\
&= cn + c' m
\end{aligned}$$

where the second to last inequality follows by setting $g = c_0 + 1$. $\square$

The grafting lemma implies that some rather complex-looking matrices are, in fact, linear.

Corollary 4.4. *The following matrices are linear:*

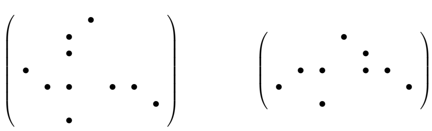

The first matrix is obtained by grafting U_1 onto the top of $U_1^{\ominus}$, then grafting $U_1^{\ominus}$ onto the left side of *that* matrix. The second is obtained by grafting $U_1^{\oplus}$ onto the top of the matrix shown to be linear by Fulek [9]. The grafting lemma can easily be adapted to the case when c_0, c_0', c_1, c_1' are not constants but functions of n and m. For example, it is shown in [27] that $\mathrm{Ex}((\cdot\,\cdot\cdot\cdot\,\cdot\,\cdot), n, m) = O(n\alpha^2(n, m) + m)$ and $\mathrm{Ex}((\cdot\,\cdot\cdot\,\cdot\,\cdot\,\cdot\cdot\,\cdot), n, m) = O(n 2^{\alpha(n,m)} + m)$. Theorem 4.5 follows from Lemma 4.3, [27, Theorems 3.3, 3.5], and the standard sequence-to-canonical matrix conversion.

Theorem 4.5. $\mathrm{Ex}(\sigma_1, n, m)$, $\mathrm{Ex}(\sigma_2, n, m)$, $\mathrm{Ex}(W_1, n, m)$, *and* $\mathrm{Ex}(W_1, n, m)$ *are all* $O((n+m)\alpha^2(n, m) 2^{\alpha(n,m)})$, *where*

$$\sigma_1 = a\,\mathrm{dbl}(bcbababcb),$$
$$\sigma_2 = ab\,\mathrm{dbl}(cdcbababcdc),$$

$W_1 =$ ,

and

$W_2 =$ 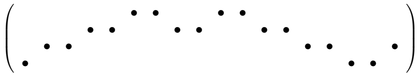.

Observe that Theorem 4.5's bounds are tight to within an $O(\alpha^2(n, m))$ factor, which is remarkable given the apparent complexity of the forbidden sequences and matrices. We now show that the grafting lemma yields linear bounds on N-shaped sequences, improving the leading constants of [17].

Theorem 4.6. *Define* $\nu_k = a_1 a_2 \cdots a_k a_{k-1} \cdots a_1 a_2 \cdots a_k$, $\hat{\nu}_k$ *to be* ν_k *with the first* a_k *and second* a_1 *doubled, and* $\tilde{\nu}_k$ *to be* ν_k *with the last two thirds doubled, i.e.,*

$$\hat{\nu}_k = a_1 a_2 \cdots a_k^2 a_{k-1} \cdots a_1^2 a_2 \cdots a_k$$
$$\tilde{\nu}_k = a_1 \cdots a_{k-1} a_k^2 a_{k-1}^2 \cdots a_2^2 a_1^2 a_2^2 \cdots a_{k-1}^2 a_k$$

The following bounds hold, for $k \ge 3$*:*

1. $\mathrm{Ex}(\nu_k, n, m) \le 2 \cdot \mathrm{Ex}(\hat{V}_{k-1}, n, m) + m \le 2 \cdot 5^{k-2} n + 5^{k-2} m$.
2. $\mathrm{Ex}(\hat{\nu}_k, n, m) \le \mathrm{Ex}(\hat{V}_k, n, m) + n \le 5^{k-1} n + 5^{k-1} m/2$.
3. $\mathrm{Ex}(\tilde{\nu}_k, n, m) \le \mathrm{Ex}(\tilde{V}_k, n, m) + n \le 7^{k-1} n + 7^{k-1} m/3$.
4. $\mathrm{Ex}(\mathrm{dbl}(\nu_k), n, m) \le \frac{3}{2} 5^{\binom{k}{2}} n + 4 \cdot 5^{\binom{k}{2}} m$

Proof. Let S be an n-letter m-block ν_k-free sequence and let $A = A(S)$ be its canonical matrix. A is clearly V_k-free and if we remove the first 1 in each column and every other 1 in each row the matrix becomes $\hat{V}_{k-1}$-free. Similarly, if S is $\hat{\nu}_k$-free and we remove the first 1 in each row of A, the resulting matrix is $\hat{V}_k$-free. Finally, if S is $\tilde{\nu}_k$-free and we remove the first 1 in each row, the resulting matrix is $\tilde{V}_k$-free. To prove Parts (1–3) it suffices to prove bounds on $\hat{V}_k$-free and $\tilde{V}_k$-free matrices. Part (4), however, requires more than a simple forbidden matrix argument.

For the induction hypothesis we actually consider something slightly stronger. Let $\hat{V}_k'$ be the pair of matrices including $\hat{V}_k$ and the matrix obtained from it by contracting its last two columns, e.g., $\hat{V}_2' = \{\hat{V}_2, (\cdot\,\cdot\,\vdots)\}$. Clearly $\mathrm{Ex}(\hat{V}_k, n, m) \le \mathrm{Ex}(\hat{V}_k', n, m) + n$. We will show $\mathrm{Ex}(\hat{V}_k', n, m) \le \hat{c}_k n + \hat{c}_k' m$ and $\mathrm{Ex}(\tilde{V}_k, n, m) \le \tilde{c}_k n + \tilde{c}_k' m$ for sequences $(\hat{c}_k), (\hat{c}_k'), (\tilde{c}_k)$, and $(\tilde{c}_k')$ to be determined. It is known that $\hat{c}_2 = 4, \tilde{c}_2 = 5$ and $\hat{c}_2' = \tilde{c}_2' = 1$; see [11]. Observe that for $k \ge 3$, $\hat{V}_k'$ can be obtained by grafting $\hat{V}_2'$ onto $\hat{V}_{k-1}'$ and removing a spurious 1 in the second row.[3] By Lemma 4.3, $\hat{c}_k = 5\hat{c}_{k-1} + 4$ and $\hat{c}_k' = 5(\hat{c}_{k-1}' + 1)$. By induction $\hat{c}_k = 5^{k-1} - 1$ and $\hat{c}_k' = (2 \cdot 5^{k-1} - 5^{k-2} - 5)/4$. This proves Parts (1) and (2). To apply Lemma 4.3 to $\mathrm{Ex}(\tilde{V}_k, n, m)$, observe that $\tilde{V}_k$ is obtained by grafting $\tilde{P} = (\cdot\,\cdot\,\cdot\,\cdot\,\cdot)$ onto $\tilde{V}_{k-1}$. It is known that $\mathrm{Ex}(\tilde{P}, n, m) < 6n + m$ [11], so $\tilde{c}_k$ and $\tilde{c}_k'$ obey the following equalities, for $k \ge 3$: $\tilde{c}_k = 7\hat{c}_{k-1} + 6$, and $\tilde{c}_k' = 7(\tilde{c}_{k-1}' + 1)$. It follows by induction that $\tilde{c}_k = 6 \cdot 7^{k-2} - 1$ and $\tilde{c}_k' = (2 \cdot 7^{k-1} - 7^{k-2} - 7)/6$. This proves Part (3).

We now turn to part (4). Let S be a $\mathrm{dbl}(\nu_k)$-free sequence and A the canonical matrix for S. We will form a sparser matrix $\hat{A}$ with a corresponding sequence $\hat{S}$ that is $\mathrm{dbl}(\nu_{k-1})$-free through the same process used in the grafting lemma. Let P_0 be $\hat{V}_k$ with the first column duplicated. It follows from Part (2) that $\mathrm{Ex}(P_0, n, m) \le (5^{k-1} + 1)n + 5^{k-1} m/2$. First, delete the 1st 1 from each row of A and the corresponding letters from S. Choose the parameter $g = 5^{k-1} + 2$. Next assign the remaining 1s in each row of A to groups of g, discarding up to $n(g-1)$ ungrouped 1s and their corresponding letters. Let $\tilde{A}$ be the matrix obtained be assigning each group to a distinct row, as in Lemma 4.3. Identify *good 1s* and *bad rows* as in Lemma 4.3, i.e., bad rows do not occur as the bottom row in an occurrence of P_0. It follows that there are at most $g(5^{k-1}/2)m$ 1s in bad rows. Let $\hat{A}$ be a submatrix of A containing exactly one representative good 1 from

[3]Definition 4.1 only defines grafting with respect to single matrices. The definition extends to the case where P_0 and P_1 are sets of matrices so long as each constituent matrix satisfies all the required properties.

each group and let $\hat{S}$ be the associated sequence. We claim $\hat{S}$ is $\mathrm{dbl}(\nu_{k-1})$-free. Suppose there were such an occurrence $a_2a_3^2\cdots a_k^2a_{k-1}^2\cdots a_3^2a_2^2a_3^2\cdots a_{k-1}^2a_k$. Observe that the 1s in $\hat{A}$ corresponding to the adjacent a_2s were good and necessarily from distinct groups, which indicates the existence of a P_0 anchored at the first 1 whose columns strictly precede the second 1. This implies that the following sequence appears in S:

$$a_2a_3^2\cdots a_k^2a_{k-1}^2\cdots a_3^2a_2 \quad [a_2b_1b_2\cdots b_{k-1}b_{k-1}\cdots b_2b_1a_2] \quad a_2a_3^2\cdots a_{k-1}^2a_k$$

Where $a_2, b_1, \ldots, b_{k-1}$ are distinct symbols and each of b_1 through b_{k-1} precedes a_2 in the canonical ordering. By the pigeonhole principle there is some symbol, call it a_1, in $\{b_1,\ldots,b_{k-1}\}\backslash\{a_3,\ldots,a_{k-1}\}$. Due to the canonical ordering and the fact that we deleted the first occurrence of each symbol, the first a_2 must be preceded by a_1a_2 in S, forming an instance of $\mathrm{dbl}(\nu_k)$ in S, a contradiction. We would like to show that $\mathrm{Ex}(\mathrm{dbl}(\nu_k), n, m) < \bar{c}_kn + \bar{c}'_km$ for constants $\bar{c}_k$ and $\bar{c}'_k$ to be determined. In the base case $\bar{c}_2 = 6$ and $\bar{c}'_2 = 1$ [11]. For $k > 2$ we can bound the size of S as follows:

$$\begin{aligned}|S| &= g|\hat{S}| + gn + g(5^{k-1}/2)m \\ &\le g\cdot \mathrm{Ex}(\mathrm{dbl}(\nu_{k-1}), n, m) + gn + g(5^{k-1}/2)m \\ &\le g(\bar{c}_{k-1}n + \bar{c}'_{k-1}m) + gn + g(5^{k-1}/2)m \\ &= (5^{k-1}+2)(\bar{c}_{k-1}+1)n + (5^{k-1}+2)(\bar{c}'_{k-1} + 5^{k-1}/2)m\end{aligned}$$

Where the last line comes from the choice of $g = 5^{k-1}+2$, hence $\bar{c}_k = (5^{k-1}+2)(\bar{c}_{k-1}+1)$ and $\bar{c}'_k = (5^{k-1}+2)(\bar{c}'_{k-1} + 5^{k-1}/2)$. These constants do not have a clean closed-form solution but are bounded as $\bar{c}_k < (3/2)\cdot 5^{\binom{k}{2}}$ and $\bar{c}'_k < 4\cdot 5^{\binom{k}{2}}$. □

THEOREM 4.7. *Let ν_k be as defined in Theorem 4.6.*

1. $\mathrm{Ex}(\nu_k, n) < 5^{\left(\frac{k}{2}\right)^2 + k\log k}n$
2. $\mathrm{Ex}(\mathrm{dbl}(\nu_k), n) \le 5^{\left(\frac{k}{2}\right)^3 + 2k\log k}n$

PROOF. Proof of Part (1): Define ν'_k and ν''_k as follows:

$$\begin{aligned}\nu'_k &= a_1a_2\cdots a_k^3a_{k-1}\cdots a_1^3a_2\cdots a_k \\ \nu''_k &= a_2\cdots a_k^3a_{k-1}\cdots a_1^3a_2\cdots a_{k-1}\end{aligned}$$

We intend to show that $\mathrm{Ex}(\nu'_k, n) \le c_kn$ for a sequence (c_k) of constants to be determined. (This bound clearly extends to ν_k-free sequences.) The base case is satisfied with $c_2 = 28$ [16]. Suppose we are given a ν'_k-free sequence S with maximum length. Greedily partition $S = s_1\cdots s_m$ into ν''_k-free sequences (s_i) and let $S' = \Sigma(s_1)\cdots\Sigma(s_m)$ be obtained by discarding all but one occurrence of each symbol in each s_i. It follows that $m < 2n$ since each block contains the first or last occurrence of some symbol. Let $A' = A(S')$ be the canonical matrix for S'. If we remove from each row the first 1, second 1, and every other 1 thereafter, the resulting matrix is clearly $\hat{V}_k$-free. If it is shown that $\mathrm{Ex}(\nu''_k, n_0) \le c''_kn_0$ then $|S| = \sum_i |s_i| \le \sum_i \mathrm{Ex}(\nu''_k, \|s_i\|) \le c''_k|S'| \le c''_k(2\cdot \mathrm{Ex}(\hat{V}_k, n, 2n) + 2n)$.

We will now bound c''_k in terms of c_{k-2}. Let S'' be an arbitrary ν''_k-free sequence. Greedily partition $S'' = s''_1\cdots s''_m$ into sequences such that $|s''_i| = 2k\|s''_i\|$ for $i < m$ and $|s''_m| \le 2k\|s''_m\|$. Let $S''' = s'''_1s'''_2\cdots s'''_m$ be the sequence where $s'''_i \precsim s''_i$ retains one occurrence of each symbol in s''_i. We claim S''' is ν'_{k-2}-free and therefore has length at most $c_{k-2}n$. Consider an occurrence of ν'_{k-2} of the form $a_2\cdots a_{k-2}a_{k-1}^3a_{k-2}\cdots a_3a_2^3a_3\cdots a_{k-1}$ in S'''. Let $s'''_{i_0}, s'''_{i_1}, s'''_{i_2}$ be the blocks containing the three adjacent a_{k-1}s in ν'_{k-1} and let $s'''_{j_0}, s'''_{j_1}, s'''_{j_2}$ be the blocks containing the three adjacent a_2s. (Observe that i_1 and j_1 are always strictly less than m.) If we can identify some $a_k \notin \{a_2,\ldots,a_{k-1}\}$ that appears thrice in s''_{i_1} and some $a_1 \notin \{a_2,\ldots,a_k\}$ that appears thrice in s''_{j_1} then we have found an occurrence of ν''_k in S''. Iteratively remove all occurrences of $a_2,\ldots,a_{k-1}$ from s''_{i_1}. Removing a_2 leaves a $(k-1)$-sparse sequence with length $|s''_{i_1}| - \lceil |s''_{i-1}|/k > \frac{k-1}{k}|s''_{i_1}| - 1$, and, in the same manner, removing $a_3,\ldots,a_{k-1}$ leaves a 2-sparse sequence with length $\frac{2}{k}|s''_{i_1}| - (k-2) = 4\|s''_{i_1}\| - (k-2)$ on an alphabet of size $\|s''_{i_1}\| - (k-2)$. This sequence clearly contains three occurrences of some symbol that is by definition different from $a_2,\ldots,a_{k-1}$. We do the same procedure on s''_{j_1}, removing $a_2,\ldots,a_k$, which leaves a sequence with length $\frac{1}{k}|s''_{j_1}| - (k-1) = 2\|s''_{j-1}\| - (k-1)$ on an alphabet of $\|s''_{j-1}\| - (k-1)$ symbols. This sequence also clearly contains a symbol different from $a_2,\ldots,a_k$ appearing thrice. Putting everything together, $|S''| \le 2k|S'''| = 2kc_{k-2}n$, hence $c''_k = 2kc_{k-2}$. Finally, we determine c_k based on $|S| = \mathrm{Ex}(\nu'_k, n)$.

$$\begin{aligned}|S| &= \sum_i |s_i| \\ &\le \sum_i \mathrm{Ex}(\nu''_k, \|s_i\|) \\ &\le 2kc_{k-2}|S'| \\ &\le 2kc_{k-2}[2\cdot \mathrm{Ex}(\hat{V}_k, n, 2n) + 2n] \\ &\le 2kc_{k-2}\cdot 2[(5^{k-1}-1)n + (5^{k-1}/2-1)2n + n] \\ &< 8kc_{k-2}5^{k-1}n \\ &\le c_kn\end{aligned}$$

Where the last line holds for

$$c_k = 8^{k/2}k(k-2)(k-4)\cdots 5^{(k/2)^2} < 5^{(k/2)^2 + k\log k}.$$

The proof of Part (2) follows similar lines. Redefine ν'_k and ν''_k as follows:

$$\begin{aligned}\nu'_k &= a_1^2a_2^2\cdots a_{k-1}^2a_k^5a_{k-1}^2\cdots a_2^2a_1^5a_2^2\cdots a_{k-1}^2a_k^2 \\ \nu''_k &= a_2^2\cdots a_{k-1}^2a_k^5a_{k-1}^2\cdots a_2^2a_1^5a_2^2\cdots a_{k-1}^2.\end{aligned}$$

We will prove that $\mathrm{Ex}(\nu'_k, n) \le c_kn$ holds for a sequence (c_k) to be determined. The base case is satisfied with $c_2 = 100$ [16]. As in Part (1), we partition a ν'_k-free sequence $S = s_1\cdots s_m$ into ν''_k-free sequences (s_i). Let $S' = \Sigma(s_1)\cdots\Sigma(s_m)$ be obtained by discarding all but one occurrence of each symbol in each s_i. It follows that $m < 4n$ since each block contains the first, second, last, or second to last occurrence of some symbol. We will prove that $\mathrm{Ex}(\nu''_k, n_0) \le c''_kn_0$ for some sequence (c''_k). If we retain only every fourth occurrence of each symbol in S', the resulting matrix is clearly $\mathrm{dbl}(\nu_k)$-free, which implies that $|S| \le 4c''_k\cdot \mathrm{Ex}(\mathrm{dbl}(\nu_k), n, 4n)$.

We bound c''_k in terms of c_{k-2} in a manner similar to Part (1). Let S'' be an arbitrary ν''_k-free sequence. Greedily partition $S'' = s''_1\cdots s''_m$ into sequences such that $|s''_i| = 4k\|s''_i\|$ for $i < m$ and $|s''_m| \le 4k\|s''_m\|$. Let $S''' = s'''_1\cdots s'''_m$

be the sequence where $s_i''' \bar{\prec} s_i''$ retains one occurrence of each symbol in s_i''. From the same argument as in Part (1) it follows that S''' is ν_{k-2}'-free and therefore has length at most $c_{k-2}n$; hence $c_k'' = 4kc_{k-2}$.

$$\begin{aligned}\mathrm{Ex}(\nu_k', n) = |S| &= \sum_i |s_i| \\ &\le \sum_i \mathrm{Ex}(\nu_k'', \|s_i\|) \\ &\le 4kc_{k-2}|S'| \\ &\le 4kc_{k-2} \cdot 4 \cdot \mathrm{Ex}(\mathrm{dbl}(\nu_k), n, 4n) \\ &\le 4kc_{k-2} \cdot 4 \cdot (\tfrac{3}{2}5^{\binom{k}{2}}n + 4 \cdot 5^{\binom{k}{2}} \cdot 4n) \\ &\le 280kc_{k-2}5^{\binom{k}{2}}n \;\; \le \;\; c_k n\end{aligned}$$

The last line holds for $c_k = 280^{k/2}k(k-2)\cdots 5^{(k/2)^3} < 5^{(k/2)^3+2k\log k}$. □

5. DISCUSSION

Davenport-Schinzel sequences have numerous applications in discrete geometry and the analysis of algorithms and data structures, but *generalized* Davenport-Schinzel sequences have been applied in fewer situations; see [24, 31] and Theorem 3.2. What accounts for this discrepancy? We believe it to be due partly to the naturalness of alternating forbidden subsequences (and their geometric manifestations) and partly to the underpublicized work on general forbidden patterns [1, 15, 17, 21, 29, 27, 26]. We expect that broader awareness of the diversity of forbidden substructure theorems will lead to broader applications.

There are now a variety of tools to analyze general forbidden subsequences, the two most powerful being Klazar [15] and Nivasch's [21] bounds on $\mathrm{Perm}(r, s)$-avoiding sequences and the reduction from sequences to canonical 0-1 matrices [11, 27, 26]. The latter benefits from a growing library of results on forbidden 0-1 matrices [11, 30, 18, 12, 14, 9, 32, 26], the *grafting* operation introduced in Section 4 being just the latest example.

We will mention only a couple open problems in generalized Davenport-Schinzel sequences. All forbidden sequences known to be linear are N-shaped ([17], Theorems 4.6,4.7), double-permutations [26], the one outlier *abcbbccac* (Theorem 2.4), or formed from the previous three by embeddings [17]. Are there other large classes of linear forbidden sequences? Which are the linear sequences over 3- and 4-letter alphabets?

6. REFERENCES

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

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

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

[4] D. Bienstock and E. Györi. An extremal problem on sparse 0-1 matrices. *SIAM J. Discr. Math.*, 4(1):17–27, 1991.

[5] H. Davenport and A. Schinzel. A combinatorial problem connected with differential equations. *American J. Mathematics*, 87:684–694, 1965.

[6] A. Efrat and M. Sharir. A near-linear algorithm for the planar segment center problem. In *Proceedings 5th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 87–97, 1994.

[7] A. Efrat and M. Sharir. A near-linear algorithm for the planar segment-center problem. *Discrete & Computational Geometry*, 16(3):239–257, 1996.

[8] E. Ezra, B. Aronov, and M. Sharir. Improved bound for the union of fat triangles. In *Proceedings 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1778–1785, 2011.

[9] R. Fulek. Linear bound on extremal functions of some forbidden patterns in 0-1 matrices. *Discrete Mathematics*, 309:1736–1739, 2009.

[10] Z. Füredi. The maximum number of unit distances in a convex n-gon. *J. Combin. Theory Ser. A*, 55(2):316–320, 1990.

[11] Z. Füredi and P. Hajnal. Davenport-Schinzel theory of matrices. *Discrete Mathematics*, 103(3):233–251, 1992.

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

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

[14] B. Keszegh. On linear forbidden submatrices. *J. Combin. Theory Ser. A*, 116(1):232–241, 2009.

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

[16] M. Klazar. Extremal functions for sequences. *Discrete Mathematics*, 150(1–3):195–203, 1996.

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

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

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

[20] J. S. B. Mitchell. L_1 shortest paths among polygonal obstacles in the plane. *Algorithmica*, 8(1):55–88, 1992.

[21] G. Nivasch. Improved bounds and new techniques for Davenport-Schinzel sequences and their generalizations. *J. ACM*, 57(3), 2010.

[22] J. Pach and M. Sharir. On vertical visibility in arrangements of segments and the queue size in the Bentley-Ottmann line sweeping algorithm. *SIAM J. Comput.*, 20(3):460–470, 1991.

[23] J. Pach and G. Tardos. Forbidden paths and cycles in ordered graphs and matrices. *Israel J. Math.*, 155:359–380, 2006.

[24] S. Pettie. Splay trees, Davenport-Schinzel sequences, and the deque conjecture. In *Proceedings 19th ACM-SIAM Symposium on Discrete Algorithms*, pages 1115–1124, 2008.

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

[26] S. Pettie. Degrees of nonlinearity in forbidden 0-1 matrix problems. In submission, 2010.

[27] S. Pettie. Generalized Davenport-Schinzel sequences and their 0-1 matrix counterparts. In submission, 2010.

[28] S. Pettie. On nonlinear forbidden 0-1 matrices: A refutation of a Füredi-Hajnal conjecture. In *Proceedings 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 875–885, 2010.

[29] S. Pettie. Origins of nonlinearity in Davenport-Schinzel sequences. *SIAM J. Discrete Mathematics*, 25(1):211–233, 2011.

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

[31] P. Valtr. Graph drawings with no k pairwise crossing edges. In *Proceedings 5th Int'l Symposium on Graph Drawing*, pages 205–218, 1997.

[32] C. Weidert. Extremal problems in ordered graphs. *arXiv*, 0907.2479v1, 2009.

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

APPENDIX

The proof of Theorem 2.4 follows the same lines as that of Theorem 2.3 but requires some significant changes. Recall that Theorem 2.4 states that $\mathrm{Ex}(abcbbccac, n) < 198n$ and $\mathrm{Ex}(\tilde{U}_3, n, m) < 11n + 7m$.

Proof. (Theorem 2.4) Let S be an *abcbbccac*-free sequence with length $\mathrm{Ex}(abcbbccac, n)$. As in the proof of Theorem 2.3, we partition $S = s_1 \cdots s_m$ into *bcbbccac*-free subsequences, where $m \le n$. Let $S' = \Sigma(s_1)\cdots\Sigma(s_m)$ and let $A = A(S')$ be the $n \times m$ canonical matrix for S'. Since, by [16], $|s_i| \le \mathrm{Ex}(bcbbccac, \|s_i\|) < 11\|s_i\|$, we have $|S| \le 11|S'| = 11|A|$. The canonical matrix argument shows that A is $\tilde{U}_3$-free. We will show that $\mathrm{Ex}(\tilde{U}_3, n, m) < 11n + 7m$ and, therefore, that $\mathrm{Ex}(abcbbccac, n) \le 11 \cdot \mathrm{Ex}(\tilde{U}_3, n, n) < 198n$.

To show that $\mathrm{Ex}(\tilde{U}_3, n, m) = O(n+m)$ we require a couple nontrivial modifications to the proof of Theorem 2.3, beginning with the construction of $\mathscr{Q}$. For j from 3 to $n-1$, examine the 1s in row j from top to bottom. Let $(i, j) \in A$ be the current 1, let $\mathscr{Q}$ be the boxes constructed so far, and let i' be maximum such that $(i, j'), (i', j'') \in A$ where $i < i'$ and $j'' < j' < j$. If (i, j) is the first 1 in its column, or if it is already contained in a box in $\mathscr{Q}$, or if i' does not exist, then skip to the next 1. Otherwise include in $\mathscr{Q}$ the box $(i, \hat{\imath}) \times (j, \infty)$, where $\hat{\imath}$ is defined as:

$$\hat{\imath} = \min\left\{ i', \min\left\{ i_0 + 1 \;\middle|\; \begin{array}{l} (i_0, i_1) \times (j_0, \infty) = Q \in \mathscr{Q} \\ \text{and } i' \in [i_0 + 2, i_1) \end{array} \right\}\right\}$$

In other words, we force the rows spanned by $\mathscr{Q}$-boxes to be laminar. The new box would naturally span rows in the interval (i, i') but if $i' \in [i_0 + 2, i_1)$ then it would only partially intersect the rows spanned by Q. In this case we artificially make the lower boundary of the new box meet the upper boundary of Q. See Figure 4(A) for an illustration. As before we let $\mathscr{R} = \{R_1, R_2, \ldots\}$ where $R_k = Q_k \backslash \bigcup_{l>k} Q_l$. Clearly $\mathscr{R}$ consists of rectangular, non-overlapping boxes. We claim the matrix $A \backslash \mathscr{R}$ is J-free, where:

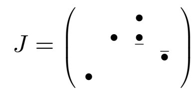

To see this, consider the moment the underlined 1 is examined during the construction of $\mathscr{Q}$. A box will be created that contains the overlined 1, which means that it cannot appear in $A \backslash \mathscr{R}$. After removing the first 1 in each row and each column of $A \backslash \mathscr{R}$ the resulting matrix is $U_1^{\circledcirc}$-free, which, by Theorem 1.1, implies $|A \backslash \mathscr{R}| < 3n + 3m$.

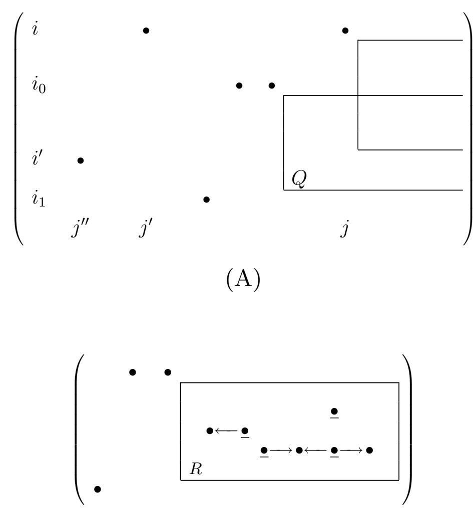

Figure 4: An instance of $U_1^{\circledcirc}$ in an $R \in \mathscr{R}$ (underlined) implies an instance of $\tilde{U}_3$ in A.

Obtain the matrix $\hat{A}$ by removing all 1s outside $\mathscr{R}$, removing the first three 1s and last 1 in each row, then removing every alternate 1 in each row. Thus, $|A| < 2|\hat{A}| + 7n + 3m$. An argument similar to that in the proof of Theorem 2.3 shows that no column or row has a non-zero intersection with two boxes in $\mathscr{R}$. Furthermore, every 1 in $\hat{A} \cap R$, for an $R \in \mathscr{R}$, is preceded by two 1s in its row in $A \cap R$. We claim each box in $\mathscr{R}$ is $U_1^{\circledcirc}$-free, which, if true, implies that $|A| < 2(\mathrm{Ex}(U_1^{\circledcirc}, n, m)) + 7n + 3m \le 11n + 7m$. Suppose that $U_1^{\circledcirc}$ appeared in $R \in \mathscr{R}$. Each 1 in $R \cap \hat{A}$ is preceded by a 1 in its row in $R \cap A$ and followed by a 1 in its row in A. Furthermore, two consecutive 1s in a row in $R \cap \hat{A}$ contain a 1 between them in A. These implied 1s and one 1 used in the formation of R give an instance of $\tilde{U}_3$. See Figure 4(B). □