Length-Four Pattern Avoidance in Inversion Sequences

Letong Hong
University of Oxford
Oxford, Oxfordshire, U.K.
clhong@alum.mit.edu

Rupert Li
Massachusetts Institute of Technology
Cambridge, MA, U.S.A.
rupertli@mit.edu

Submitted: Dec 30, 2021; Accepted: Nov 15, 2022; Published: Dec 2, 2022

Abstract

Inversion sequences of length n are integer sequences $e_1, \ldots, e_n$ with $0 \leqslant e_i < i$ for all i, which are in bijection with the permutations of length n. In this paper, we classify all Wilf equivalence classes of pattern-avoiding inversion sequences of length-4 patterns except for one case (whether $3012 \equiv 3201$) and enumerate some of the length-4 pattern-avoiding inversion sequences that are in the OEIS.

Mathematics Subject Classifications: 05A05

1 Introduction

Pattern avoidance for permutations is a robust and well-established branch of enumerative combinatorics. We refer readers to Stanley [18] for an overview of this field, and to Simion and Schmidt [16] in 1985 for the first systematic study of pattern avoidance on permutations. Classical pattern avoidance represents permutations using one-line notation $\pi = \pi_1 \cdots \pi_n$; an alternative representation for permutations is using *inversion sequences* $e = e_1 \cdots e_n$, sequences of integers such that $0 \leqslant e_i < i$ for all i. Inversion sequences are in natural bijection with permutations via the well-known Lehmer code [8], an example of an inversion table: one can biject an inversion sequence e to a permutation π via ensuring that for each i, there exist e_i values $j < i$ such that $\pi_j > \pi_i$. Inversion sequences have been studied in many contexts and fields, not just pattern avoidance; for example, see Savage and Schuster [15].

The study of pattern avoidance on inversion sequences was concurrently initiated by Mansour and Shattuck [12] in 2015 and Corteel, Martinez, Savage, and Weselcouch [6] in 2016. The former obtained the explicit number and/or generating function of inversion sequences avoiding any element of S_3; the latter further enumerated the number of pattern-avoiding sequences for all patterns of length 3 and related these quantities to well-known combinatorial sequences including the Bell numbers, Euler up/down numbers, Fibonacci numbers, and Schröder numbers. For patterns of length 4, Chern [4] proved the exact formula for 0012-avoiding inversions, answering a conjecture by Lin and Ma (see the end of [9]) in 2020. However, the enumeration, or even the determination of the Wilf equivalence classes, for all other patterns of length 4 remains open. For simultaneous avoidance of multiple patterns, Lin and Yan [10] in 2020 studied inversion sequences avoiding certain combinations of two length-3 patterns by establishing correspondences with objects enumerated by the Bell numbers, Fishburn numbers, powered Catalan numbers, semi-Baxter numbers, and 3-noncrossing partitions. In 2018, Martinez and Savage [13] rephrased and generalized the question, investigating the avoidance of triples of binary relations, that is, no simultaneous appearances $e_iR_1e_j$, $e_jR_2e_k$, and $e_iR_3e_k$ are allowed to appear with $i < j < k$ for some given $R_1, R_2, R_3 \in \{<, >, \leqslant, \geqslant, =, \neq, -\}$. On the other hand, in 2019 Auli and Elizalde [1] enumerated the length-3 consecutive pattern-avoiding inversion sequences as well as classified consecutive patterns up to length 4 according to the corresponding Wilf equivalence relations. In a following 2021 paper [3], the same authors gave a complete list of generalized Wilf equivalences between hybrid vincular patterns of length 3, completing the classification of Wilf equivalence classes for all vincular patterns of length 3. They further built on Martinez and Savage's framework and extended the enumeration to inversion sequences avoiding $e_iR_1e_{i+1}R_2e_{i+2}$ configurations [2] in 2019.

Our main result classifies all Wilf equivalence classes for length-4 patterns, except one unresolved case of 3012 possibly belonging to the last class, as demarcated below by a question mark. A computer search for lengths $n \leqslant 10$ demonstrates that no other Wilf equivalences are possible: in particular, 2001 agrees with the second Wilf equivalence class $2110 \equiv 2101 \equiv 2011$ for all lengths $n \leqslant 9$, but diverges at $n = 10$.

$$\begin{gathered}
1011 \equiv 1101 \equiv 1110 \\
2110 \equiv 2101 \equiv 2011 \\
0221 \equiv 0212 \\
0312 \equiv 0321 \\
1102 \equiv 1012 \\
2201 \equiv 2210 \\
2301 \equiv 2310 \\
3201 \equiv 3210 \stackrel{?}{\equiv} 3012.
\end{gathered} \tag{1}$$

Theorem 1. *Length-4 patterns satisfy the Wilf equivalences listed in Eq.* (1)*, with possible exception as demarcated with a question mark.*

The paper is organized as follows. In Section 2, we introduce necessary definitions and notation. In Section 3, we establish the aforementioned equivalences using techniques including double induction and direct characterization, together with explicitly constructed correspondences. In Section 4, we enumerate the 0000 and 0111-avoiding inversion sequences. Aside from 0012 as addressed by Chern [4], the only other length-4 pattern that is on the OEIS [17] is 0021, whose enumeration we leave as an open question, which after the writing of the original version of this paper has been resolved (see Remark 24).

2 Preliminaries

For a positive integer n, let $[n]$ denote the set $\{1, 2, \ldots, n\}$. An inversion sequence of length n is a sequence $e = e_1 \cdots e_n$ of integers such that $0 \leqslant e_i < i$ for all $i \in [n]$. We denote the length of e by $|e| = n$. The set of inversion sequences of length n is denoted by $\mathbf{I}_n$, where we use the convention that $\mathbf{I}_0$ contains exactly one sequence, the empty sequence.

Two sequences of integers $\pi = \pi_1 \cdots \pi_k$ and $\sigma = \sigma_1 \cdots \sigma_k$ of the same length are said to be *order isomorphic*, denoted $\pi \sim \sigma$, if both $\pi_i R \pi_j$ and $\sigma_i R \sigma_j$ have the same relation $R \in \{<, =, >\}$, for all $1 \leqslant i, j \leqslant n$. For example, $0212 \sim 5969$. A *pattern* refers to such a sequence of integers $\pi = \pi_1 \cdots \pi_k$.

For an inversion sequence $e \in \mathbf{I}_n$ and a pattern $\pi = \pi_1 \cdots \pi_k$ for $k \leqslant n$, we say e *contains* π as a pattern if there is a not necessarily consecutive subsequence e' of e with $|e'| = k$ such that $e' \sim \pi$. For $S \subseteq [n]$, we let e_S denote the subsequence of e consisting of the elements e_i for $i \in S$, sorted in ascending order of i. Using this notation, e contains π if there exists $S \subseteq [n]$ with $|S| = k$ such that $e_S \sim \pi$. If e does not contain π, it is said to *avoid* π. In particular, if $|e| < |\pi|$, we also say e avoids π. The same definition can be used to define pattern avoidance on permutations.

The *avoidance class* of π is

$$\mathbf{I}_n(\pi) = \{e \in \mathbf{I}_n \mid e \text{ avoids } \pi\}.$$

We say two patterns π and σ are *Wilf equivalent*, denoted $\pi \equiv \sigma$, if for all $n \geqslant 1$, we have $|\mathbf{I}_n(\pi)| = |\mathbf{I}_n(\sigma)|$.

We now define the following generalization of an inversion sequence, as originally introduced by Savage and Schuster [15].

Definition 2. For a finite set of positive integers $S \subset \mathbb{Z}_+$ enumerated in increasing order $s_1 < \cdots < s_n$, an *S-inversion sequence* is a sequence $e = e_1 \cdots e_n$ of length n such that $0 \leqslant e_i < s_i$ for all $i \in [n]$. The set of S-inversion sequences is denoted by $\mathbf{I}_S$.

Notice that for $S = [n]$, we recover the original definition of an inversion sequence of length n, i.e., $\mathbf{I}_{[n]} = \mathbf{I}_n$. We continue to use the same notation for pattern avoidance on S-inversion sequences as on inversion sequences: $\mathbf{I}_S(\pi)$ is the set of S-inversion sequences that avoid π. We note that we define $\mathbf{I}_\emptyset$ to contain the empty sequence, which avoids all patterns, consistent with the previous observation that $\mathbf{I}_{[n]} = \mathbf{I}_n$ when $n = 0$.

Finally, define $e \cdot f$ as the concatenation of sequences e and f. For example, $e \cdot f = 142857$ for $e = 14$ and $f = 2857$.

3 Wilf equivalences of length-4 patterns

Before we prove the Wilf equivalences of length-4 patterns, we present the following useful result.

Theorem 3. *For any $n \geqslant 1$ and a pattern $\pi = \pi_1 \cdots \pi_k$ where $\pi_1 = 0$ and $\pi_i > 0$ for all $i > 1$,*

$$|\mathbf{I}_n(\pi)| = \sum_{S \subseteq [n-1]} |\mathbf{I}_S(\pi_2 \cdots \pi_k)|.$$

Proof. Define $\pi' = \pi_2 \cdots \pi_k$. Notice that all elements of π' are positive. Consider $e \in \mathbf{I}_n$, and let S be the set of indices $i \in [2, n]$ such that $e_i > 0$. We claim that e_S avoids π' if and only if e avoids π. If e_S contains π', then suppose $e_{i_1} \cdots e_{i_{k-1}} \sim \pi'$ for $i_1, \ldots, i_{k-1} \in S$. By definition of S, we have $e_{i_1}, \ldots, e_{i_{k-1}} > 0$, and thus adding in $e_1 = 0$ yields $e_1 e_{i_1} \cdots e_{i_{k-1}} \sim \pi$. Conversely, if e contains π, suppose $e_{i_0} \cdots e_{i_{k-1}} \sim \pi$ for $1 \leqslant i_0 < i_1 < \cdots < e_{i_{k-1}} \leqslant n$. Then as all elements of π are positive except for $\pi_1 = 0$, we have $e_{i_1}, \ldots, e_{i_{k-1}} > e_{i_0} \geqslant 0$, so $e_{i_1} \cdots e_{i_{k-1}} \sim \pi'$ is a subsequence of e_S, which thus contains π'.

Let $S^- = \{s - 1 \mid s \in S\}$. Notice that e_S^- is a S^--inversion sequence, and there is a natural bijection between $\mathbf{I}_{S^-}$ and the elements of $\mathbf{I}_n$ with $e_i > 0$ if and only if $i \in S$. Thus, for a fixed subset $S \subseteq [2, n]$, the number of elements of $\mathbf{I}_n(\pi)$ with $e_i > 0$ if and only if $i \in S$ is equal to $|\mathbf{I}_{S^-}(\pi')|$. Summing over all such subsets S and re-indexing over S^- instead yields the result. □

The following result was initially stated for inversion sequences rather than S-inversion sequences, but we note that the same proof works to obtain the following stronger result.

Theorem 4 ([6, Theorem 5])**.** *For any finite set S of positive integers,*

$$|\mathbf{I}_S(210)| = |\mathbf{I}_S(201)|.$$

This allows us to prove that 0312 and 0321 are Wilf equivalent.

Theorem 5. *For $n \geqslant 1$,*

$$|\mathbf{I}_n(0312)| = |\mathbf{I}_n(0321)| = \sum_{S \subseteq [n-1]} |\mathbf{I}_S(210)|.$$

Proof. The result follows from applying Theorem 4 to Theorem 3. □

3.1 Wilf equivalences by double induction

The following lemma is useful for many of our later results. A *binary word* of length n is an element of $\{0, 1\}^n$, i.e., a string of n zeros and ones. Pattern avoidance on binary words is defined analogously.

Lemma 6. *Let $\pi = \pi_1 \cdots \pi_\ell$ be a pattern of length $\ell \geqslant 2$ such that $\pi_i \in \{0, 1\}$ for all i, and there exists exactly one j such that $\pi_j = 0$. Then for any two integers $j, k \geqslant 0$, the number of binary words of length $j + k$ with j zeros and k ones that avoid π is $\binom{j+\min\{k,\ell-2\}}{j}$.*

Proof. There are $\binom{j+k}{k}$ binary words of length $j+k$ with j zeros and k ones, corresponding to choosing the positions of the ones. If $k \leqslant \ell-2$, as π has $\ell-1$ ones, all of these binary words avoid π, so there are $\binom{j+k}{k} = \binom{j+\min\{k,\ell-2\}}{j}$ valid binary words. If $k \geqslant \ell-1$, suppose π_j for $j \in [\ell]$ is the unique zero in π. If $j=1$ or $j=\ell$, then all but $\ell-2$ of the k ones must be at the beginning or end of the binary word, respectively, which yields $\binom{j+\ell-2}{\ell-2} = \binom{j+\min\{k,\ell-2\}}{j}$ valid binary words. Otherwise $2 \leqslant j \leqslant \ell-1$, and then the i-th and $(i+1)$-th ones of any valid binary word must be consecutive, for all $j-1 \leqslant i \leqslant k+j-\ell$. If this were not the case, then suppose there exists such an i where the i-th and $(i+1)$-th ones are not consecutive. Then the $(i-j+2)$-th through i-th ones, a zero between the i-th and $(i+1)$-th ones, and the $(i+1)$-th through $(i+\ell-j)$-th ones, form a π pattern. Notice that $i-j+2 \geqslant 1$ and $i+\ell-j \leqslant k$, so the indices are valid. Note that this condition is a necessary and sufficient condition for the binary word to avoid π. Thus, the $(j-1)$-th through $(k+j-\ell+1)$-th ones in the binary word are consecutive; viewing this block of $k-\ell+3$ ones as a single entity allows us to determine that there are $\binom{j+k-(k-\ell+2)}{j} = \binom{j+\min\{k,\ell-2\}}{j}$ such valid binary words. $\square$

Lemma 6 allows us to prove the following result.

Theorem 7. *Let $\pi = \pi_1 \cdots \pi_\ell$ and $\sigma = \sigma_1 \cdots \sigma_\ell$ be two patterns of length $\ell \geqslant 3$ such that $\pi_i, \sigma_i \in \{0,1\}$ for all $2 \leqslant i \leqslant \ell$ and $\pi_1 = \sigma_1 = 1$. If there exists exactly one j such that $\pi_j = 0$ and exactly one j' such that $\sigma_{j'} = 0$, then for any finite set S of positive integers, $|\mathbf{I}_S(\pi)| = |\mathbf{I}_S(\sigma)|$.*

Proof. Let $x_{S,j,k}$ denote the number of π-avoiding S-inversion sequences with j zeros and k ones, and similarly define $y_{S,j,k}$ for σ-avoidance. We will prove the refinement that $x_{S,j,k} = y_{S,j,k}$ for all S, j, and k by induction on $|S|$.

When $|S| < \ell$, the result trivially holds as all S-inversion sequences avoid all patterns of length ℓ. For the inductive step, assume the result holds for all S with $|S| = n-1$; we will show the result holds for all S with $|S| = n$ via a second induction on $\min S$. For the base case $\min S = 1$, any $e \in \mathbf{I}_S$ has $e_1 = 0$, which cannot be part of a π or σ pattern, so $e = e_1 \cdots e_n \in \mathbf{I}_S$ avoids π if and only if $e_2 \cdots e_n$ avoids π, and similarly for σ. Hence, $x_{S,j,k} = x_{S\setminus\{1\},j-1,k} = y_{S\setminus\{1\},j-1,k} = y_{S,j,k}$.

Now assume the result holds for all S of size n with $\min S = m-1 \geqslant 1$; we will show the result holds for all S with $S = m$. Consider a given S of size n with $\min S = m \geqslant 2$, and let $S^- = \{s-1 \mid s \in S\}$. Define $\phi : \mathbf{I}_S \to \mathbf{I}_{S^-}$ by $\phi(e_1 \cdots e_n)_i = \max\{e_i - 1, 0\}$. Notice that π-avoidance and σ-avoidance are both preserved under ϕ.

Consider a π-avoiding S-inversion sequence e' with j zeros and k ones. Then $\phi(e')$ has $k+j$ zeros. We claim that for any $d \in \mathbf{I}_{S^-}(\pi)$ with $k+j$ zeros, there exist exactly $\binom{j+\min\{k,\ell-2\}}{j}$ sequences $e \in \mathbf{I}_S(\pi)$ with j zeros and k ones such that $\phi(e) = d$. This would then show that

$$x_{S,j,k} = \binom{j+\min\{k,\ell-2\}}{j} \sum_{i=0}^{n-k-j} x_{S^-,k+j,i}.$$

Consider some $d \in \mathbf{I}_{S^-}(\pi)$ with $k+j$ zeros. As $\phi(e) = d$ and e has j zeros and k ones, we find e is completely determined apart from selecting which k of the $k+j$ zeros in d

become ones in e. As d avoids π, we find that e avoids π if and only if the zeros and ones of e avoid π. By Lemma 6, we find there are $\binom{j+\min\{k,\ell-2\}}{j}$ such choices of zeros and ones of e that avoid π, each yielding a distinct valid $e \in \mathbf{I}_S(\pi)$.

As S^- satisfies the conditions of the inductive hypothesis, it now suffices to similarly show that for any $d \in \mathbf{I}_{S^-}(\sigma)$ with $k+j$ zeros, there exist exactly $\binom{j+\min\{k,\ell-2\}}{j}$ sequences $e \in \mathbf{I}_S(\sigma)$ with j zeros and k ones such that $\phi(e) = d$. The argument is identical to that for π, as σ also satisfies the conditions of Lemma 6. This implies $x_{S,j,k} = y_{S,j,k}$ for all j and k, and by induction, for all S. This completes the proof. $\square$

Corollary 8. *For any finite set S of positive integers,*

$$|\mathbf{I}_S(1011)| = |\mathbf{I}_S(1101)| = |\mathbf{I}_S(1110)|.$$

This implies 1011, 1101, and 1110 are Wilf equivalent over inversion sequences.

Corollary 9. *For $n \geqslant 1$,*

$$|\mathbf{I}_n(0221)| = |\mathbf{I}_n(0212)| = \sum_{S \subseteq [n-1]} |\mathbf{I}_S(110)|.$$

Proof. Theorem 7 implies $|\mathbf{I}_S(110)| = |\mathbf{I}_S(101)|$ for any finite set S of positive integers, from which the result follows by applying Theorem 3. $\square$

The following result augments the method of Theorem 7 to prove that $\pi \cdot \rho$ and $\sigma \cdot \rho$ are Wilf equivalent over S-inversion sequences, where π and σ satisfy the assumptions of Theorem 7 and ρ consists only of twos.

Theorem 10. *Let $\rho = \rho_1 \cdots \rho_h$ be a pattern of length $h \geqslant 0$ such that $\rho_i = 2$ for all i. Let $\pi = \pi_1 \cdots \pi_\ell$ and $\sigma = \sigma_1 \cdots \sigma_\ell$ be two patterns of length $\ell \geqslant 3$ such that $\pi_i, \sigma_i \in \{0,1\}$ for all $2 \leqslant i \leqslant \ell$ and $\pi_1 = \sigma_1 = 1$. If there exists exactly one j such that $\pi_j = 0$ and exactly one j' such that $\sigma_{j'} = 0$, then for any finite set S of positive integers, $|\mathbf{I}_S(\pi\cdot\rho)| = |\mathbf{I}_S(\sigma\cdot\rho)|$.*

Proof. When $h = 0$, the result follows from Theorem 7.

For convenience, define $\pi' = \pi \cdot \rho$ and $\sigma' = \sigma \cdot \rho$. We use a similar double induction approach as in Theorem 7. Let the *terminal h-repeat statistic* of an S-inversion sequence e be the largest integer r such that there are at least r zeros in e, and letting z denote the index of the r-th zero in e, then there exist positive integers $z < i_1 < i_2 < \cdots < i_h \leqslant |S|$ where $e_{i_1} = e_{i_2} = \cdots = e_{i_h} > 0$; if no such r exists, define the terminal h-repeat statistic to be 0. For example, the terminal 1-repeat statistic is simply the number of non-terminal zeros, where a terminal zero only has zeros after it, if anything.

Let $x_{S,j,k,r}$ denote the number of π'-avoiding S-inversion sequences with j zeros, k ones, and terminal h-repeat statistic r, and similarly define $y_{S,j,k,r}$ for σ'-avoidance. We will prove the refinement that $x_{S,j,k,r} = y_{S,j,k,r}$ for all S, j, k, r by induction on $|S|$.

When $|S| < \ell + h$, the result trivially holds as all S-inversion sequences avoid all patterns of length $\ell + h \geqslant 3$. For the inductive step, assume the result holds for all S with $|S| = n-1$; we will show the result holds for all S with $|S| = n$ via a second induction on

$\min S$. For the base case $\min S = 1$, any $e \in \mathbf{I}_S$ has $e_1 = 0$, which cannot be part of a π' or σ' pattern, so $e = e_1 \cdots e_n \in \mathbf{I}_S$ avoids π' if and only if $e_2 \cdots e_n$ avoids π', and similarly for σ'. Hence, $x_{S,j,k,r} = y_{S,j,k,r}$ using the inductive hypothesis for $S \setminus \{1\}$.

Now assume the result holds for all S of size n with $\min S = m - 1 \geqslant 1$; we will show the result holds for all S with $S = m$. Consider a given S of size n with $\min S = m \geqslant 2$, and let $S^- = \{s - 1 \mid s \in S\}$. Define ϕ as in Theorem 7, and notice that π'-avoidance and σ'-avoidance are both preserved under ϕ.

Suppose $d \in \mathbf{I}_{S^-}(\pi')$, and consider the S-inversion sequences $e \in \mathbf{I}_S$ such that $\phi(e) = d$. Similarly, suppose $d' \in \mathbf{I}_{S^-}(\sigma')$, and consider the S-inversion sequences $e' \in \mathbf{I}_S$ such that $\phi(e') = d'$. Suppose d and d' both have $j + k$ zeros and terminal h-repeat statistic r. By the inductive hypothesis, the number of such d equals the number of such d'.

As d and d' both have $j + k$ zeros, e and e' must each have $j + k$ total zeros and ones. Now restrict consideration to those e and e' that have j zeros and k ones. As d is π'-avoiding, e avoids π' if and only if no π pattern occurs within its first r zero and one entries. Similarly, e' avoids σ' if and only if no σ pattern occurs within its first r zero and one entries.

Consider some choice of zeros and ones for the last $j + k - r$ zeros of d and d', i.e., some binary sequence in $\{0, 1\}^{j+k-r}$. We claim that the number of e whose last $j + k - r$ zeros and ones follow this binary sequence equals the number of e' whose last $j + k - r$ zeros and ones follow this binary sequence. Notice that all such e and e' have the same terminal h-repeat statistic r': if the binary sequence contains at least h ones, then r' is the number of zeros before the h-th-to-last one in e or respectively e'; otherwise, r' is the number of zeros within the first r zeros and ones of e or respectively e'. As e and e' both have k ones and j zeros, and this binary sequence fixes the terminal h-repeat statistic, this would be a stronger refinement that implies $x_{S,j,k,r'} = y_{S,j,k,r'}$ for all j, k, r'.

Suppose this binary sequence has j' zeros and $k' = j + k - r - j'$ ones, where we may assume $j' \leqslant j$ and $k' \leqslant k$, as otherwise no valid e or e', with k ones and j zeros, exist. Hence both e and e' must have $j - j'$ zeros and $k - k'$ ones among the positions of the first r zeros in d and d', respectively. These zeros and ones in e must avoid π, and these zeros and ones in e' must avoid σ. Then Lemma 6 implies that the number of such e equals the number of such e', namely equaling $\binom{j-j'+\min\{k-k',\ell-2\}}{j-j'}$.

This implies $x_{S,j,k,r'} = y_{S,j,k,r'}$ for all j, k, r', and by induction, for all S, which completes the proof. $\square$

Corollary 11. *For any finite set S of positive integers,*

$$|\mathbf{I}_S(1012)| = |\mathbf{I}_S(1102)|.$$

This implies 1012 and 1102 are Wilf equivalent over inversion sequences.

Similar to Theorem 10, which appends twos to π and σ that satisfy the assumptions of Theorem 7, the following result augments the method of Theorem 7 to prove that $\rho \cdot \pi$ and $\rho \cdot \sigma$ are Wilf equivalent over S-inversion sequences, where π and σ satisfy the assumptions of Lemma 6 and ρ consists only of twos.

Theorem 12. *Let $\rho = \rho_1 \cdots \rho_h$ be a pattern of length $h \geqslant 1$ such that $\rho_i = 2$ for all i. Let $\pi = \pi_1 \cdots \pi_\ell$ and $\sigma = \sigma_1 \cdots \sigma_\ell$ be two patterns of length $\ell \geqslant 2$ such that $\pi_i, \sigma_i \in \{0, 1\}$ for all $i \in [\ell]$, and there exists exactly one j such that $\pi_j = 0$ and exactly one j' such that $\sigma_{j'} = 0$. Then for any finite set S of positive integers, $|\mathbf{I}_S(\rho \cdot \pi)| = |\mathbf{I}_S(\rho \cdot \sigma)|$.*

Proof. For convenience, define $\pi' = \rho{\cdot}\pi$ and $\sigma' = \rho{\cdot}\sigma$. We use an almost identical approach as in Theorem 10, reversing the definition of the terminal h-repeat statistic. Let the *initial h-repeat statistic* of an S-inversion sequence e be the largest integer r such that there are at least r zeros in e, and letting z denote the index of the r-th-to-last zero in e, then there exist positive integers $1 \leqslant i_1 < i_2 < \cdots < i_h < z$ where $e_{i_1} = e_{i_2} = \cdots = e_{i_h} > 0$; if no such r exists, define the initial h-repeat statistic to be 0. For example, the initial 1-repeat statistic is simply the number of non-initial zeros, where an initial zero only has zeros before it, if anything.

Let $x_{S,j,k,r}$ denote the number of π'-avoiding S-inversion sequences with j zeros, k ones, and initial h-repeat statistic r, and similarly define $y_{S,j,k,r}$ for σ'-avoidance. We will prove the refinement that $x_{S,j,k,r} = y_{S,j,k,r}$ for all S, j, k, r by induction on $|S|$.

The proof then follows the same reasoning as that of Theorem 10. For sake of brevity and clarity, we comment on some of the minor differences between the proofs. We assume $h \geqslant 1$ so that neither π' and σ' start with a 0, allowing the base case $\min S = 1$ for the second induction to hold; this in turn allows us to lift the restriction that $\pi_1 = \sigma_1 = 1$. Using the same notation as in the proof of Theorem 10, the characterization of e becomes as follows: e avoids π' if and only if no π pattern occurs within its last r zero and one entries, and similarly for e' avoiding σ'. We then consider some binary sequence for the first $j + k - r$ zeros of d and d', as opposed to the last, where fixing this binary sequence fixes the initial h-repeat statistic of e and e', so the proof proceeds identically. ☐

Corollary 13. *For any finite set S of positive integers,*

$$|\mathbf{I}_S(2011)| = |\mathbf{I}_S(2101)| = |\mathbf{I}_S(2110)|.$$

This implies 2011, 2101, and 2110 are Wilf equivalent over inversion sequences.

Corollary 14. *For any finite set S of positive integers,*

$$|\mathbf{I}_S(2201)| = |\mathbf{I}_S(2210)|.$$

This implies 2201 and 2210 are Wilf equivalent over inversion sequences.

Theorem 15. *Let $\pi = \pi_1 \cdots \pi_\ell$ and $\sigma = \sigma_1 \cdots \sigma_\ell$ be two patterns of length $\ell \geqslant 2$ such that $\pi_i, \sigma_i \in \{0, 1\}$ for all $i \in [\ell]$, and there exists exactly one j such that $\pi_j = 0$ and exactly one j' such that $\sigma_{j'} = 0$. Then for any finite set S of positive integers, $|\mathbf{I}_S(23{\cdot}\pi)| = |\mathbf{I}_S(23{\cdot}\sigma)|$.*

Proof. For convenience, define $\pi' = 23 \cdot \pi$ and $\sigma' = 23 \cdot \sigma$. We use a similar double induction approach as in Theorem 12.

Let the *initial non-inversion statistic* of an S-inversion sequence e be the largest integer z such that there are at least z zeros in e, and there does not exist two elements $0 < e_{i_1} <$

e_{i_2} of e where $i_1 < i_2$ and both come before the z-th zero in e; this statistic can equal zero. Furthermore, let the *initial positive set* of an S-inversion sequence e be the set P of integers i for $0 \leqslant i < z$, where z is the initial non-inversion statistic of e, such that there exists a positive element between the i-th and $(i+1)$-th zeros of e, where "between the zeroth and first zeros of e" is interpreted to mean before the first zero of e.

Let $x_{S,j,k,z,P}$ denote the number of π'-avoiding S-inversion sequences with j zeros, k ones, initial non-inversion statistic z, and initial positive set P, and similarly define $y_{S,j,k,z,P}$ for σ'-avoidance. We will prove the refinement that $x_{S,j,k,z,P} = y_{S,j,k,z,P}$ for all S, j, k, z, P by induction on $|S|$.

The initial argument, from the base cases of $|S| < \ell + 2$ up to the beginning of the second inductive step using ϕ, follow the same reasoning as in Theorem 12. We use the same definitions for S^- and ϕ, where we note that π'-avoidance and σ'-avoidance are both preserved under ϕ.

Suppose $d \in \mathbf{I}_{S^-}(\pi')$, and consider the S-inversion sequences $e \in \mathbf{I}_S$ such that $\phi(e) = d$. Similarly, suppose $d' \in \mathbf{I}_{S^-}(\sigma')$, and consider the S-inversion sequences $e' \in \mathbf{I}_S$ such that $\phi(e') = d'$. Suppose d and d' both have $j+k$ zeros, initial non-inversion statistic z, and initial positive set P. By the inductive hypothesis, the number of such d equals the number of such d'.

As d and d' both have $j+k$ zeros, e and e' must each have $j+k$ total zeros and ones. Now restrict consideration to those e and e' that have j zeros and k ones. As d is π'-avoiding, e avoids π' if and only if no π pattern occurs within its last $j+k-z$ zero and one entries. Similarly, e' avoids σ' if and only if no σ pattern occurs within its last $j+k-z$ zero and one entries.

Consider some choice of zeros and ones for the first z zeros of d and d', i.e., some binary sequence in $\{0,1\}^z$. We claim that the number of e whose first z zeros and ones follow this binary sequence equals the number of e' whose first z zeros and ones follow this binary sequence. However, we first show that all such e and e' have the same initial non-inversion statistic z' and initial positive set P', as then the claim yields a stronger refinement that implies $x_{S,j,k,z',P'} = y_{S,j,k,z',P'}$ for all j, k, z', P'.

If the binary sequence contains no ones, then $z' = z$ and $P' = P$. Otherwise, the ones in this binary sequence will cause $z' < z$ and may cause P' to change. If P is empty, then notice z' is simply the number of zeros in this binary sequence, and P' is defined according to which zeros in the binary sequence have ones in between them. If P is non-empty, suppose i_1 is the index of the first one in the binary sequence; and let i_2 be the minimum element of $(P \cup \{z\}) \cap [i_1, \infty)$. Then z' is the number of zeros within the first i_2 elements of the binary sequence, and P' is uniquely determined from P and the binary sequence. Hence, given z, P, and the binary sequence, z' and P' are uniquely determined, as desired.

We conclude the proof by proving our claim that the number of such e equals the number of such e'. Suppose this binary sequence has $j' \leqslant j$ zeros and $k' = z - j' \leqslant k$ ones. Then both e and e' have $j - j'$ zeros and $k - k'$ ones among the positions of the last $j+k-z$ zeros in d and d', respectively. These zeros and ones in e must avoid π, and these zeros and ones in e' must avoid σ. Lemma 6 implies that the number of such e

equals the number of such e', namely equaling $\binom{j-j'+\min\{k-k',\ell-2\}}{j-j'}$.

This implies $x_{S,j,k,z',P'} = y_{S,j,k,z',P'}$ for all j, k, z', P', and by induction, for all S, which completes the proof. $\square$

Corollary 16. *For any finite set S of positive integers,*

$$|\mathbf{I}_S(2301)| = |\mathbf{I}_S(2310)|.$$

This implies 2301 and 2310 are Wilf equivalent over inversion sequences.

3.2 Wilf equivalences by characterization

For a sequence $e_1 \cdots e_n$ of nonnegative integers, a position $j \in [n]$ is a *weak left-to-right maximum* if $e_i \leqslant e_j$ for all $i < j$. We use this definition to characterize 3210 and 3201-avoiding inversion sequences, allowing us to construct an explicit bijection between $\mathbf{I}_n(3210)$ and $\mathbf{I}_n(3201)$. First, we characterize 3210-avoiding inversion sequences.

Lemma 17. *The 3210-avoiding inversion sequences are precisely those that can be partitioned into three weakly increasing subsequences.*

Proof. Suppose $e \in \mathbf{I}_n$ has such a partition $e_{x_1} \leqslant e_{x_2} \leqslant \cdots \leqslant e_{x_t}$, $e_{y_1} \leqslant e_{y_2} \leqslant \cdots \leqslant e_{y_r}$, and $e_{z_1} \leqslant e_{z_2} \leqslant \cdots \leqslant e_{z_{n-t-r}}$. If there exist $i < j < k < \ell$ such that $e_i > e_j > e_k > e_\ell$, then no two of i, j, k, ℓ can both be in any of the three sets $\{x_1, \ldots, x_t\}$, $\{y_1, \ldots, y_r\}$, and $\{z_1, \ldots, z_{n-t-r}\}$, but this is impossible due to the Pigeonhole principle. Therefore, e avoids 3210. Conversely, if e is 3210-avoiding, let $x = (x_1, \ldots, x_t)$ be the sequence of weak left-to-right maxima of e. Then $e_{x_1} \leqslant e_{x_2} \leqslant \cdots \leqslant e_{x_t}$. We then let $y = (y_1, \ldots, y_r)$ be the sequence of weak left-to-right maxima of the sequence obtained by deleting positions $\{x_1, \ldots, x_t\}$ from e, and in general we call these positions *weak 2nd left-to-right maxima.* Similarly $e_{y_1} \leqslant e_{y_2} \leqslant \cdots \leqslant e_{y_r}$. We then consider the remaining terms of the sequence, and take $i, j \notin (\{x_1, \ldots, x_t\}) \cup (\{y_1, \ldots, y_r\})$ where $i < j$. The fact that i is not included in $\{y_1, \ldots, y_r\}$ implies there exists some $v \in \{y_1, \ldots, y_r\}$ such that $v < i$ and $e_v > e_i$. The fact that v is not a weak left-to-right maxima implies there exists some u such that $u < v$ and $e_u > e_v$. Now we have $e_u > e_v > e_i$ with $u < v < i < j$. Thus, to avoid 3210, we must have $e_i \leqslant e_j$. Both directions are thus concluded. $\square$

Now, we characterize 3201-avoiding inversion sequences.

Lemma 18. *Let $(e_1, e_2, \ldots, e_n) \in \mathbf{I}_n$. For any $i \in [n]$, let M_i^1 and M_i^2 be the largest and second largest value among $\{e_1, e_2, \ldots, e_{i-1}\}$, respectively. Then $e \in \mathbf{I}_n(3201)$ if and only if for every $i \in [n]$, the entry e_i is either a weak left-to-right maximum, a weak 2nd left-to-right maximum, or for every $j > i$, we have $e_j \leqslant e_i$ or $e_j \geqslant M_i^2$.*

Proof. Let $e \in \mathbf{I}_n$ satisfy the conditions of Lemma 18 and, for the sake of contradiction, assume that there exists $i < j < k < \ell$ such that $e_k < e_\ell < e_j < e_i$. Notice that we have $M_k^1 \geqslant e_i$ and thus $M_k^2 \geqslant e_j$. Then $e_k < e_\ell < e_j \leqslant M_k^2$, a contradiction to our assumption.

Conversely, suppose e is 3201-avoiding. If e_i is neither a weak left-to-right maximum nor a weak 2nd left-to-right maximum, then there exists some 2nd maximum value M_i^2

such that $M_i^2 = e_v > e_i$ for some $v < i$. By definition of 2nd maximum, there is some maximum value $M_v^1 = e_u > e_v$ for some $u < v$. To avoid 3201, we must have that for all $j > i$, $e_j \leqslant e_i$ or $e_j \geqslant e_v = M_i^2$. $\square$

Combining these two results allows us to prove $3210 \equiv 3201$.

Theorem 19. *For $n \geqslant 1$,*

$$|\mathbf{I}_n(3210)| = |\mathbf{I}_n(3201)|.$$

Proof. We exhibit a bijection based on the characterizations in Lemma 17 and Lemma 18.

Given $e \in \mathbf{I}_n(3210)$, we define $f \in \mathbf{I}_n(3201)$ as follows. Let $e_{x_1} \leqslant e_{x_2} \leqslant \cdots \leqslant e_{x_t}$ and $e_{y_1} \leqslant e_{y_2} \leqslant \cdots \leqslant e_{y_r}$ be the subsequences of weak left-to-right maxima and weak 2nd left-to-right maxima of e, respectively, and let $e_{z_1} \leqslant e_{z_2} \leqslant \cdots \leqslant e_{z_{n-t-r}}$ be the remaining entries.

For $i \in [t]$, we set $f_{x_i} = e_{x_i}$ and for $i \in [r]$, we set $f_{y_i} = e_{y_i}$. For each $j \in [n - t - r]$, we extract an element of the multiset $Z = \{e_{z_1}, e_{z_2}, \ldots, e_{z_{n-t-r}}\}$ and assign it to $f_{z_1}, f_{z_2}, \ldots, f_{z_{n-t-r}}$, one at a time in order, as follows:

$$f_{z_j} := \max\{k \mid k \in Z - \{f_{z_1}, f_{z_2}, \ldots, f_{z_{j-1}}\} \text{ and } k < M_{z_j}^2\}.$$

By definition, f will satisfy the characterization property in Lemma 18 of $\mathbf{I}_n(3201)$. One can see that this is invertible, hence a bijection. $\square$

A computer search proves no other length-4 patterns are Wilf equivalent, except possibly 3012 and the aforementioned Wilf equivalence class $3210 \equiv 3201$. We leave this last case as an open question.

Conjecture 20. For $n \geqslant 1$,

$$|\mathbf{I}_n(3201)| = |\mathbf{I}_n(3012)|.$$

This has been verified for all $n \leqslant 12$.

4 Enumeration of inversion sequences avoiding patterns of length 4

Define a *label-increasing tree* on n vertices to be a rooted unordered tree in which each vertex is labeled with a distinct label from the set $\{0, \ldots, n-1\}$ and labels increase along any path from the root to a leaf. Then define a *label-increasing tree with branching bounded by k* to be a label-increasing tree such that each vertex has at most k children. Let $L_{n,k}$ denote the set of n-vertex label-increasing trees with branching bounded by k.

Kuznetsov, Pak, and Postnikov [7] showed that $L_{n,2}$ is in bijection with the *up/down permutations*, that is, the permutations π of $[n]$ such that $\pi_1 < \pi_2 > \pi_3 < \pi_4 > \cdots$; the number of up/down permutations is the Euler number E_n, whose exponential generating function is well-known, namely

$$\sum_{n\geqslant 0} E_n \frac{x^n}{n!} = \tan(x) + \sec(x).$$

Corteel, Martinez, Savage, and Weselcouch [6] proved that $|\mathbf{I}_n(000)| = E_{n+1}$ via a bijection between $\mathbf{I}_n(000)$ and $L_{n+1,2}$. We generalize their result to patterns $00\cdots 0$ of any length k.

Theorem 21. *For $k \geqslant 1$, let $\pi = 00\cdots 0$ be the pattern consisting of k zeros. Then for all $n \geqslant 1$,*

$$|\mathbf{I}_n(\pi)| = |L_{n+1,k-1}|.$$

Proof. Notice that $\mathbf{I}_n(\pi)$ is the set of inversion sequences of length n where each entry occurs at most $k-1$ times, and $L_{n+1,k-1}$ is the set of label-increasing trees of $n+1$ vertices labeled $0,\ldots,n$, with branching bounded by k. Then it is easy to see that the mapping sending $T \in L_{n+1,k-1}$ to $e \in \mathbf{I}_n(\pi)$, where e_i is the parent of i in T, is a bijection between $L_{n+1,k-1}$ and $\mathbf{I}_n(\pi)$. □

Theorem 21 implies $\mathbf{I}_n(0000)$ is in bijection with the label-increasing trees with branching bounded by 3, which is OEIS sequence A297196 [17]. Theorem 21 also enables us to determine the exponential generating function for $|\mathbf{I}_n(00\cdots 0)|$, as Riordan [14] showed the exponential generating function

$$T_k(x) = \sum_{n\geqslant 0} |L_{n,k}| \frac{x^n}{n!} \tag{2}$$

satisfies the differential equation

$$T_k'(x) = \sum_{i=0}^{k} \frac{(T_k(x)-1)^i}{i!}.$$

In other words, $T_k(x)$ satisfies $T_k(0) = 1$ and

$$k!\,T_k'(x) = (T_k(x))^k + \sum_{m=0}^{k-2} c_{m,k}(T_k(x))^m,$$

where

$$c_{m,k} = \frac{1}{m!}\left(\sum_{j=0}^{k-m}(-1)^j k(k-1)\cdots(j+1)\right).$$

These are the same coefficients satisfying

$$k!\sum_{j=0}^{k}\frac{x^j}{j!} = (x+1)^k + \sum_{m=0}^{k-2} c_{m,k}(x+1)^m$$

coming from the differential equation.

Let $L'_{n,k}$ denote the set of n-vertex label-increasing trees with unbounded root degree and branching bounded by k at all other nodes. An alternative way to think about these combinatorial objects is to consider the possible ways how n sufficiently large boxes can contain each other under the condition that each box may contain at most k (themselves possibly nested) boxes. Similar to Theorem 21, we have the following result for patterns of the form $011\cdots 1$.

Theorem 22. *For $k \geqslant 1$, let $\pi = 011\cdots 1$ be the pattern consisting of a zero and k ones. Then for all $n \geqslant 1$,*

$$|\mathbf{I}_n(\pi)| = |L'_{n+1,k-1}|.$$

Proof. Notice that $\mathbf{I}_n(\pi)$ is the set of inversion sequences of length n where each entry except 0 occurs at most $k-1$ times, and $L'_{n+1,k-1}$ is the set of label-increasing trees of $n+1$ vertices labeled $0, 1, \ldots, n$, with branching bounded by k except at the root. Then it is easy to see that the mapping sending $T \in L'_{n+1,k-1}$ to $e \in \mathbf{I}_n(\pi)$, where e_i is the parent of i in T, is a bijection between $L'_{n+1,k-1}$ and $\mathbf{I}_n(\pi)$. □

Theorem 22 implies $\mathbf{I}_n(0111)$ is in bijection with the label-increasing trees (of unbounded root degree) with branching bounded by 2, which is OEIS sequence A000772 [17].

More generally, it is well-established that $|L'_{n,k}|$ equals $D^n(\exp(x))$ evaluated at $x = 0$, where the operator D is defined by

$$D = \left(\sum_{j=0}^{k} \frac{x^j}{j!} \right) \frac{d}{dx}.$$

Therefore, we have the general formulae of exponential generating function

$$R_k(x) := \sum_{n \geqslant 0} |L'_{n,k}| \frac{x^n}{n!} = \exp\big(T_k(x) - 1\big),$$

where $T_k(x)$ is defined above in Eq. (2). When $k = 1$, the exponential generating function is $R_1(x) = \exp(\exp(x) - 1)$, whose coefficients yield OEIS sequence A000110 [17]. When $k = 2$, the exponential generating function is

$$R_2(x) = \exp\big(\tan(x) + \sec(x) - 1\big).$$

It is hard to explicitly write down $R_3(x)$, whose coefficients form OEIS sequence A094198.

Next, we present the following conjecture.

Conjecture 23. Let $A_n = |\mathbf{I}_n(0021)|$. We have that $A(x) = \sum_{n \geqslant 1} A_n x^n$ satisfies

$$\frac{1}{\big(1 - A(x)\big)\big(1 + A(x)\big)^2} = 1 - x.$$

In other words, $|\mathbf{I}_n(0021)|$ corresponds to the OEIS sequence A218225 [17]. This has been verified for all $n \leqslant 11$.

Remark 24. Since the writing of the original version of this paper, Conjecture 23 has been simultaneously proven by Chern, Fu, and Lin [5] and Mansour [11].

Acknowledgments

The research was conducted at the 2021 University of Minnesota Duluth REU (NSF–DMS Grant 1949884 and NSA Grant H98230-20-1-0009) and fully supported by the generosity of the CYAN Mathematics Undergraduate Activities Fund. The authors are deeply thankful to Professor Joseph Gallian for his long-lasting dedication in running the wonderful program. The authors are also grateful to Amanda Burcroff for her editing feedback. Lastly, we thank the anonymous referees who made many suggestions improving this article.

References

[1] Juan S. Auli and Sergi Elizalde. Consecutive patterns in inversion sequences. *Discrete Math. Theor. Comput. Sci.*, 21(2):6, 2019.

[2] Juan S. Auli and Sergi Elizalde. Consecutive patterns in inversion sequences II: Avoiding patterns of relations. *J. Integer Seq.*, 22(7):19.7.5, 2019.

[3] Juan S. Auli and Sergi Elizalde. Wilf equivalences between vincular patterns in inversion sequences. *Appl. Math. Comput.*, 388(C):125514, 2021.

[4] Shane Chern. On 0012-avoiding inversion sequences and a conjecture of Lin and Ma. *Quaest. Math.*, pages 1–14, 2022.

[5] Shane Chern, Shishuo Fu, and Zhicong Lin. Burstein's permutation conjecture, Hong and Li's inversion sequence conjecture, and restricted Eulerian distributions. arXiv:2209.12137, 2022.

[6] Sylvie Corteel, Megan A. Martinez, Carla D. Savage, and Michael Weselcouch. Patterns in inversion sequences I. *Discrete Math. Theor. Comput. Sci.*, 18(2):2, 2016.

[7] Alexander G. Kuznetsov, Igor M. Pak, and Alexander E. Postnikov. Increasing trees and alternating permutations. *Russ. Math. Surv.*, 49(6):79, 1994.

[8] Derrick H. Lehmer. Teaching combinatorial tricks to a computer. In *Combinatorial Analysis, Proc. Sympos. Appl. Math.*, volume 10, pages 179–193, 1960.

[9] Zhicong Lin and Chunyan Yan. Inversion sequences avoiding pairs of patterns. *Discrete Math. Theor. Comput. Sci.*, 22(1):23, 2020.

[10] Zhicong Lin and Sherry H. F. Yan. Vincular patterns in inversion sequences. *Appl. Math. Comput.*, 364: 124672, 2020.

[11] Toufik Mansour. Generating trees, 0021-avoiding inversion sequences, and a conjecture of Hong and Li. 2022.

[12] Toufik Mansour and Mark Shattuck. Pattern avoidance in inversion sequences. *Pure Math. Appl.*, 25:157–176, 2015.

[13] Megan A. Martinez and Carla D. Savage. Patterns in inversion sequences II: Inversion sequences avoiding triples of relations. *J. Integer Seq.*, 21(2):18.2.2, 2018.

[14] John Riordan. Forests of label-increasing trees. *J. Graph Theory*, 3(2):127–133, 1979.

[15] Carla D. Savage and Michael J. Schuster. Ehrhart series of lecture hall polytopes and Eulerian polynomials for inversion sequences. *J. Combin. Theory Ser. A*, 119(4):850–870, 2012.

[16] Rodica Simion and Frank W. Schmidt. Restricted permutations. *European J. Combin.*, 6(4):383–406, 1985.

[17] Neil J. A. Sloane et al. The On-line Encyclopedia of Integer Sequences. published electronically at `https://oeis.org`, 2021.

[18] Richard P. Stanley. A survey of alternating permutations. *Contemp. Math.*, 531:165–196, 2010.