

PATTERN AVOIDANCE IS NOT P-RECURSIVE

SCOTT GARRABRANT* AND IGOR PAK*

ABSTRACT. Let $\mathcal{F} \subset S_k$ be a finite set of permutations and let $C_n(\mathcal{F})$ denote the number of permutations $\sigma \in S_n$ avoiding the set of patterns \mathcal{F} . The *Noonan–Zeilberger conjecture* states that the sequence $\{C_n(\mathcal{F})\}$ is P-recursive. We disprove this conjecture.

The proof uses Computability Theory and builds on our earlier work [GP1]. We also give two applications of our approach to complexity of computing $\{C_n(\mathcal{F})\}$.

1. INTRODUCTION

Combinatorial sequences have been studied for centuries, with results ranging from minute properties of individual sequences to broad results on large classes of sequences. Even just listing the tools and ideas can be exhausting, which range from algebraic to bijective, to probabilistic and number theoretic [Rio]. The existing technology is so strong, it is rare for an open problem to remain unresolved for more than a few years, which makes the surviving conjectures all the more interesting and exciting.

The celebrated *Noonan–Zeilberger Conjecture* is one such open problem. It is in the area of *pattern avoidance* which has been very popular in the past few decades, see e.g. [Bóna, Kit]. The problem was first raised as a question by Gessel in 1990, see [Ges, §10]. In 1996, it was upgraded to a conjecture and further investigated by Noonan and Zeilberger [NZ], see also §10.11. There are now hundreds of papers in the area with positive results for special sets of patterns, all in support of the conjecture, and only two recent experimental clouds against it (see §10.4). Here we disprove the conjecture.

Let $\sigma \in S_n$ and $\omega \in S_k$. Permutation σ is said to *contain* the *pattern* ω if there is a subset $X \subseteq \{1, \dots, n\}$, $|X| = k$, such that $\sigma|_X$ has the same relative order as ω . Otherwise, σ is said to *avoid* ω . Fix a set of patterns $\mathcal{F} \subset S_k$. Denote by $C_n(\mathcal{F})$ the number of permutations $\sigma \in S_n$ avoiding the patterns $\omega \in \mathcal{F}$. The sequence $\{C_n(\mathcal{F})\}$ is the main object in the area, extensively analyzed from analytic, asymptotic and combinatorial points of view (see §10.1).

An integer sequence $\{a_n\}$ is called *polynomially recursive*, or *P-recursive*, if it satisfies a nontrivial linear recurrence relation of the form

$$q_0(n)a_n + q_1(n)a_{n-1} + \dots + q_k(n)a_{n-k} = 0,$$

for some $q_i(x) \in \mathbb{Z}[x]$, $0 \leq i \leq k$. The study of P-recursive sequences plays a major role in modern Enumerative and Asymptotic Combinatorics (see e.g. [FS, Odl, Sta]). They have *D-finite* (also called *holonomic*) generating series and various asymptotic properties (see §10.2).

Conjecture 1.1 (Noonan–Zeilberger). *Let $\mathcal{F} \subset S_k$ be a fixed set of patterns. Then the sequence $\{C_n(\mathcal{F})\}$ is P-recursive.*

The following is the main result of the paper.

Theorem 1.2. *The Noonan–Zeilberger conjecture is false. More precisely, there exists a set of patterns $\mathcal{F} \subset S_{80}$, such that the sequence $\{C_n(\mathcal{F})\}$ is not P-recursive.*

We should mention that in contrast with most literature in the area which studies small sets of patterns, our set \mathcal{F} is enormously large and we make no effort to decrease its size (cf. §10.6). To

Date: May 26, 2015.

*Department of Mathematics, UCLA, Los Angeles, CA, 90095. Email: {coscott, pak}@math.ucla.edu.

be precise, we construct two large sets $\mathcal{F}, \mathcal{F}' \subset S_{80}$ and show that at least one of them gives a counterexample to the conjecture. In fact, it is conceivable that a single permutation pattern $\omega = (1324)$ may be sufficient (see §10.4).

The proof of Theorem 1.2 is based on the following idea. Roughly, we show that every *two-stack automaton* M can be emulated by a finite set of permutation patterns. More precisely, we show that the number of accepted paths of M is equal to $C_n(\mathcal{F}) \bmod 2$, for a subset of integers n forming an arithmetic progression (see Main Lemma 3.2). This highly technical construction occupies much of the paper. The rest of the proof is based on our approach in [GP1], where we resolved *Kontsevich's problem* on the P-recursiveness of certain numbers of words in linear groups. The ability to emulate *any* two-stack automaton M in the weak sense described above is surprisingly powerful (see below).

Let us now present two applications of our results to the complexity of computing $\{C_n(\mathcal{F})\}$. Two sets of patterns \mathcal{F}_1 and \mathcal{F}_2 are called *Wilf-equivalent*, denoted $\mathcal{F}_1 \sim \mathcal{F}_2$, if $C_n(\mathcal{F}_1) = C_n(\mathcal{F}_2)$ for all n . In [V2], Vatter asked whether it is decidable when two patterns are Wilf-equivalent. Here we resolve a mod-2 version of this problem (cf. Section 8 and §10.10).

Theorem 1.3. *The problem whether $C_n(\mathcal{F}_1) = C_n(\mathcal{F}_2) \bmod 2$ for all $n \in \mathbb{N}$, is undecidable.*

For the second application, we say that an integer sequence $\{a_n\}$ has a *Wilfian formula*, if it can be computed in time polynomial in n . This notion was introduced by Wilf in a different context [Wilf]. Recently, both Klazar [Kla] and Vatter [V2] asked whether every sequence $\{C_n(\mathcal{F})\}$ has a Wilfian formula.

Theorem 1.4. *If $\text{EXP} \neq \oplus\text{EXP}$, then there exists a finite set of patterns \mathcal{F} , such that the sequence $\{C_n(\mathcal{F})\}$ cannot be computed in time polynomial in n .*

This also seems to give a good answer to Wilf's original question "Can one describe a reasonable and natural family of combinatorial enumeration problems for which there is provably no polynomial-in- n time formula or algorithm to compute $f(n)$?" [Wilf] (see also §10.9).

Here EXP is the complexity class of problems decidable in exponential time, and $\oplus\text{EXP}$ is the exponential time analogue of the $\oplus\text{P}$, so $\text{EXP} \subseteq \oplus\text{EXP} \subseteq \text{EXPSPACE}$. We refer to Section 9 for the definitions, details and references. It is believed that $\text{EXP} \neq \oplus\text{EXP}$ (see §9.1), so the answer to the question by Klazar and Vatter is likely negative (cf. §10.8).

The rest of the paper is structured as follows. We begin with an explicit construction of a two-stack automata with a non-P-recursive number of accepted paths (Section 2). In Section 3, we reduce the proof of Theorem 1.2 to the Main Lemma 3.2 on embedding two-stack automata into pattern avoidance problems. The proof of the Main Lemma spans the next four sections. We first present the construction in Section 4. In Section 5, we prove the lemma modulo a number of technical results. We illustrate the construction in a lengthy example in Section 6, and prove the technical results in Section 7. We proceed to prove theorems 1.3 and 1.4 in sections 8 and 9, respectively. We conclude with final remarks and open problems in Section 10.

2. TWO-STACK AUTOMATA

In this section we construct an automaton with a non-P-recursive number of accepted paths. The construction is technical, but elementary. Although it is more natural if the reader is familiar with basic Automata Theory (see e.g. [HMU, Sip]), the construction is completely self-contained and is given in the language of elementary Graph Theory. The proof, however, is not self-contained and follows a similar proof in [GP1].

To be precise, we give an explicit construction of a graph, where the vertices have certain variables as weights. We then count the number a_n of paths of length n between two fixed vertices, where only certain weight sequences are allowed (we call these *balanced paths*). The non-P-recursiveness of $a_n \bmod 2$ is explained below.

2.1. The motivation. It is relatively easy to present a construction of an automaton which produces a non-P-recursive sequence $\{a_n\}$ of balanced paths. Our goal in this section is stronger – the sequence $\{a_n\}$ we get is not equal to *any* P-recursive sequence modulo 2. Somewhat informally, we call such automaton *non-P-recursive*. The advantages afforded by the modulo 2 property are technical and will become clear later in this paper.

Our main tool for building a non-P-recursive two-stack automaton is the following result.

Theorem 2.1 ([GP1, Lemma 2]). *Let $\{a_n\}$ be a P-recursive integer sequence. Consider an infinite binary word $\bar{\alpha} = (\alpha_1\alpha_2\dots)$, defined by $\alpha_n := a_n \bmod 2$. Then, there exists a finite binary word which is not a subword of $\bar{\alpha}$.*

What follows is a construction of a two-stack automaton such that the corresponding binary sequence $\bar{\alpha}$ contains every finite binary subword by design. There are many such automata, in fact. We give a complete description of this one as we need both its notation and additional properties of the construction later on.

2.2. The setup. Let Γ be a finite directed graph with vertices v_1, \dots, v_m . Let X denote the set of labels of the form x_i and let X^{-1} denote the set of labels of the form x_i^{-1} , where i is any integer. Define Y and Y^{-1} similarly. Label each vertex of Γ with an element of $X \cup X^{-1} \cup Y \cup Y^{-1} \cup \{\varepsilon\}$.

Let $\rho(v)$ denote the label on vertex v . We say that $w_1 \sim w_2$ if $w_1, w_2 \in X \cup X^{-1}$ or if $w_1, w_2 \in Y \cup Y^{-1}$. If Γ has an edge from v_i to v_j , we say that $v_i \rightarrow v_j$.

Contrary to standard notation, we refer to a path $\gamma = \gamma_1 \dots \gamma_n$, where each γ_i is a vertex, not an edge, and we say that such a path is of length n , even though it only has $n - 1$ edges.

We further require that $\rho(v_1) = \rho(v_2) = \varepsilon$, and that there is no edge $v_i \rightarrow v_j$, with $\rho(v_i) \sim \rho(v_j)$. A graph Γ satisfying all of the above conditions is called a *two-stack automaton*.

As we traverse a path γ , we keep track of two words $w_X \in X^*$ and $w_Y \in Y^*$, which start out empty. Whenever we enter a vertex with label x_i , we append x_i to the end of w_X . When we enter a vertex with label x_i^{-1} , we remove x_i from the end of w_X . We modify w_Y similarly when entering vertices with label y_i or y_i^{-1} . When we enter a vertex with label ε , we do nothing. A path is called *balanced* if every step of this process is well defined and both w_X and w_Y are empty at the end of the path. Let $G(\Gamma, n)$ denote the number of balanced paths in Γ from v_1 to v_2 of length n .

Define an involution $\pi_\gamma \in S_n$ as follows. If the above process writes an instance of a label to w_X or w_Y at some time t_i , and removes the same instance of that label for the first time at time t_j , then $\pi(t_i) = t_j$ and $\pi(t_j) = t_i$. If the process does not write or remove anything at a time step t_k , then $\pi(t_k) = t_k$. For example, if $\rho(\gamma_1)\rho(\gamma_2)\dots\rho(\gamma_9) = \varepsilon x_1 y_1 x_1 y_1^{-1} x_1^{-1} \varepsilon x_1^{-1} \varepsilon$, then $\pi_\gamma = (2\ 8)(3\ 5)(4\ 6)$. This gives the following alternate characterization of balanced paths.

Proposition 2.2. *A path γ is balanced if and only if there exists an involution $\pi_\gamma \in S_n$ such that:*

- (1) $\rho(\gamma_i) = \varepsilon$ for all $\pi_\gamma(i) = i$,
- (2) $\rho(\gamma_i) \in X \cup Y$ and $\rho(\gamma_{\pi_\gamma(i)}) = \rho(\gamma_i)^{-1}$, for all $\pi_\gamma(i) > i$, and
- (3) There are no i and j with $\rho(\gamma_i) \sim \rho(\gamma_j)$ such that $i < j < \pi_\gamma(i) < \pi_\gamma(j)$.

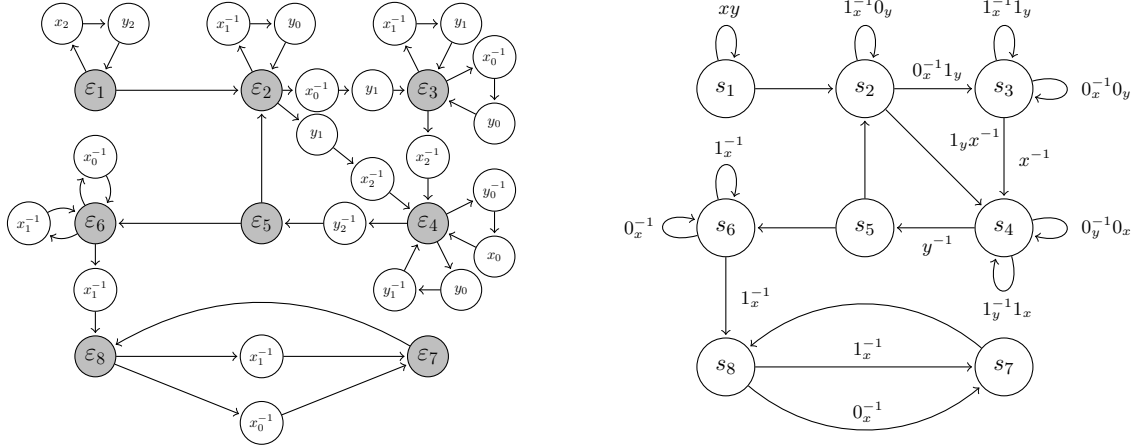
Further, this involution π_γ is uniquely defined for each balanced γ .

The proof is straightforward.

2.3. Non-P-recursive automaton. We are now ready to present a construction of such automaton Γ_1 , which is given in Figure 1. The construction is based on a smaller automaton Γ_2 we introduced in [GP1].

Lemma 2.3. *There exists a two-stack automaton Γ_1 , such that $\alpha_n := G(\Gamma_1, n) \in \{0, 1\}$ for all n , and such that the word $\bar{\alpha} = (\alpha_1\alpha_2\dots)$ is an infinite binary word which contains every finite binary word as a subword.*

Proof. We give an explicit automaton Γ_1 in Figure 1. This automaton is formed by modifying the automaton Γ_2 , given in [GP1]. Here we use $\varepsilon_1, \dots, \varepsilon_8$ to denote the same trivial label ε ; we make this

FIGURE 1. The automata Γ_1 (left), and Γ_2 (right).

distinction only for the purpose of illustration. The vertex v_1 is the shaded vertex labelled ε_1 , and the vertex v_2 is the shaded vertex labelled ε_8 . The vertex labelled ε_i in Γ_1 corresponds to the vertex labelled s_i in Γ_2 .

The primary difference between Γ_1 and Γ_2 is that Γ_1 has labels on vertices while Γ_2 has labels on edges. The labels were also changed by replacing 0_x , 1_x , and x with x_0 , x_1 , and x_2 respectively, and similarly for y . Since the lengths of the paths change slightly, we get a slightly different formula, but the analysis is similar.

In counting paths we follow the proof of Lemma 3 in [GP1]. The valid paths through Γ_2 have

$$\mu = j + 6k + 2 \sum_{i=1}^k \lfloor \log_2 i \rfloor \quad \text{edges,}$$

for some positive integers j and k such that the j -th binary digit of k is a 1. Every path through Γ_1 will similarly have $(\mu + 1)$ vertices with label ε .

Such paths will also have a total of $4k$ vertices labelled x_2 , x_2^{-1} , y_2 or y_2^{-1} , since k copies each of x_2 and y_2 are written and removed in the computation. Similarly, every binary integer from 1 to k is written and removed from both tapes, so the vertices with the remaining 8 labels are used

$$\nu = 4k + 4 \sum_{i=1}^k \lfloor \log_2 i \rfloor \quad \text{times in total.}$$

In summary, we have $G(\Gamma_1, n) = 1$ for all $n = (\mu + 1) + \nu + 4k$, where j and k are positive integers such that the j -th binary digit of k is a 1, and $G(\Gamma_1, n) = 0$ otherwise. The word $\bar{\alpha} = (\alpha_1 \alpha_2 \dots)$ then contains the positive integer k written out in binary starting at location

$$n = 2 + 14k + 6 \sum_{i=1}^k \lfloor \log_2 i \rfloor,$$

and will therefore contain every finite binary word as a subword.

Finally, note that in the notion of “valid path” from [GP1], it was possible for some instance of x^{-1} to cancel with a later instance of x . For our purposes, this difference is irrelevant since the words defined by paths in Γ_2 do not have such cancellations. \square

3. MAIN LEMMA AND THE PROOF OF THEOREM 1.2

In this section we first change our setting from pattern avoidance to slightly more general but equivalent notion of *partial pattern avoidance*. We state the Main Lemma 3.2 and show that it implies Theorem 1.2.

3.1. Partial patterns. A 0-1 matrix is called a *partial pattern* if every row and column contains at most one 1. Clearly, every permutation pattern is also a partial pattern. We say that a permutation matrix M *contains* a partial pattern L , if L can be obtained from M by deleting some rows and columns; we say that M *avoids* L otherwise. Given a set \mathcal{F} of partial patterns, let $\mathcal{C}_n(\mathcal{F})$ denote the set of $n \times n$ matrices M which avoid all partial patterns in \mathcal{F} . By analogy with the usual permutation patterns, let $C_n(\mathcal{F}) = |\mathcal{C}_n(\mathcal{F})|$.

Proposition 3.1. *Let \mathcal{F}_1 be a finite set of partial patterns. Then there exists a finite set of the usual permutation patterns \mathcal{F}_2 , such that $C_n(\mathcal{F}_1) = C_n(\mathcal{F}_2)$ for all $n \in \mathbb{N}$.*

Proof. First, let us prove the result for a single partial pattern. Let L be a partial pattern of size $p \times q$, and let $k = p + q$. Denote by $\mathcal{P}_n(L)$ be the set of $n \times n$ permutation matrices containing L . Let us show by induction that for all $n \geq k$, every permutation matrix $M \in \mathcal{P}_n(L)$ contains a matrix in $\mathcal{P}_k(L)$. Indeed, the claim is trivially true for $n = k$. For larger n , observe that every $n \times n$ permutation matrix M which contains L must also contain some i -th row and j -th column, such that $M_{p,q} = 1$, and neither i -th row nor j -th column intersect L . This follows from the fact that otherwise $\text{rank}(M) \leq i + j < n$. Deleting these row and column gives a smaller permutation matrix which contains L , proving the induction claim.

We conclude that $\mathcal{F}_2(L) := \cup_{\ell \leq k} \mathcal{P}_\ell(L)$ is the desired set of matrices for $\mathcal{F} = \{L\}$. In full generality, take $\mathcal{F}_2 = \cup_{L \in \mathcal{F}} \mathcal{F}_2(L)$. The details are straightforward. \square

It therefore suffices to disprove the Noonan–Zeilberger Conjecture 1.1 for partial patterns.

Lemma 3.2 (Main Lemma). *Let Γ be a two-stack automaton. Then there exist sets \mathcal{F} and \mathcal{F}' of partial patterns, and some integers $c, d \geq 1$, such that*

$$C_{cn+d}(\mathcal{F}) - C_{cn+d}(\mathcal{F}') = G(\Gamma, n) \pmod{2}, \quad \text{for all } n \in \mathbb{N}.$$

The proof of the Main Lemma is given in Section 5.

3.2. Proof of Theorem 1.2. By Lemma 2.3, there exists a two-stack automaton Γ_1 , such that the infinite binary word $\bar{\alpha} = (\alpha_1\alpha_2\dots)$ given by $\alpha_n = G(\Gamma_1, n)$, contains every finite binary word as a subword. By Lemma 3.2, there exist integers c and d and two sets \mathcal{F} and \mathcal{F}' of partial patterns such that $C_{cn+d}(\mathcal{F}) - C_{cn+d}(\mathcal{F}') = G(\Gamma_1, n) \pmod{2}$, for all n .

If Conjecture 1.1 is true, then both $\{C_n(\mathcal{F})\}$ and $\{C_n(\mathcal{F}')\}$ are P-recursive sequences. Since P-recursive sequences are closed under taking the differences and subsequences with indices in arithmetic progressions (see e.g. [Sta, §6.4]), this means that the sequence $\{a_n\}$, defined as $a_n = \{C_{cn+d}(\mathcal{F}) - C_{cn+d}(\mathcal{F}')\}$, is also P-recursive. On the other hand, from above, we have $\alpha_n = a_n \pmod{2}$. This gives a contradiction with Theorem 2.1.

The second part of the theorem requires a quantitative form of the Main Lemma and is given as Corollary 4.4. \square

4. THE CONSTRUCTION OF AN AUTOMATON IN THE MAIN LEMMA

4.1. Notation. The construction of sets of matrices $\mathcal{F}, \mathcal{F}'$ has two layers and is quite involved, so we try to simplify it by choosing a clear notation. We use \mathcal{A}_g to denote a certain subset of $g \times g$ matrices, which we call an *alphabet* and use as building blocks. We use English capitals with various decorations, notably $A, A', B, B', E, L, P, Q, R, S, T_k$ and Z_p , to denote 0-1 matrices of size at most $g \times g$. We use script capital letters $\mathcal{F}_i, \mathcal{F}'_i, \mathcal{W}_i, \mathcal{W}'_i$, to denote the sets of larger matrices (partial patterns) which form sets $\mathcal{F}, \mathcal{F}'$. Each is of size at most $8g \times 8g$, and some of the matrices are denoted W_i and W'_i .

On a bigger scale, we use $M = M(*, *)$ to denote large block matrices, with individual blocks $M^{i,j}$ being either zero or matrices in the alphabet \mathcal{A}_g . For the proof of Theorem 1.2 we take $g = 10$, but for theorems 1.3 and 1.4 we need larger g . When writing matrices, we use a dot (\cdot) within a matrix to represent a single 0 entry, and a circle (\circ) to represent a $g \times g$ submatrix of zeros.

4.2. The alphabet. A permutation matrix is called *simple* if it contains no permutation matrix as a proper submatrix consisting of consecutive rows and columns, other than the trivial 1×1 permutation matrix.

Define an alphabet \mathcal{A}_g of all $g \times g$ simple permutation matrices which contain the following matrix as a submatrix:

$$L = \begin{pmatrix} \dots \dots 1 \dots \\ \dots \dots \dots 1 \dots \\ \dots \dots \dots \dots 1 \dots \\ \dots \dots 1 \dots \dots \\ 1 \dots \dots \dots \\ \dots 1 \dots \dots \dots \\ \dots \dots 1 \dots \dots \dots \end{pmatrix}.$$

Proposition 4.1. *We have: $|\mathcal{A}_g| \rightarrow g!/e^2$ as $g \rightarrow \infty$.*

Proof. It was shown in [AAK] that the probability that a random $g \times g$ permutation matrix M is simple tends to $1/e^2$ as $g \rightarrow \infty$ (see also [OEIS, A111111]). On the other hand, the probability that M avoids L tends to 0 as $g \rightarrow \infty$. Thus, the probability that $M \in \mathcal{A}_g$ tends to $1/e^2$, as desired. \square

By the proposition, we can fix an integer g large enough that $|\mathcal{A}_g| > 5 + m + r$, where m is the number of vertices in Γ and r is the number of distinct labels in $X \cup Y$ on vertices of Γ .

We build our forbidden partial patterns out of elements of \mathcal{A}_g as follows. Choose five special matrices $P, Q, B, B', E \in \mathcal{A}_g$, as well as two classes of matrices, $T_1, \dots, T_m \in \mathcal{A}_g$, and $Z_1, \dots, Z_r \in \mathcal{A}_g$. Here the matrices T_1, \dots, T_m represent the m vertices in Γ . Let T_i denote the matrix corresponding to v_i . The matrices Z_1, \dots, Z_r represent the r labels in $X \cup Y$. Let $s(Z_p)$ denote the label which corresponds to Z_p . Let us emphasize that these choices are arbitrary as the only important properties of these $(5 + m + r)$ matrices is that they are all in \mathcal{A}_g and distinct.

4.3. Forbidden matrices. Let \mathcal{F}_1 denote the set of all $g \times (g + 1)$ or $(g + 1) \times g$ partial patterns formed by taking a matrix A in \mathcal{A}_g , and inserting a row or column of all zeros somewhere in the middle of A .

Let \mathcal{F}_2 denote the set of all $(2g + 1) \times (5g + 1)$ or $(5g + 1) \times (2g + 1)$ partial patterns whose bottom left $g \times g$ consecutive submatrix is a B or B' and whose top right $g \times g$ consecutive submatrix is T_j for some j .

Let \mathcal{F}_3 denote the four element set consisting of the $(2g + 1) \times g$ and $g \times (2g + 1)$ partial pattern formed by inserting $g + 1$ rows of zeros below Q , inserting $g + 1$ rows of zeros above P , inserting $g + 1$ columns of zeros to the right of Q , or inserting $g + 1$ columns of zeros to the left of P .

Let $\mathcal{F}_4 = \mathcal{W}_1 \cup \mathcal{W}_2 \cup \mathcal{W}_3 \cup \mathcal{W}_4 \cup \mathcal{W}_5$, where \mathcal{W}_i are defined as follows. Let $\mathcal{W}_1, \mathcal{W}_2$ and \mathcal{W}_3 denote the sets of matrices of the form W_1, W_2 and W_3 , respectively:

$$W_1 = \begin{pmatrix} \circ \circ T_i \circ \circ \circ \circ \circ \\ \circ \circ \circ \circ \circ T_j \circ \circ \\ L \circ \circ \circ \circ \circ \circ \circ \\ \circ \circ \circ \circ \circ \circ \circ Z_p \\ \circ \circ \circ \circ \circ \circ T_k \circ \\ \circ B' \circ \circ \circ \circ \circ \circ \\ \circ \circ \circ \circ R \circ \circ \circ \\ \circ \circ \circ Z_p \circ \circ \circ \circ \end{pmatrix}, W_2 = \begin{pmatrix} \circ \circ \circ \circ Z_p \circ \circ \circ \\ \circ \circ \circ T_i \circ \circ \circ \circ \\ \circ \circ \circ \circ \circ \circ T_j \circ \\ \circ L \circ \circ \circ \circ \circ \circ \\ Z_p \circ \circ \circ \circ \circ \circ \circ \\ \circ \circ \circ \circ \circ \circ T_k \\ \circ \circ B' \circ \circ \circ \circ \circ \\ \circ \circ \circ \circ \circ R \circ \circ \circ \end{pmatrix}, W_3 = \begin{pmatrix} \circ \circ T_i \circ \circ \circ \circ \\ \circ \circ \circ \circ \circ T_j \circ \\ L \circ \circ \circ \circ \circ \circ \\ \circ \circ \circ E \circ \circ \circ \\ \circ \circ \circ \circ \circ \circ T_k \\ \circ B' \circ \circ \circ \circ \circ \\ \circ \circ \circ \circ R \circ \circ \circ \end{pmatrix}.$$

In all three cases, we require $L, R \in \{B, B'\}$ and $v_i \rightarrow v_j \rightarrow v_k$. In W_1 , we require $\rho(v_j) = s(Z_p)$. In W_2 , we require $\rho(v_j) = (s(Z_p))^{-1}$. In W_3 , we require $\rho(v_j) = \varepsilon$.

Similarly, let \mathcal{W}_4 and \mathcal{W}_5 denote the set of all matrices of the form W_4 and W_5 respectively:

$$W_4 = \begin{pmatrix} \circ & \circ & \circ & \circ & T_1 & \circ \\ \circ & P & \circ & \circ & \circ & \circ \\ \circ & \circ & E & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & T_k \\ B' & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & R & \circ & \circ \end{pmatrix}, \quad W_5 = \begin{pmatrix} \circ & \circ & T_i & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & T_2 \\ L & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & E & \circ & \circ \\ \circ & \circ & \circ & \circ & Q & \circ \\ \circ & B' & \circ & \circ & \circ & \circ \end{pmatrix}.$$

Here, in W_4 , we require $R \in \{B, B'\}$ and $v_1 \rightarrow v_k$, and in W_5 , we require $L \in \{B, B'\}$ and $v_i \rightarrow v_2$. Finally, let \mathcal{F}_5 denote the set of all patterns of the form

$$\begin{pmatrix} \circ & Z_p & \circ \\ \circ & \circ & Z_q \\ B & \circ & \circ \end{pmatrix}, \quad \begin{pmatrix} \circ & Z_p & \circ \\ \circ & \circ & Z_q \\ B' & \circ & \circ \end{pmatrix} \quad \text{or} \quad \begin{pmatrix} \circ & \circ & T_j \\ Z_p & \circ & \circ \\ \circ & Z_q & \circ \end{pmatrix}, \quad \text{where } s(Z_p) \sim s(Z_q).$$

Lemma 4.2 (Explicit Construction). *Given a two-stack automaton Γ , let*

$$\mathcal{F} := \mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3 \cup \mathcal{F}_4 \cup \mathcal{F}_5 \quad \text{and} \quad \mathcal{F}' := \mathcal{F} \cup \{B, B'\},$$

where $\mathcal{F}_1, \dots, \mathcal{F}_5, B, B'$ are defined as above. Then, for all n , we have:

$$C_m(\mathcal{F}) - C_m(\mathcal{F}') = G(\Gamma, n) \pmod{2}, \quad \text{where } m = (3n + 2)g.$$

The Main Lemma 3.2 follows immediately from this result.

4.4. Counting Partial Patterns. We will now analyze the above construction in the specific case of Γ_1 .

Theorem 4.3. *There exists a set \mathcal{F} of at most 6854 partial patterns of size at most 80×80 such that $\{C_n(\mathcal{F})\}$ is not P-recursive.*

Converting these partial patterns into the usual permutation patterns would require many more patterns. However all the patterns avoided would still be size at most 80×80 .

Corollary 4.4. *There exists a set \mathcal{F} of 80×80 permutation matrices such that the sequence $\{C_n(\mathcal{F})\}$ is not P-recursive. In particular, $|\mathcal{F}| < 80! < 10^{119}$.*

Proof of Theorem 4.3. Observe that Γ_1 has 31 vertices and uses 6 labels in $X \cup Y$. Therefore we need $5 + 31 + 6 = 42$ matrices in \mathcal{A}_g . Let $g = 10$. Consider the following simple 9×9 pattern

$$L' = \begin{pmatrix} \dots & \dots & 1 & \dots \\ \cdot & 1 & \dots & \dots \\ \dots & \dots & \dots & 1 \dots \\ \dots & \dots & \dots & \dots & 1 \\ \dots & \dots & 1 & \dots \\ 1 & \dots & \dots & \dots \\ \dots & 1 & \dots & \dots \\ \dots & \dots & \dots & 1 \dots \\ \dots & \dots & 1 & \dots \end{pmatrix}.$$

Note that there are 60 ways to insert 1 into L' to form a simple 10×10 pattern. Indeed, the 1 may be inserted anywhere other than the 4 corners or the 36 locations that would form a 2×2 consecutive submatrix. All 60 of these 10×10 are distinct and in \mathcal{A}_g .

For \mathcal{F}_1 , we actually only need to include the 42 matrices in \mathcal{A}_{10} which are actually used, so $|\mathcal{F}_1| = 42 \cdot 9 \cdot 2 = 756$. Similarly, for \mathcal{F}_2 , there are 2 choices for the bottom left 10×10 consecutive submatrix and 31 choices for the top right 10×10 consecutive submatrix. There are 41 entries in the middle and at most one of them can be a 1, which can be satisfied in 42 ways. Therefore, $|\mathcal{F}_2| = 2 \cdot 31 \cdot 2 \cdot 42 = 5208$. Clearly, $|\mathcal{F}_3| = 4$.

Let us show that $|\mathcal{F}_4| = 292$. Indeed, a matrix in $\mathcal{W}_1 \cup \mathcal{W}_2 \cup \mathcal{W}_3$ is defined by the path $v_i \rightarrow v_j \rightarrow v_k$ and the choices for L and R . There are 71 paths in Γ of length 3, so we have $|\mathcal{W}_1 \cup \mathcal{W}_2 \cup \mathcal{W}_3| = 71 \cdot 4 = 284$. A matrix in \mathcal{W}_4 is defined by the vertex v_k and the choices for R , so we have $|\mathcal{W}_4| = 4$. Similarly, a matrix in \mathcal{W}_5 is defined by the vertex v_i and the choices for L , so $|\mathcal{W}_5| = 4$.

Finally, for \mathcal{F}_5 , there are 6 choices for Z_p , 3 choices for Z_q and $31 + 2$ choices for the B, B' or T_j . Therefore, $|\mathcal{F}_5| = 6 \cdot 3 \cdot 33 = 594$. In total, \mathcal{F} consists of $|\mathcal{F}| = 756 + 5208 + 4 + 292 + 594 = 6854$

partial patterns of dimensions at most 80×80 . The set \mathcal{F}' has two extra matrices, but can be made smaller than \mathcal{F} since avoiding B' makes all matrices $W_i \in \mathcal{F}_4$ redundant. \square

5. PROOF OF THE EXPLICIT CONSTRUCTION LEMMA 4.2

In this section we give a proof of Lemma 4.2 by reducing it to three technical lemmas which are proved in Section 7. Briefly, since $\mathcal{F} \subset \mathcal{F}'$, we have $\mathcal{C}_n(\mathcal{F}') \subseteq \mathcal{C}_n(\mathcal{F})$ for all n . Denote $\mathcal{D}_n = \mathcal{C}_n(\mathcal{F}') \setminus \mathcal{C}_n(\mathcal{F})$. We construct an explicit involution ϕ on \mathcal{D}_n and analyze the set of fixed points \mathcal{D}'_n . We show that the set \mathcal{D}'_n has a very rigid structure emulating the working of a given two-stack automaton Γ .

5.1. Preliminaries. The key idea of an involution ϕ defined below is a switch $B \leftrightarrow B'$ between submatrices B and B' , in such a way that the fixed points \mathcal{D}'_n of ϕ avoid B' . The remaining copies of B create a general diagonal structure of the matrices in \mathcal{D}'_n , and enforce the location of all other submatrices from the alphabet. We invite the reader to consult the example in the next section to have a visual understanding of our approach.

We also need a convenient notion of a *marked submatrix*. Such marked submatrix will always be a B , and is located at a specific position in forbidden matrices W_i . This is best illustrated in the matrix formulas below, where marked submatrix B is boxed.

Let $\mathcal{F}'_4 = \mathcal{W}'_1 \cup \mathcal{W}'_2 \cup \mathcal{W}'_3 \cup \mathcal{W}'_4 \cup \mathcal{W}'_5$, where $\mathcal{W}'_i \subseteq \mathcal{W}_i$ are defined to have no submatrices B' (so $L = R = B$ in the notation above), and where the elements of \mathcal{W}'_4 have a unique marked submatrix B . Precisely, let $\mathcal{W}'_1, \mathcal{W}'_2$ and \mathcal{W}'_3 denote the set of matrices of the form W'_1, W'_2 and W'_3 , respectively:

$$W'_1 = \begin{pmatrix} \circ & \circ & T_i & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & T_j & \circ & \circ \\ B & \circ & \circ & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & \circ & Z_p & \circ \\ \circ & \circ & \circ & \circ & \circ & \circ & \circ & T_k \\ \circ & \boxed{B} & \circ & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & B & \circ & \circ & \circ \\ \circ & \circ & \circ & Z_p & \circ & \circ & \circ & \circ \end{pmatrix}, W'_2 = \begin{pmatrix} \circ & \circ & \circ & \circ & Z_p & \circ & \circ & \circ \\ \circ & \circ & \circ & T_i & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & \circ & T_j & \circ \\ \circ & B & \circ & \circ & \circ & \circ & \circ & \circ \\ Z_p & \circ & \circ & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & \circ & \circ & T_k \\ \circ & \circ & \boxed{B} & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & B & \circ & \circ \end{pmatrix}, W'_3 = \begin{pmatrix} \circ & \circ & T_i & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & T_j & \circ \\ B & \circ & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & E & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & \circ & T_k \\ \circ & \circ & \boxed{B} & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & B & \circ & \circ \end{pmatrix}.$$

Similarly, let \mathcal{W}'_4 , and \mathcal{W}'_5 denote the set of matrices of the form W'_4 and W'_5 , respectively:

$$W'_4 = \begin{pmatrix} \circ & \circ & \circ & \circ & T_1 & \circ \\ \circ & P & \circ & \circ & \circ & \circ \\ \circ & \circ & E & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & T_k \\ \circ & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & B & \circ \end{pmatrix}, W'_5 = \begin{pmatrix} \circ & \circ & T_i & \circ & \circ & \circ \\ \circ & \circ & \circ & \circ & \circ & T_2 \\ B & \circ & \circ & \circ & \circ & \circ \\ \circ & \circ & \circ & E & \circ & \circ \\ \circ & \circ & \circ & \circ & Q & \circ \\ \circ & \circ & \boxed{B} & \circ & \circ & \circ \end{pmatrix}.$$

Of course, all these W'_i satisfy the same conditions as W_i in the previous section.

5.2. Construction of the involution ϕ . From this point on, let $m = (3n + 2)g$. Given a $m \times m$ permutation matrix M , let $M^{i,j}$ refer to the $g \times g$ submatrix in rows $g(i - 1) + 1$ through gi , and columns $g(j - 1) + 1$ through gj .

First, observe that $\mathcal{D}_n := \mathcal{C}_m(\mathcal{F}') \setminus \mathcal{C}_m(\mathcal{F})$ is the set of all $m \times m$ matrices avoiding \mathcal{F} with at least one submatrix B or B' . Since \mathcal{D}_n avoids \mathcal{F}_1 , every submatrix B or B' in a matrix in \mathcal{D}_n must be a consecutive $g \times g$ block. A submatrix B or B' in a matrix M of \mathcal{D}_n is called *blocked* if replacing it with B' or B , respectively, would result in a matrix not in \mathcal{D}_n .

Lemma 5.1. *Consider the map ϕ on \mathcal{D}_n which takes the leftmost unblocked submatrix B or B' and replaces it with B' or B respectively. The map ϕ is an involution on \mathcal{D}_n . Furthermore, the fixed points of ϕ are the $m \times m$ matrices M such that:*

- (1) matrix M avoids \mathcal{F} ,
- (2) matrix M avoids B' ,
- (3) matrix B is a submatrix of M , and
- (4) every submatrix B inside M is a marked submatrix of a matrix in \mathcal{F}'_4 .

Let $\mathcal{D}'_n = \text{Fix}(\phi)$ denote the set of all fixed points of ϕ . Since ϕ is an involution, we conclude that $|\mathcal{D}_n| = |\mathcal{D}'_n| \pmod 2$, so it suffices to show that $|\mathcal{D}'_n| = G(\Gamma, n) \pmod 2$.

5.3. The structure of \mathcal{D}'_n . Let γ be a path from v_1 to v_2 which is not necessarily balanced, and let $\pi \in S_n$. Denote by $M := M(\gamma, \pi)$ the $m \times m$ permutation matrix given by:

- (1) $M^{2,2} = P$,
- (2) $M^{3n+1,3n+1} = Q$,
- (3) $M^{3i+2,3i-2} = B$ for all i ,
- (4) $M^{3i-2,3i+2} = T_j$ for all i , where $\gamma_i = v_j$,
- (5) $M^{3i,3j} = E$ whenever $\rho(\gamma_i) = \varepsilon$ and $\pi(i) = j$,
- (6) $M^{3i,3j} = Z_p$ whenever $\rho(\gamma_i) = s(Z_p)$ and $\pi(i) = j$,
- (7) $M^{3i,3j} = Z_p$ whenever $\rho(\gamma_i) = (s(Z_p))^{-1}$ and $\pi(i) = j$,
- (8) $M^{i,j} = 0$ is a zero matrix otherwise.

Lemma 5.2. *Every matrix in \mathcal{D}'_n is of the form $M(\gamma, \pi)$, where γ is a path of length n in Γ , and $\pi \in S_n$.*

Note, however, not every matrix $M(\gamma, \pi)$ is in \mathcal{D}'_n . The following lemma gives a complete characterization. Recall that given a balanced path γ , there is a unique permutation π_γ associated with γ given in Proposition 2.2.

Lemma 5.3. *Let γ be a path in Γ of length n , and let $\pi \in S_n$. Then, $M(\gamma, \pi) \in \mathcal{D}'_n$ if and only if γ is balanced and $\pi = \pi_\gamma$.*

Lemmas 5.1, 5.2 and 5.3 easily imply the Main Lemma.

5.4. Proof of Lemma 4.2. We have $C_m(\mathcal{F}) - C_m(\mathcal{F}') = |\mathcal{D}_n|$ by definition. Lemma 5.1 shows that ϕ is an involution, so $|\mathcal{D}_n| = |\text{Fix}(\phi)| = |\mathcal{D}'_n| \pmod 2$. Combining lemmas 5.2 and 5.3, we get $|\mathcal{D}'_n| = G(\Gamma, n)$. Thus, $C_{cn+d}(\mathcal{F}) - C_{cn+d}(\mathcal{F}') = G(\Gamma, n) \pmod 2$, as desired. \square

6. EXAMPLE

Let us illustrate the construction in a simple case. Consider a two-stack automaton Γ_3 given in Figure 2. Note that Γ_3 has a unique balanced path $\gamma = v_1 v_3 v_5 v_3 v_6 v_4 v_2 v_4 v_2$.




FIGURE 2. The two-stack automaton Γ_3 .

Let us show that the following matrix $M = M(\gamma, \pi_\gamma)$ is unique in the set of fixed points \mathcal{D}'_9 . Here in the matrix, we have $s(Z_1) = x_1$ and $s(Z_2) = y_1$ (see below).

As in the definition of $M(\gamma, \pi_\gamma)$, observe that M has a diagonal of B entries below the main diagonal, and a diagonal of T_i entries above the main diagonal. The T_i entries give the vertices of the path γ , in order. Observe that M also has a P in the top left and Q in the bottom right, as in the definition of $M(\gamma, \pi_\gamma)$.

We have here the involution $\pi_\gamma = (2\ 8)(3\ 5)(4\ 6)$, and the locations of the E , Z_1 and Z_2 matrices form the permutation matrix for π_γ . Each matrix E corresponds to a time when the path γ visited a vertex labelled ε . Similarly, each Z_p above the diagonal corresponds to a pair of times when a given instance of a label was written and removed from one of the stacks.

(4) every submatrix B in M is a marked submatrix in some matrix \mathcal{F}'_4 .

Observe that (1) and (3) imply that M avoids \mathcal{F} but not $\mathcal{F}' = \mathcal{F} \cup \{B, B'\}$. Therefore, we have $M \in \mathcal{D}_n$. Further, (2) and (4) imply that M is a fixed point of ϕ , since M has no B' , and replacing any B with a B' will create a matrix in \mathcal{F}_4 .

Next, we show that every matrix in \mathcal{D}_n which violates the above criteria is fixed by ϕ^2 but not by ϕ . Observe that a matrix $M \in \mathcal{D}_n$ which satisfies (2) also satisfies (1) and (3). Consider now a matrix $M \in \mathcal{D}_n$ which violates either (2) or (4). Clearly, we have $\phi(M) \in \mathcal{D}_n$. It suffices to show that $\phi(M) \neq M$ and that $\phi^2(M) = M$.

Let N be a matrix in \mathcal{D}_n . Denote by A a submatrix B or B' in N . Let N' be the matrix formed by replacing A with B' or B , respectively. Similarly, let A' denote this submatrix B or B' in N' .

Assume that N' did not avoid \mathcal{F}_1 . Then there would exist a submatrix $S \in \mathcal{A}_g$ of N' which is not a consecutive $g \times g$ block. Note that S must intersect A' . Consider the submatrix $S \cap A'$. This must be a permutation matrix, since S and A' are permutation matrices. Furthermore, $S \cap A'$ must be a consecutive submatrix of S , since A' is a consecutive submatrix of N' . Since S is simple, we have that $S \cap A'$ is a single 1 entry. This entry might be a 0 in N , but it does not matter because we could replace it with another 1 entry from A , to form a non-consecutive submatrix S in N . Thus, N does not avoid \mathcal{F}_1 , contradicting the assumption that $N \in \mathcal{D}_n$.

Since N' avoids \mathcal{F}_1 , any instance of a matrix from $\mathcal{F}_2, \mathcal{F}_3$, or \mathcal{F}_5 intersecting A' in N' must contain A' or intersect A' in a single row or column. In either case, replacing A with A' gives another matrix in $\mathcal{F}_2, \mathcal{F}_3$, or \mathcal{F}_5 , contradicting the fact that N' avoids $\mathcal{F}_2, \mathcal{F}_3$ and \mathcal{F}_5 .

Therefore, if the submatrix A is blocked it must be because A' is contained in a matrix in \mathcal{F}_4 . Since \mathcal{F}_4 is closed under replacing any B with a B' , it must be that $A = B$ and $A' = B'$. Any matrix which violates (2) therefore contains a instance of B' , which is necessarily unblocked, and so is not fixed by ϕ . For matrices in \mathcal{D}_n which satisfy (2), containing an unblocked instance of B is equivalent to violating (4). Therefore, any matrix which violates (4) is not fixed by ϕ , so $\phi(M) \neq M$.

Let $\eta(N)$ denote the leftmost unblocked instance of B or B' in N , and let $A = \eta(N)$. Observe that A' is also clearly unblocked in N' . Assume that $A' \neq \eta(N')$. There must be another unblocked instance S of B or B' , which is further to the left than A' . However, S would also be unblocked in N , contradicting the assumption that $A = \eta(N)$. Therefore, $A' = \eta(N')$, which implies that $\phi(N') = N$. We conclude that $\phi(\phi(M)) = M$ for all M .

In summary, every matrix M satisfying (1) through (4) is fixed by ϕ , and every matrix in \mathcal{D}_n not satisfying (1) through (4) is fixed by ϕ^2 but not ϕ . Therefore, map ϕ is an involution whose fixed points are exactly the matrices satisfying (1) through (4). \square

7.2. Proof of Lemma 5.2. Let M be a matrix in \mathcal{D}'_n . Lemma 5.1 shows that

- (1) matrix M avoids \mathcal{F} ,
- (2) matrix M avoids B' ,
- (3) matrix B is a submatrix of M , and
- (4) every instance of B in M is the marked submatrix in some matrix in \mathcal{F}'_4 .

Observe that every matrix in \mathcal{F}'_4 has a B or a P above the marked submatrix \boxed{B} . Therefore, for every instance of B in M , there must be another B or a P somewhere above it. Since there is at least one B in M , there must be at least one P in M . This P must have at least g rows above it and g columns to the left, since it is contained in some matrix in \mathcal{W}'_4 . Note that matrix P cannot have any more rows above it or columns to the left, since M avoids \mathcal{F}_3 . Therefore, $M^{2,2}$ is the unique instance of P in M . Similar analysis shows that $M^{3n+1,3n+1}$ is the unique instance of Q in M .

Similarly, every matrix in \mathcal{F}'_4 has a T_j at least $4g$ rows above and $4g$ columns to the right of the \boxed{B} . Since M avoids \mathcal{F}_2 , there can be no more rows or columns inserted in between this B and T_j . Therefore, for every instance of B in M , there must be some T_j exactly $4g$ rows above and $4g$ columns to the right. In particular, this means that if $M^{i,j} = B$, then $M^{i-4,j+4} = T_j$ for some j .

The T_1 above the P must be $4g$ rows above and $4g$ columns to the right of the B to the left of the P . This is only possible if this B is in location $M^{5,1}$, and the T_1 is in location $M^{1,5}$.

If $M^{i,j} = B$ and $i \neq n$, then there must another instance of B exactly $3g$ columns to the right, and some T_k exactly g rows above, since every matrix in $\mathcal{W}'_1, \mathcal{W}'_2, \mathcal{W}'_3$ or \mathcal{W}'_4 has this pattern. We know that these submatrices cannot be further away because they are squeezed between the B and the T_j . The new T_k must be $4g$ rows below and $4g$ columns to the right of the new instance of B . This can only be achieved by letting $M^{i+3,j+3} = B$ and $M^{i-1,j+7} = T_k$.

Thus, the instances of B in a matrix M in D'_n are exactly the matrices $M^{3i+2,3i-2}$, for $1 \leq i \leq n$. Similarly, we have $M^{3i-2,3i+2} = T_j$ for some j . Let γ_i denote the vertex in Γ corresponding to the matrix $M^{3i-2,3i+2}$. The restrictions in \mathcal{F}'_4 about adjacent v_i ensure that $\gamma = \gamma_1 \dots \gamma_n$ is a path from v_1 to v_2 .

When we specify this path γ , we uniquely define all of the submatrices $M^{i,j}$ except for those of the form $M^{3i,3j}$. Further, the restrictions from \mathcal{F}'_4 tell us that each $M^{3i,3j}$ is either a matrix of zeros or equal to E or Z_p . If the matrix at location $M^{3i,3j}$ is nonzero, then it is uniquely determined by the restrictions from \mathcal{F}'_4 on the rows of M .

Let π be the permutation such that $\pi(i)$ is the unique j , such that $M^{3i,3j}$ is E or Z_p . Let us prove that $M = M(\gamma, \pi)$. We already showed that $M^{2,2} = P$, $M^{3n+1,3n+1} = Q$, and $M^{3i+2,3i-2} = B$ for all i . We know that $M^{3i-2,3i+2} = T_j$ where $\gamma_i = v_j$, by the definition of γ .

By definition, submatrix $M^{3i,3j}$ is non-zero if and only if $\pi(i) = j$. The restrictions on E and Z_p from \mathcal{F}'_4 imply that $M^{3i,3j} = E$ for $\rho(\gamma_i) = \varepsilon$, $M^{3i,3j} = Z_p$ for $\rho(\gamma_i) = s(Z_p)$, and $M^{3i,3j} = Z_p$ for $\rho(\gamma_i) = (s(Z_p))^{-1}$. We have a total of m entries 1 's, so all remaining $M^{i,j}$ are zeros matrices. This implies that $M = M(\gamma, \pi)$. \square

7.3. Proof of the only if part of Lemma 5.3. First, we show that if $M(\gamma, \pi) \in \mathcal{D}'_n$, then γ is balanced and $\pi = \pi_\gamma$. By Proposition 2.2, it suffices to show that the permutation π satisfies the following four properties:

- (1) $\rho(\gamma_i) = \varepsilon$ for all $\pi(i) = i$,
- (2) $\rho(\gamma_i) \in X \cup Y$ and $\rho(\gamma_{\pi(i)}) = \rho(\gamma_i)^{-1}$, for all $\pi(i) > i$,
- (3) there are no i and j with $\rho(\gamma_i) \sim \rho(\gamma_j)$ such that $i < j < \pi(i) < \pi(j)$, and
- (4) permutation π is an involution.

Proof of (1): Consider the case where $\pi(i) = i$. The submatrix $M^{3i,3i}$ is non-zero then. In this case, the B at $M^{3i+2,3j-2}$ must be the marked submatrix \boxed{B} in some instance of a matrix in $\mathcal{W}'_3, \mathcal{W}'_4$ or \mathcal{W}'_5 . Thus, $M^{3i,3i} = E$ and $M^{3i-2,3j-2} = T_j$, where $\rho(\gamma_i) = \rho(v_j) = \varepsilon$, so π satisfies (1).

Proof of (2): Consider the case where $\pi(i) > i$, so $M^{3i,3\pi(i)}$ is non-zero. In this case, the submatrix B at $M^{3i+2,3i-2}$ must be \boxed{B} in some instance of a matrix in \mathcal{W}'_1 . Thus, $M^{3i,3\pi(i)} = Z_p$, and $M^{3i-2,3i+2} = T_j$, where $\rho(\gamma_i) = \rho(v_j) = s(Z_p)$. Similarly, the submatrix B at $M^{3\pi(i)+2,3\pi(i)-2}$ must be a marked submatrix in some matrix in \mathcal{W}'_2 . Thus, $M^{3\pi(i)-2,3\pi(i)+2} = T_k$, where $\rho(\gamma_{\pi(i)}) = \rho(v_k) = (s(Z_p))^{-1}$. We conclude that $\rho(\gamma_i) \in X \cup Y$ and $\rho(\gamma_{\pi(i)}) = \rho(\gamma_i)^{-1}$, so π satisfies (2). Similar analysis shows that if $\pi(i) < i$, then $\rho(\gamma_{\pi(i)}) \in X \cup Y$ and $\rho(\gamma_i) = \rho(\gamma_{\pi(i)})^{-1}$, so π^{-1} satisfies (2) as well.

Proof of (3): Assume that there exist i and j with $\rho(\gamma_i) \sim \rho(\gamma_j)$ such that $i < j < \pi(i) < \pi(j)$. Let $s(Z_p) = \rho(\gamma_i)$ and $s(Z_q) = \rho(\gamma_j)$. We then have $M^{3i,3\pi(i)} = Z_p$ and $M^{3j,3\pi(j)} = Z_q$. Observe that $M^{3j+2,3j-2} = B$. Since $3i < 3j < 3j+2$ and $3j-2 < 3\pi(i) < 3\pi(j)$, these three matrices together form a pattern which is in \mathcal{F}_5 , a contradiction. This implies that π satisfies (3). Similar analysis shows that there are no i and j with $\rho(\gamma_i) \sim \rho(\gamma_j)$, such that $\pi(i) < \pi(j) < i < j$, so π^{-1} satisfies (3) as well.

Proof of (4): Note that the conditions on π only mention values i for which $\pi(i) \geq i$. Therefore, even if π were not an involution, the permutation π would have to agree with π_γ at all i such that $\pi(i) \geq i$, by the uniqueness of the involution π_γ . However, a similar analysis shows that π^{-1} must also satisfy conditions (1), (2) and (3), so π^{-1} also agrees with π_γ at all i such that $\pi^{-1}(i) \geq i$. Combining these two observations, we conclude that $\pi = \pi^{-1} = \pi_\gamma$. \square

7.4. Proof of the *if* part of Lemma 5.3. It remains to prove that $M(\gamma, \pi_\gamma) \in \mathcal{D}'_n$. Clearly, matrix B is a submatrix of $M(\gamma, \xi\pi_\gamma)$. In addition, matrix $M(\gamma, \pi_\gamma)$ has exactly n submatrices B , and each of them is the marked submatrix \boxed{B} in some matrix in \mathcal{F}'_4 .

Recall that there is a unique P and Q in $M(\gamma, \pi_\gamma)$, and they are sufficiently close to the edge to ensure that $M(\gamma, \pi_\gamma)$ avoids \mathcal{F}_2 . Similarly, we know the locations of all submatrices B and T_i , and they ensure that $M(\gamma, \pi_\gamma)$ avoids \mathcal{F}_3 . Clearly, matrix $M(\gamma, \pi_\gamma)$ avoids B' , and therefore also avoids \mathcal{F}_4 . Also, if $M(\gamma, \pi_\gamma)$ did not avoid \mathcal{F}_5 , this would contradict the fact that there are no i and j with $\rho(\gamma_i) \sim \rho(\gamma_j)$ such that $i < j < \pi_\gamma(i) < \pi_\gamma(j)$.

Finally, we need to show that $M(\gamma, \pi_\gamma)$ avoids \mathcal{F}_1 . Assume that $M(\gamma, \pi_\gamma)$ does not avoid \mathcal{F}_1 , and that A is a $g \times g$ submatrix of $M(\gamma, \pi_\gamma)$, such that $A \in \mathcal{A}_g$ and A is not a consecutive $g \times g$ block in $M(\gamma, \pi_\gamma)$. Since A is simple, every 1 in A comes from a different block $M^{i,j}$. By definition of \mathcal{A}_g , matrix L is a submatrix of A (see §4.1). Assume the 1 in the center of L is above or on the main diagonal of M . It must therefore be in an E , T_j , or Z_p . The top three 1's in L must be in matrices of the form Z_q , since there is not enough room for any of them to be in a T_k .

Since the relation “ \sim ” partitions the set of Z_p 's into two equivalence classes, two of these 1's must be in some Z_p and Z_q with $Z_p \sim Z_q$. However, there is a submatrix B that is below and to the left of both Z_p and Z_q . Together they form a matrix in \mathcal{F}_5 , contradicting the fact that $M(\gamma, \pi_\gamma)$ was shown above to avoid \mathcal{F}_5 . Similar analysis gives a contradiction if the 1 in the center of L is below the main diagonal.

We conclude that $M(\gamma, \pi_\gamma)$ satisfies all four conditions from Lemma 5.1. This implies $M(\gamma, \pi_\gamma) \in \mathcal{D}'_n$, as desired. \square

8. DECIDIBILITY

8.1. Simulating Turing Machines. It is well known and easy to see that two-stack automata can simulate (nondeterministic) *Turing machines*. This simulation works by using the two-stacks to represent the tape of the Turing machine (see e.g. [HMU]). To recap, let one stack contain everything written on the tape to the left of the head, while the other stack contains everything written to the right. Moving the head of the Turing machine left or right corresponds to removing a symbol from one stack and writing a (possibly different) symbol on the other stack. This simulation is direct and can be done in polynomial time. We refer to [HMU, Sip] for the definitions and details.

8.2. Proof of Theorem 1.3. Let M be an arbitrary deterministic Turing machine, with no input. Consider the two-stack automaton which simulates M . If M does not halt, $G(\Gamma, n) = 0$ for all n . If M eventually halts, then Γ will have a single balanced path, and $G(\Gamma, n) = 1$ for some n .

Construct the \mathcal{F} and \mathcal{F}' associated with this Γ as in Lemma 3.2. By analogy, we then have $C_n(\mathcal{F}) = C_n(\mathcal{F}') \bmod 2$ for all $n = d \bmod c$ if and only if M does not halt. Recall also that the *halting problem* is undecidable (see e.g. [Sip]).

Therefore, by the argument above, the construction in the proof of Lemma 3.2 can be emulated by a Turing machine, and the construction satisfies $C_n(\mathcal{F}) = C_n(\mathcal{F}') \bmod 2$ for all $n \neq d \bmod c$. Thus, given M , we can construct finite sets \mathcal{F} and \mathcal{F}' permutation patterns, such that $C_n(\mathcal{F}) = C_n(\mathcal{F}')$ for all $n \geq 1$ if and only if M does not halt. Therefore, it is undecidable whether $C_n(\mathcal{F}) = C_n(\mathcal{F}')$ for all $n \geq 1$. \square

8.3. Implications. Theorem 1.3 has rather interesting implications for our understanding of pattern avoidance (cf. §10.10). For example, a standard argument (see e.g. [Poon, p. 2]), gives the following surprising result:

Corollary 8.1. *There exist two finite sets of patterns \mathcal{F} and \mathcal{F}' , such that the problem of whether $C_n(\mathcal{F}) = C_n(\mathcal{F}') \bmod 2$ for all $n \in \mathbb{N}$, is independent of ZFC.*

We would be curious to see an explicit bound on the size of such sets of patterns, but it is likely to be quite large. Here is another immediate application of the theorem:

Corollary 8.2. *For all k large enough, there exist two finite sets of patterns \mathcal{F} and \mathcal{F}' , such that the smallest n for which $C_n(\mathcal{F}) \neq C_n(\mathcal{F}') \pmod{2}$, satisfies $n > A(n)$.*

Here $A(n)$ is the Ackermann function (see e.g. [AF]), but any other computable function can be used in its place.

8.4. Speculations. These results give us confidence in the following open problems, further extending Theorem 1.3.

Conjecture 8.3 (Parity problem). *The problem of whether $C_n(\mathcal{F}) = 0 \pmod{2}$ for all $n \in \mathbb{N}$, is undecidable.*

We believe that our tools will prove useful to establish this conjecture, likely with a great deal more effort. This goes beyond the scope of the paper.

Conjecture 8.4 (Wilf-equivalence problem). *The problem of whether $C_n(\mathcal{F}_1) = C_n(\mathcal{F}_2)$ for all $n \in \mathbb{N}$ is undecidable.*

This conjecture is partly motivated by the recently found unusual examples of Wilf-equivalence [BP]. It is more speculative than Conjecture 8.3, since in our construction $C_n(\mathcal{F}) - C_n(\mathcal{F}')$ grows rather rapidly. It is thus conceivable that in contrast with Corollary 8.2, the smallest n for which $C_n(\mathcal{F}) \neq C_n(\mathcal{F}')$ is constructible, perhaps even quite small in the size of the patterns. This would imply that the mod-2 problem is computationally harder than the Wilf-equivalence problem, a surprising conclusion.

In a different direction, it would be also very interesting to resolve the parity and the Wilf-equivalence problems when only one permutation is avoided. In this case, we believe that both are likely to be decidable.

Conjecture 8.5. *For forbidden sets with a single permutation $|\mathcal{F}| = |\mathcal{F}'| = 1$, both the parity and the Wilf-equivalence problems are decidable.*

In order to clarify and contrast the conjectures, let us note that non-P-recursiveness is quite possible and rather likely for permutation class avoiding one or two permutations (see §10.4). However, the undecidability is a much stronger condition and we do not believe one permutation has enough room to embed all logical axioms. This “small size is hard to achieve” phenomenon is somewhat similar to aperiodicity and undecidability for plane tilings (see §10.7).

9. WILFIAN FORMULAS

9.1. Complexity preliminaries. Recall the $\oplus P$ complexity class which is a parity version of the class of counting problem $\#P$, see e.g. [Pap]. For example, $\oplus 2SAT$ and the parity of the number of Hamiltonian cycles are $\oplus P$ -complete, see [Val]. It is strongly believed that $\oplus P \neq P$. In fact, if $P = \oplus P$, then $PH = BPP$ by Toda’s theorem. In particular, this implies $NP = BPP$, contradicting widely held beliefs that $P = BPP$ and $P \neq NP$.

Recall the exponential complexity classes EXP , $NEXP$ and its counting counterpart $\#EXP$. Let us now formally define an exponential time version of $\oplus P$. A language D is in $\oplus EXP$ if there exists a nondeterministic Turing machine M such that M has an *odd number* of accepting paths on input n if and only if $n \in D$, and every path in M halts in time $O(2^{p(\ell)})$ where p is a polynomial, and ℓ is the length of the input in bits. For example, computing parity of the number of non-isomorphic Hamiltonian graphs on n vertices is in $\oplus EXP$, see [OEIS, A003216]. Similarly, for a fixed \mathcal{F} , computing $C_n(\mathcal{F}) \pmod{2}$ is in $\oplus EXP$ by definition (cf. 10.9).

Now, it is believed that $EXP \neq \oplus EXP$ for reasons similar to $P \neq \oplus P$. First, note that $P = \oplus P$ implies $EXP = \oplus EXP$, see e.g. [BBF]. In the opposite direction, it follows from the proof of the main result in [HIS], that $EXP = \oplus EXP$ implies that there are no *polynomially sparse languages* in $\oplus P \setminus P$ (see [Pap] for definitions). This is an unlikely conclusion given how far apart P and $\oplus P$ are.¹ For more discussion of $\oplus EXP$ in relation to other complexity classes, see [BBF] (see also §10.8).

¹This argument and further evidence against $EXP = \oplus EXP$ was communicated to us by Joshua Grochow on *CS Theory Stack Exchange*: <http://tinyurl.com/m79o4dv>.

9.2. Proof of Theorem 1.4. Let D be a language in $\oplus\text{EXP}$. There exists a nondeterministic Turing machine M and a polynomial p such that M has an odd number of accepted paths on input n if and only if $n \in D$, and every path in M halts in time at most $2^{p(\log n)}$ on input n .

Let us form M' modifying M , by artificially extending every accepted path in such a way that it instead accepts at exactly time $\lfloor 2^{p'(\log n)} + n \rfloor$ for some $p' \geq p$. We further assume that p' is increasing, so $\lfloor 2^{p'(\log n)} + n \rfloor$ is injective on \mathbb{N} . Therefore, all accepted paths of M' of length $\lfloor 2^{p'(\log n)} + n \rfloor$ must come from input n .

Let us form M'' by modifying M' in such a way that instead of taking an input n , it nondeterministically chooses a natural number n . The number of accepted paths of length $\lfloor 2^{p'(\log n)} + n \rfloor$ will still be odd if and only if $n \in D$. We can now build a two-stack automaton Γ which simulates M'' so that $G(\Gamma, \lfloor 2^{p''(\log n)} + n \rfloor)$ is equal to the number of accepted paths in M of length $\lfloor 2^{p'(\log n)} + n \rfloor$.

Finally, we construct the sets of patterns \mathcal{F} and \mathcal{F}' associated with this Γ , following the construction in Lemma 3.2. We have then:

$$G(\Gamma, \lfloor 2^{p''(\log n)} + n \rfloor) = C_m(\mathcal{F}) - C_m(\mathcal{F}') \pmod{2},$$

where $m = c\lfloor 2^{p''(\log n)} + n \rfloor + d$. This implies that $n \in D$ if and only if $C_m(\mathcal{F}) - C_m(\mathcal{F}')$ is odd.

Assume now that \mathcal{F} and \mathcal{F}' have a Wilfian formula. Then we can compute $C_m(\mathcal{F}) - C_m(\mathcal{F}')$ and thus decide if $n \in D$ in time polynomial in m . Clearly, time polynomial in m is time exponential in $(\log n)$. This implies that $\text{EXP} = \oplus\text{EXP}$, a contradiction. \square

10. FINAL REMARKS AND OPEN PROBLEMS

10.1. Despite a large literature on pattern avoidance, relatively little is known about general sets of patterns. Notably, the *Stanley–Wilf conjecture* proved by Marcus and Tardos shows that $C_n(\mathcal{F})$ are at most exponential [MT], improving on an earlier near-exponential bound by Alon and Friedgut [AF]. Most recently, Fox showed that for a fixed k and a random permutation $\omega \in S_k$, we have $C_n(\mathcal{F}) = \exp(k^{\Theta(1)}n)$ [Fox] (see also [V3, §2.5]).

10.2. For an integer P-recursive sequence $\{a_n\}$, the generating series $\mathcal{A}(t) = \sum_{n=0}^{\infty} a_n t^n$ is *D-finite* (*holonomic*), i.e. satisfies a linear ODE (see e.g. [FS, Sta]). In the case $a_n = e^{O(n)}$, the series $\mathcal{A}(t)$ is also a *G-function*, an important notion in Analytic Number Theory (see e.g. [Gar]).

10.3. The most celebrated example of Wilf-equivalence is $(123) \sim (213)$. It follows from results of MacMahon (1915) and Knuth (1973), that $C_n(123) = C_n(213) = \frac{1}{n+1} \binom{2n}{n}$, the *n-th Catalan number*, see e.g. [Kit, Sta]. Many other Wilf-equivalent classes are known now, see e.g. [Kit].

10.4. There are three Wilf-equivalence classes of patterns in S_4 . Two of them are known to give P-recursive sequences, but whether $\{C_n(1324)\}$ is P-recursive remains a long-standing open problem in the area (see e.g. [Ste, V3]). The problem proved so challenging, Zeilberger himself seems to have abandoned Conjecture 1.1 because of it, see [EV].

We should mention that there seems to be some recent strong experimental evidence *against* sequence $\{C_n(1324)\}$ being P-recursive. Numerical analysis of the values for $n \leq 36$ given in [CG] suggest the asymptotic behavior

$$C_n(1324) \sim A \cdot \lambda^n \cdot \mu \sqrt{n} \cdot n^\alpha,$$

where $A \approx 8$, $\lambda \approx 11.6$, $\mu \approx 0.04$, and $\alpha \approx -1.1$. If this asymptotics holds, this would imply non-P-recursive by Theorem 7 in [GP1], which forbids $\mu \sqrt{n}$ terms.

For larger sets, the leading candidates for non-P-recursive are the sequences $\{C_n(4231, 4123)\}$ and $\{C_n(4123, 4231, 4312)\}$ which were recently computed for $n \leq 1000$ and 5000, respectively [A+].

10.5. There is a bit of a controversy over a question whether the integer $N := C_{1000}(1324)$ is humanly or supernaturally computable, with Steingrímsson and Zeilberger occupying the opposite sides of the argument, see [EV, Ste]. From a formal Theoretical CS point of view, this is really an argument over the existence of a Wilfian formula to compute $\{C_n(1324)\}$ (cf. [MR]). This raises a curious question whether computing $\{C_n(1324) \bmod 2\}$ is in EXP, which is perhaps easier to resolve. Leaving aside philosophical and theological issues, the proof of Theorem 1.4 suggests a surprising possibility that even the *parity* of N might be hard to compute.

Note also that $C_n(1324) = V_n(1324) \bmod 2$, where $V_n(\omega)$ is the number of *involutions* $\sigma \in S_n$ avoiding ω . The sequence $\{V_n(1324)\}$ is still difficult to analyze (see [BHPV]). It is known to have a much smaller growth, so perhaps easier to compute.

10.6. The bound $|\mathcal{F}| < 10^{119}$ in Corollary 4.4 can easily be improved down to a bit under 30,000 by refining the proof of Theorem 4.3 and using small technical tricks. While it is possible that a single permutation suffices (see above), it is unlikely that the tools from this paper can be used.

10.7. It is worth comparing the non-P-recursiveness of pattern avoidance with the history of *aperiodic tilings* in the plane. Originally conjectured by Wang to be impossible [Wang], such tilings were first constructed by Berger in 1966 by an undecidability argument [Ber]. Note that Berger’s construction used about 20,000 tiles. In 1971, by a different undecidability argument, Robinson was able to reduce this number to six [Rob]. Most recently, Ollinger established the current record of five polyomino tiles [Oll]. Without undecidability, Penrose famously used only two tiles to construct aperiodic tilings (see e.g. [GS]). Whether there exists one tile which forces aperiodicity remains open. This is the so-called *einstein problem* (cf. [ST] and [Yang, §7.5]). See also [Boas, GI, P1] for the complexity of several other related general tiling problems.

10.8. Note that we do not claim that computing the parity of $C_n(\mathcal{F})$ is \oplus EXP-complete. This is because our argument would have \mathcal{F} and \mathcal{F}' as an input, which is tangential to the Klazar–Vatter question in the introduction. Still, it would be interesting to prove the following complexity variation on Theorem 1.4:

Conjecture 10.1. *There exists a finite set of patterns \mathcal{F} , such that computing $\{C_n(\mathcal{F})\}$ is $\#$ EXP-complete, and computing $\{C_n(\mathcal{F}) \bmod 2\}$ is \oplus EXP-complete.*

To give an example of a \oplus EXP-complete problem, one can take a parity version of a known $\#$ EXP-complete problem in cases where the reduction is parsimonious. For example, the parity of $\frac{1}{6}$ times the number of 3-colorings in *succinct graphs* [GW] is a natural candidate (see [Bar] for a parsimonious reduction).² Perhaps the closest in spirit to the conjecture are the NEXP-complete tiling problems considered in [GI], where the input is a single integer n , rather than a set of tiles (cf. [Boas, P1]).

10.9. *Wilfian formulas* were introduced by Wilf in 1982, see [Wilf]. They have been christened and received some attention only very recently [V2, Zei].³ We believe this is the first paper which examines them from the computational complexity point of view. Of course, pattern avoidance itself is a special case of the *pattern matching problem*, with a large algorithmic and complexity literature. See e.g. [GM] for the recent breakthrough and [Kit, §8.2.1] for further references.

We should mention that all P-recursive sequences trivially have Wilfian formulas, but the opposite is false. For example, the sequence of primes is not P-recursive [FGS], but the n -th prime is clearly computable in exponential time. Same with the number $p(n)$ of partitions of n . The asymptotics $p(n) \sim C \cdot n^{-1} \cdot \mu \sqrt{n}$ implies that it is not P-recursive (see above), but Euler’s recurrence (see e.g. [P2]), shows it is computable in time polynomial in n .

A quick warning to the reader. Just because the sequence $\{a_n\}$ grows exponentially and has no known $o(a_n)$ time algorithm to compute it, does not mean that $\{a_n \bmod 2\}$ is also hard. Typical

²This approach was suggested to us by Greg Kuperberg on *CS Theory Stack Exchange*: <http://tinyurl.com/m79o4dv>.

³In the original paper [Wilf], Wilf uses two notions of a *formula* which led to some confusion in the definitions. In particular, Klazar in [Kla] uses a more restrictive notion which is equivalent to the one we use for exponentially growing sequences.

examples mentioned in the literature are the number of self-avoiding paths [Zei] (see also [OEIS, A001411]), and the number of unlabeled graphs on n vertices [Wilf] (see also [OEIS, A000088]). In both cases, the sequences are even for n large enough (the first is obvious, the second is proved in [CR]).

10.10. In [EV], Zeilberger asks to characterize pattern avoiding permutation classes (of single permutations) with g.f.'s rational, algebraic, or D-finite. We believe a complete characterization of the latter class is impossible.

Open Problem 10.2. *The problem of whether $\{C_n(\mathcal{F})\}$ is P-recursive is undecidable.*

Let us note that our results offer only a weak evidence in favor of the open problem, since Theorem 2.1 gives only a necessary condition on P-recursiveness, and (cf. [GP1]). It would be interesting to see if any of Zeilberger's problems are decidable. A rare decidability result in this direction is given in [BRV] (see also a discussion in [V3, §3]).

10.11. The original form of the Noonan–Zeilberger Conjecture makes a stronger claim: for every $\mathcal{F} = \{\omega_1, \dots, \omega_\ell\}$ and an integer vector $\mathbf{r} = (r_1, \dots, r_\ell) \in \mathbb{N}^\ell$, the sequence $\{C_n(\mathcal{F}, \mathbf{r})\}$ is P-recursive, where $C_n(\mathcal{F}, \mathbf{r})$ is the number of permutations $\sigma \in S_n$ which contain exactly r_i copies of pattern ω_i , for all $1 \leq i \leq \ell$. It was shown by Atkinson in 1999, that this stronger claim is equivalent to the $r_1 = \dots = r_\ell = 0$ case, which is our Conjecture 1.1, see [Atk].

Other variations on the Noonan–Zeilberger Conjecture have also been studied in the literature, including the case of consecutive patterns [Eli], set partitions [Sag], and patterns with infinite support (see [V3]). Theorem 1.2 can perhaps be further extended to show that for some set \mathcal{F} the g.f. for $\{C_n(\mathcal{F})\}$ is not ADE (see e.g. [Sta]). We plan to explore this problem in [GP3].

10.12. Some basic ideas in this paper have been first developed in a much easier case of Wang tilings of the square, studied in [GP2] (see also [Kit, §9.8.2]). Curiously, a key tool is Theorem 2.1, which we proved in [GP1] by a short self-contained argument. Note that this lemma is used in [GP1] also to prove the non-P-recursiveness of the probabilities of returns in linear groups, but is applied in a very different way.

Acknowledgements. We are very grateful to Michael Albert, Matthias Aschenbrenner, Milkós Bóna, Alexander Burstein, Ted Dokos, Lance Fortnow, Joanna Garrabrant, Ira Gessel, Joshua Grochow, Greg Kuperberg, Andrew Marks, Sam Miner, Alejandro Morales, Greta Panova, Andrew Soffer, Richard Stanley, Vince Vatter and Jed Yang for helpful conversations and useful comments at different stages of the project. Special thanks to Doron Zeilberger for telling us about importance of the conjecture and his ambivalence about its validity. The first author was partially supported by the University of California Eugene V. Cota-Robles Fellowship; the second author was partially supported by the NSF.

REFERENCES

- [AAK] M. H. Albert, M. D. Atkinson and M. Klazar, The enumeration of simple permutations, *J. Integer Seq.* **6** (2003), Art. 03.4.4, 18 pp.
- [A+] M. H. Albert, C. Homberger, J. Pantone, N. Shar and V. Vatter, Generating Permutations With Restricted Containers, in preparation.
- [AF] N. Alon and E. Friedgut, On the number of permutations avoiding a given pattern, *J. Combin. Theory, Ser. A* **89** (2000), 133–140.
- [Atk] M. D. Atkinson, Restricted permutations, *Discrete Math.* **195** (1999), 27–38.
- [Bar] R. Barbanchon, On unique graph 3-colorability and parsimonious reductions in the plane, *Theor. Comput. Sci.* **319** (2004), 455–482.
- [BBF] R. Beigel, H. Buhrman and L. Fortnow, NP might not be as easy as detecting unique solutions, in *Proc. 30th STOC* (Dallas, TX, 1998), ACM, New York, 1999, 203–208.
- [Ber] R. Berger, The undecidability of the domino problem, *Memoirs AMS* **66** (1966), 72 pp.
- [Boas] P. van Emde Boas, The convenience of tilings, in *Complexity, logic, and recursion theory*, Dekker, New York, 1997, 331–363.
- [Bóna] M. Bóna, *Combinatorics of permutations*, CRC Press, Boca Raton, FL, 2012.
- [BHPV] M. Bóna, C. Homberger, J. Pantone and V. Vatter, Pattern-avoiding involutions: exact and asymptotic enumeration; [arXiv:1310.7003](https://arxiv.org/abs/1310.7003).
- [BRV] R. Brignall, N. Ruškuc and V. Vatter, Simple permutations: decidability and unavoidable substructures, *Theor. Comput. Sci.* **391** (2008), 150–163.
- [BP] A. Burstein and J. Pantone, Two examples of unbalanced Wilf-equivalence, *J. Combin.*, to appear.
- [CR] S. C. Cater and R. W. Robinson, Exponents of 2 in the numbers of unlabeled graphs and tournaments, *Congr. Numer.* **82** (1991), 139–155.
- [CG] A. R. Conway and A. J. Guttmann, On 1324-avoiding permutations, *Adv. Applied Math.* **64** (2015), 50–69.
- [EV] M. Elder and V. Vatter, Problems and Conjectures presented at Third Permutation Patterns Conference (2005); [arXiv:math/0505504](https://arxiv.org/abs/math/0505504).
- [Eli] S. Elizalde, A survey of consecutive patterns in permutations, to appear in *Recent Trends in Combinatorics*, IMA, 2015.
- [FGS] P. Flajolet, S. Gerhold and B. Salvy, On the non-holonomic character of logarithms, powers, and the n th prime function, *Electron. J. Combin.* **11** (2004/06), A2, 16 pp.
- [FS] P. Flajolet and R. Sedgewick, *Analytic Combinatorics*, Cambridge Univ. Press, Cambridge, 2009.
- [Fox] J. Fox, Stanley–Wilf limits are typically exponential, to appear in *Advances in Mathematics*.
- [GW] H. Galperin and A. Wigderson, Succinct representations of graphs, *Inf. Control* **56** (1983), 183–198.
- [Gar] S. Garoufalidis, G -functions and multisum versus holonomic sequences, *Advances Math.* **220** (2009), 1945–1955.
- [GP1] S. Garrabrant and I. Pak, Words in Linear Groups, Automata and P-recursiveness; [arXiv:1502.06565](https://arxiv.org/abs/1502.06565)
- [GP2] S. Garrabrant and I. Pak, Counting with Wang tiles, in preparation.
- [GP3] S. Garrabrant and I. Pak, How to be non-ADE, in preparation.
- [Ges] I. Gessel, Symmetric Functions and P-Recursiveness, *J. Combin. Theory, Ser. A* **53** (1990), 257–285.
- [GI] D. Gottesman and S. Irani, The quantum and classical complexity of translationally-invariant tiling and Hamiltonian problems, *Theory Comput.* **9** (2013), 31–116.
- [GS] B. Grünbaum and G. C. Shephard, *Tilings and Patterns*, W. H. Freeman, New York, 1989.
- [GM] S. Guillemot and D. Marx, Finding small patterns in permutations in linear time, in *Proc. 25th SODA*, SIAM, 2014, 82–101.
- [HIS] J. Hartmanis, N. Immerman and V. Sewelson, Sparse Sets in NP-P: EXPTIME versus NEXPTIME, *Inf. Control* **65** (1985), 158–181.
- [HMU] J. E. Hopcroft, R. Motwani and J. D. Ullman, *Introduction to Automata Theory, Languages, and Computation* (Third Ed.), Addison Wesley, Reading, MA, 2003.
- [Kit] S. Kitaev, *Patterns in Permutations and Words*, Springer, Berlin, 2011.
- [Kla] M. Klazar, Some general results in combinatorial enumeration, in *Proc. Permutation Patterns (St. Andrews 2007)*, Cambridge Univ. Press, 2010, 3–40; [arXiv:0803.4292](https://arxiv.org/abs/0803.4292).
- [MT] A. Marcus and G. Tardos, Excluded permutation matrices and the Stanley-Wilf conjecture, *J. Combin. Theory, Ser. A* **107** (2004), 153–160.
- [MR] D. Marinov and R. Radoičić, Counting 1324-avoiding permutations, *Electron. J. Combin.* **9** (2002), RP13, 9 pp.
- [NZ] J. Noonan and D. Zeilberger, The enumeration of permutations with a prescribed number of forbidden patterns, *Adv. Appl. Math.* **17** (1996), 381–407.
- [Odl] A. M. Odlyzko, Asymptotic enumeration methods, in *Handbook of Combinatorics*, Vol. 2, Elsevier, Amsterdam, 1995, 1063–1229.
- [Oll] N. Ollinger, Tiling the plane with a fixed number of polyominoes, in *Lecture Notes Comput. Sci.* **5457**, Springer, Berlin, 2009, 638–647.
- [P1] I. Pak, Tile Invariants: New Horizons, *Theor. Comp. Sci.* **303** (2003), 303–331.
- [P2] I. Pak, Partition bijections, a survey, *Ramanujan J.* **12** (2006), 5–75.

- [Pap] C. H. Papadimitriou, *Computational complexity*, Addison-Wesley, Reading, MA, 1994.
- [Poon] B. Poonen, Undecidable problems: a sampler, in *Interpreting Gödel: Critical essays* (J. Kennedy, editor), Cambridge Univ. Press, 2014, 211–241.
- [Rob] R. M. Robinson, Undecidability and nonperiodicity for tilings of the plane, *Invent. Math.* **12** (1971), 177–209.
- [Rio] J. Riordan, *Combinatorial Identities*, John Wiley, New York, 1968.
- [Sag] B. E. Sagan, Pattern avoidance in set partitions, *Ars Combin.* **94** (2010), 79–96.
- [Sip] M. Sipser, *Introduction to the Theory of Computation*, PWS, Boston, 1997.
- [OEIS] N. J. A. Sloane, *The On-Line Encyclopedia of Integer Sequences*, <http://oeis.org>.
- [ST] J. E. S. Socolar and J. M. Taylor, An aperiodic hexagonal tile, *J. Combin. Theory, Ser. A* **118** (2011), 2207–2231.
- [Sta] R. P. Stanley, *Enumerative combinatorics*, Vol. 1, 2, Cambridge Univ. Press, Cambridge, UK, 1997/9.
- [Ste] E. Steingrímsson, Some open problems on permutation patterns, in *Surveys in combinatorics*, Cambridge Univ. Press, Cambridge, 2013, 239–263.
- [Val] L. Valiant, Accidental algorithms, in *Proc. 47th FOCS* (2006), Berkeley, CA, IEEE Press, 509–517.
- [V1] V. Vatter, Problems and conjectures presented at the problem session, in *Permutation Patterns*, LMS Lecture Notes Series 376, Cambridge Univ. Press, 2010, 339–345.
- [V2] V. Vatter, Enumeration schemes for restricted permutations, *Combin. Probab. Comput.* **17** (2008), 137–159.
- [V3] V. Vatter, Permutation Classes, in *The Handbook of Enumerative Combinatorics*, CRC Press, Boca Raton, FL, 2015, 753–834.
- [Wang] H. Wang, Proving theorems by pattern recognition II, *Bell System Tech. J.* **40** (1961), 1–41.
- [Wilf] H. S. Wilf, What is an answer?, *Amer. Math. Monthly* **89** (1982), 289–292.
- [Yang] J. Yang, Rectangular tileability and complementary tileability are undecidable, *Europ. J. Combin.* **41** (2014), 20–34.
- [Zei] D. Zeilberger, Enumerative and Algebraic Combinatorics, in *The Princeton companion to mathematics* (T. Gowers, J. Barrow-Green and I. Leader, editors), Princeton Univ. Press, Princeton, NJ, 2008, 550–561.