On the Rearrangement Conjecture for Generalized Factor Order Over $\mathbb{P}$

Jay Pantone and Vincent Vatter*

Department of Mathematics, University of Florida, Gainesville, Florida, USA

Abstract. The Rearrangement Conjecture states that if two words over $\mathbb{P}$ are Wilf-equivalent in the factor order on $\mathbb{P}^*$ then they are rearrangements of each other. We introduce the notion of strong Wilf-equivalence and prove that if two words over $\mathbb{P}$ are strongly Wilf-equivalent then they are rearrangements of each other. We further conjecture that Wilf-equivalence implies strong Wilf-equivalence.

Résumé. La Conjecture de Réarrangements dit que si deux mots sont Wilf-équivalents dans l'ordre partiel de facteurs sur l'ensemble $\mathbb{P}^*$, alors l'un est un réarrangement de l'autre. Nous introduisons la notion d'équivalence forte de Wilf et démontrons que si deux mots sont équivalents dans ce sens, alors l'un est un réarrangement de l'autre. En plus, nous conjecturons que l'équivalence de Wilf, elle en fait implique l'équivalence forte de Wilf.

Keywords: composition, generalized factor order, Wilf-equivalence

1 Introduction

For ordinary words (finite sequences) u and w, we say that u is a *factor* of w if $w = w^{(1)} u w^{(2)}$ for possibly empty words $w^{(1)}$ and $w^{(2)}$. We are concerned with the *generalized factor order*, which extends the factor order to words over an arbitrary poset P. Given words $u, v \in P^*$, we say that v *dominates* u if they have the same length and $v_i \geq_P u_i$ for all i. We say that u is a *factor* of w if $w = w^{(1)} v w^{(2)}$ for a word v which dominates u, and in this case we write $w \geq_{\mathrm{gfo}} u$. For example, when $P = \mathbb{P}$, the positive integers, we have $1423314 \geq_{\mathrm{gfo}} 3123$ because 4233 dominates 3123. This paper is solely concerned with the case where $P = \mathbb{P}$.

The primary generating function we are interested in enumerates the set of all words in $\mathbb{P}^*$ (which could also be thought of as compositions) according to their length $|w|$, the sum of their entries $\|w\|$, and the number of factors dominating u they contain:

$$A_u(x, y, z) = \sum_{w \in \mathbb{P}^*} x^{|w|} y^{\|w\|} z^{\text{\# of factors dominating } u}.$$

The generalized factor order was introduced by Kitaev, Liese, Remmel, and Sagan [4], who defined the words $u, v \in \mathbb{P}^*$ to be *Wilf-equivalent* if $A_u(x, y, 0) = A_v(x, y, 0)$. They also made what has become known as the Rearrangement Conjecture. To state this conjecture, we need another definition. The words u and v of equal length are said to be *rearrangements* if u and v have the same multiset of values, i.e., if there exists a permutation π of the set $\{1, 2, \ldots, |u|\}$ such that $v = u_{\pi(1)} u_{\pi(2)} \cdots u_{\pi(|u|)}$.

The Rearrangement Conjecture *If two words in $\mathbb{P}^*$ are Wilf-equivalent then they are rearrangements of each other.*

Note that the converse to the Rearrangement Conjecture is false. Following the methods of [4], we can construct automata which recognize words avoiding a given factor. From this we find that

$$A_{122}(x, y, 0) = \frac{1 - 2y + (1 + x)y^2 - xy^3 + x^2 y^4}{1 - (2 + x)y + (1 + 2x)y^2 - (x + x^2)y^3 + x^2 y^4},$$

while

$$A_{212}(x, y, 0) = \frac{1 - 2y + (1 + x)y^2 - (x - x^2)y^3 + x^3 y^5}{(1 - y + x^2 y^3)(1 - (1 + x)y + xy^2 - x^2 y^3)}.$$

In particular, $[x^4 y^7]A_{122}(x, y, 0) = 13$, while $[x^4 y^7]A_{212}(x, y, 0) = 12$.

In [4], the authors also introduced a stronger notion of Wilf-equivalence. Given words $u, w \in \mathbb{P}^*$ we define $\mathrm{Em}(u, w)$ to consist of the set of indices of letters in w which begin a factor dominating u. The words u and v are said to be *super-strongly Wilf-equivalent* if there is a bijection $f : \mathbb{P}^* \to \mathbb{P}^*$ such that $|f(w)| = |w|$, $\|f(w)\| = \|w\|$, and $\mathrm{Em}(u, f(w)) = \mathrm{Em}(v, w)$ for all $w \in \mathbb{P}^*$. (Kitaev, Liese, Remmel, and Sagan had called this property "strong Wilf-equivalence", but we use that term to mean something different, in keeping with related literature.) This is such a stringent condition that while the words 2143 and 3412 are trivially Wilf-equivalent (because one is the reverse of the other), they are not super-strongly Wilf-equivalent (see [4]).

We focus on a condition which lies between these two; we say that u and v are *strongly Wilf-equivalent* if $A_u(x, y, z) = A_v(x, y, z)$. Note that super-strong Wilf-equivalence implies strong Wilf-equivalence, which in turn implies Wilf-equivalence. Our main result is the following.

Theorem 1.1 *If two words in $\mathbb{P}^*$ are strongly Wilf-equivalent then they are rearrangements of each other.*

Moreover, we conjecture that the following is true, which would imply the full Rearrangement Conjecture.

Conjecture 1.2 *If two words in $\mathbb{P}^*$ are Wilf-equivalent then they are strongly Wilf-equivalent.*

We have tested Conjecture 1.2 by finding all Wilf-equivalence classes for factors of weight up to 11 contained in words of weight up to 20 and verifying that each such class is also a strong Wilf-equivalence class (again, checking words of weight up to 20).

It is interesting that a related conjecture has recently been made for consecutive permutation patterns. Given a permutation β of length k and another permutation π of length n, π contains a consecutive occurrence of β if there is an index i such that the subsequence $\pi(i), \pi(i + 1), \ldots, \pi(i + k - 1)$ is in the same relative order as β. The analogue of the A_u generating function in this context is then

$$A_\beta(x, z) = \sum_{\text{permutations } \pi} x^{|\pi|} z^{\text{\# of occurrences of } \beta \text{ in } \pi}.$$

Fig. 1: A 2-cluster and a 3-cluster of 3123 which share the same underlying element of $\mathbb{P}^*$.

The permutations β and γ are said to be *c-Wilf-equivalent* (the "c" is to denote that this is for consecutive pattern containment) if $A_\beta(x, 0) = A_\gamma(x, 0)$, and *strongly c-Wilf-equivalent* if $A_\beta(x, z) = A_\gamma(x, z)$. Nakamura [5, Conjecture 5.6] has conjectured that c-Wilf-equivalence implies strong c-Wilf-equivalence.

2 The Cluster Method

The easiest way to compute the generating function $A_u(x, y, z)$ for a given word u is probably to construct an automaton, as detailed in [4]. However, to prove Theorem 1.1, we instead use the *cluster method*. This method is originally due to Goulden and Jackson [3]. The authors owe their knowledge of the method to the recent work of Elizalde and Noy [2] and Elizalde [1] on consecutive patterns in permutations.

Given a word $u \in \mathbb{P}^*$ of length k (which will be the forbidden factor), an *m-cluster* of u is a word $c \in \mathbb{P}^*$ consisting of m overlapping factors of length k containing u, which are marked. The overlapping condition requires that when labeled from left to right, each pair of consecutive factors must share at least one entry. Additionally, the first and last letters of c must be contained in marked factors.

Note that an m- and m'-cluster can share the same underlying word. For example, Figure 1 shows a 2- and a 3-cluster of $u = 3123$ defined on the same word in $\mathbb{P}^*$.

The cluster generating function for u is defined as

$$C_u(x, y, z) = \sum_{m \geq 1} z^m \sum_{\substack{m\text{-clusters} \\ c \text{ of } u}} x^{|c|} y^{\|c\|}.$$

Now fix a forbidden factor u. We may view an arbitrary word w (which may or may not contain u as a factor) as a sequence of letters and clusters of u. The generating function for an arbitrary letter of $\mathbb{P}$ is

$$\frac{xy}{1 - y} = xy + xy^2 + \cdots,$$

while the generating function for clusters is $C_u(x, y, z)$. We first claim that

$$A_u(x, y, 0) = \frac{1}{1 - \frac{xy}{1-y} - C_u(x, y, -1)}. \tag{\dagger}$$

To see why ($\dagger$) is true, consider the expansion of the right-hand side, which includes terms for every word $w \in \mathbb{P}^*$. If the word w has s factors which dominate u, then each such factor may or may not be chosen to participate in a cluster. Thus w corresponds to 2^s terms in the expansion of the right-hand side of ($\dagger$), half with positive signs and half with negative signs. Therefore what remains after cancellation is indeed $A_u(x, y, 0)$. In particular, if u and v are Wilf-equivalent, then $C_u(x, y, -1) = C_v(x, y, -1)$.

To investigate strong Wilf-equivalence, we need to use what Wilf [6, Section 4.2] calls the Sieve Method. Consider the generating function

$$\frac{1}{1 - \frac{xy}{1-y} - C_u(x, y, z)}.$$

If the word w contains s factors which dominate u then the same analysis as above shows that the contribution of w to the expansion of the above generating function is

$$x^{|w|}y^{\|w\|}\left(\binom{s}{0}z^0 + \binom{s}{1}z^1 + \cdots + \binom{s}{s}z^s\right) = x^{|w|}y^{\|w\|}(z+1)^s.$$

We therefore deduce by substitution that

$$A_u(x,y,z) = \frac{1}{1 - \frac{xy}{1-y} - C_u(x,y,z-1)}.$$

Thus the words u and v are strongly Wilf-equivalent if and only if $C_u(x,y,z) = C_v(x,y,z)$. In light of this, we may view Conjecture 1.2 as stating that if $C_u(x,y,-1) = C_v(x,y,-1)$ then $C_u(x,y,z) = C_u(x,y,z)$.

In the context of the factor order on $\mathbb{P}^*$, increasing the entries of an m-cluster still leaves an m-cluster. Therefore we say that the m-cluster c for u is *minimal* if none of its entries can be decreased without destroying a marked factor. (Note that the clusters in Figure 1 are *not* minimal.) We define the generating function for these clusters to be

$$M_u(x,y,z) = \sum_{m \geq 1} z^m \sum_{\substack{\text{minimal} \\ m\text{-clusters} \\ c \text{ of } u}} x^{|c|}y^{\|c\|},$$

and it is easy to see that

$$C_u(x,y,z) = M_u\left(\frac{x}{1-y}, y, z\right).$$

Therefore the words u and v are Wilf-equivalent if and only if $M_u(x,y,-1) = M_v(x,y,-1)$ and they are strongly Wilf-equivalent if and only if $M_u(x,y,z) = M_u(x,y,z)$. Note that $|u|$ is the smallest exponent of x in $M_u(x,y,-1)$, so if two words are Wilf-equivalent they must have the same length. Similarly, $\|u\|$ is the smallest exponent of y in $M_u(x,y,-1)$, so if u and v are Wilf-equivalent, $\|u\| = \|v\|$.

We can now outline our proof of Theorem 1.1. By sorting the entries of a word $u \in \mathbb{P}^*$ we obtain a partition $\lambda = (\lambda_1 \geq \cdots \geq \lambda_k)$ of $\|u\|$. We prove our theorem by showing that one can compute λ by examining the coefficients of $M_u(x,y,z)$. This proves the theorem because it means that if the words u and v are strongly Wilf-equivalent, then they are rearrangements of the same partition. In the next Section, we illustrate our proof by presenting the special case of a forbidden factor of length 4. The proof of the general case is presented in the section after that. We conclude by showing how the cluster method can be applied to derive short proofs of a result and two conjectures of Kitaev, Liese, Remmel, and Sagan [4].

3 The Case $k = 4$

We begin by establishing terminology which will be used in the proof of the general case. A minimal m-cluster of u is built by aligning m overlapping (but not necessarily mutually overlapping) copies of u in an array and then taking the maximum of each column. For example, below are all minimal 2-clusters of a word u of length $k = 4$, where for a set $I \subseteq \{1,2,3,4\}$ we use the notation $u_I = \max\{u_i \ : \ i \in I\}$.

u_1	u_2	u_3	u_4				u_1	u_2	u_3	u_4				u_1	u_2	u_3	u_4			
	u_1	u_2	u_3	u_4					u_1	u_2	u_3	u_4					u_1	u_2	u_3	u_4
u_1	$u_{1,2}$	$u_{2,3}$	$u_{3,4}$	u_4			u_1	u_2	$u_{1,3}$	$u_{2,4}$	u_3	u_4		u_1	u_2	u_3	$u_{1,4}$	u_2	u_3	u_4

Using bracket notation for coefficient extraction, we see that for a word u of length 4,

$$[z^2]M_u = x^5 y^{u_1+u_{1,2}+u_{2,3}+u_{3,4}+u_4} + x^6 y^{u_1+u_2+u_{1,3}+u_{2,4}+u_3+u_4} + x^7 y^{u_1+u_2+u_3+u_{1,4}+u_2+u_3+u_4}.$$

Of course we don't know what the u_i are, but we can examine the terms of M_u to determine these sums.

There will typically be more than one minimal m-cluster of each length for $m \geq 3$. We summarize the information on 2-clusters in the following chart, where the number corresponding to length m and entry u_I represents the number of occurrences of u_I among all 2-clusters of length m.

length	u_1	u_2	u_3	u_4	$u_{1,2}$	$u_{1,3}$	$u_{1,4}$	$u_{2,3}$	$u_{2,4}$	$u_{3,4}$
5	1			1	1			1		1
6	1	1	1	1		1			1	
7	1	2	2	1			1			

While we selected $m = 2$ in order to explain our approach, this data does not lead to any conclusions. Instead, we need to consider the $m = 3$ and $m = 4$ cases, for which we only present the charts. For $m = 3$, we have the following.

length	u_1	u_2	u_3	u_4	$u_{1,2}$	$u_{1,3}$	$u_{1,4}$	$u_{2,3}$	$u_{2,4}$	$u_{3,4}$	$u_{1,2,3}$	$u_{1,2,4}$	$u_{1,3,4}$	$u_{2,3,4}$	$u_{1,2,3,4}$
6	1			1	1					1	1			1	
7	2	1	1	2	1	1		2	1	1		1	1		
8	3	3	3	3	2	2	2	2	2	2					
9	2	4	4	2		2	2		2						
10	1	3	3	1			2								

Because every minimal cluster has precisely one letter equal to u_1 (its first), we can see from the above table that there are nine 3-clusters of a pattern of length 4. The data for $m = 4$ is displayed below.

length	u_1	u_2	u_3	u_4	$u_{1,2}$	$u_{1,3}$	$u_{1,4}$	$u_{2,3}$	$u_{2,4}$	$u_{3,4}$	$u_{1,2,3}$	$u_{1,2,4}$	$u_{1,3,4}$	$u_{2,3,4}$	$u_{1,2,3,4}$
7	1			1	1					1	1			1	1
8	3	1	1	3	2	1		2	1	2	2	2	2	2	
9	6	4	4	6	5	4	3	5	4	5	2	2	2	2	
10	7	9	9	7	4	7	6	6	7	4		2	2		
11	6	12	12	6	3	6	9	3	6	3					
12	3	9	9	3		3	6		3						
13	1	4	4	1			3								

We notice from the tables above that the shortest minimal 3-cluster and the shortest minimal 4-cluster, which have lengths 6 and 7 respectively, are almost identical except for the presence of $u_{1,2,3,4}$ in the (exponent of y corresponding to the relevant) 4-cluster. In terms of generating functions, this means that

$$\frac{d}{dy}\left(\left([x^7 z^4] - [x^6 z^3]\right) M_u \right)\Bigg|_{y=1} = u_{1,2,3,4}.$$

Of course, $u_{1,2,3,4}$ is the largest entry of u, so we have just computed λ_1 from M_u. Comparing the length 8, $m = 4$ data with the lengths 6 and 7, $m = 3$ data, we see that

$$\frac{d}{dy}\left(\left([x^8 z^4] - [x^7 z^3] - [x^6 z^3]\right) M_u \right)\Bigg|_{y=1} = u_{1,2,3} + u_{1,2,4} + u_{1,3,4} + u_{2,3,4}.$$

Of these four terms, three are equal to the greatest entry of u, while one is equal to the second greatest entry. Therefore

$$\frac{d}{dy}\left(\left([x^8 z^4] - [x^7 z^3] - [x^6 z^3]\right) M_u\right)\bigg|_{y=1} = 3\lambda_1 + \lambda_2.$$

As we have previously determined λ_1, this allows us to compute λ_2. Next we compare the length 9, $m = 4$ data to the lengths 6, 7, and 8, $m = 3$ data to see that

$$\frac{d}{dy}\left(\left([x^9 z^4] - [x^8 z^3] - [x^7 z^3] - [x^6 z^3]\right) M_u\right)\bigg|_{y=1} =$$

$$u_{1,2} + u_{1,3} + u_{1,4} + u_{2,3} + u_{2,4} + u_{3,4} + u_{1,2,3} + u_{1,2,4} + u_{1,3,4} + u_{2,3,4}.$$

We then see that

$$u_{1,2} + u_{1,3} + u_{1,4} + u_{2,3} + u_{2,4} + u_{3,4} = 3\lambda_1 + 2\lambda_2 + \lambda_3,$$

and thus, by our observation above about $u_{1,2,3} + u_{1,2,4} + u_{1,3,4} + u_{2,3,4}$,

$$\frac{d}{dy}\left(\left([x^9 z^4] - [x^8 z^3] - [x^7 z^3] - [x^6 z^3]\right) M_u\right)\bigg|_{y=1} = 6\lambda_1 + 3\lambda_2 + \lambda_3,$$

enabling us to compute λ_3. It remains only to compute λ_4, but given that we know λ_1, λ_2, and λ_3, we can compute λ_4 by looking at the smallest exponent of y in M_u, which is equal to $\|u\| = \lambda_1 + \lambda_2 + \lambda_3 + \lambda_4$.

4   The General Case

Now suppose that the word u has arbitrary length k. We aim to show that for every $1 \leq i \leq k - 1$, the quantity

$$\frac{d}{dy}\left(\left([x^{2k+i-2} z^k] - [x^{2k+i-3} z^{k-1}] - \cdots - [x^{2k-2} z^{k-1}]\right) M_u\right)\bigg|_{y=1} \tag{‡}$$

is a linear combination (with strictly positive coefficients) of $\lambda_1, \ldots, \lambda_i$. This will allow us to compute $\lambda_1, \ldots, \lambda_{k-1}$. We then compute λ_k by examining the smallest exponent of y in M_u which is equal to $\|u\| = \lambda_1 + \cdots + \lambda_k$. We must first refine our terminology.

As we began the previous section, note that minimal m-clusters of u are obtained by aligning m overlapping copies of u in an array and then taking the maximum of each column. For example, below we show a 5-cluster of a word u of length 5.

u_1	u_2	u_3	u_4	u_5					
	u_1	u_2	u_3	u_4	u_5				
		u_1	u_2	u_3	u_4	u_5			
			u_1	u_2	u_3	u_4	u_5		
				u_1	u_2	u_3	u_4	u_5	
u_1	$u_{1,2}$	$u_{2,3}$	$u_{1,3,4}$	$u_{1,2,4,5}$	$u_{2,3,5}$	$u_{1,3,4}$	$u_{2,4,5}$	$u_{3,5}$	u_4	u_5

We refer to the word below the line as a *symbolic m-cluster*, owing to the fact that we think about it throughout our proof as a word over the letters u_I for $I \subseteq \{1, \ldots, k\}$. We call the array above the line which produced the symbolic cluster an *m-pre-cluster*. Thus every entry of a symbolic cluster comes from

a column of its associated pre-cluster. Note that the presence of a given column in a pre-cluster uniquely determines the contents and relative position of the rows which have nonempty entries in that column.

In particular, in the expansion of (‡), each term corresponds to a column of a k- or $(k-1)$-pre-cluster. We say that the *height* of a column of a pre-cluster is the number of nonempty entries it contains. A *top column* of a pre-cluster is one that includes a nonempty entry of the top row. Similarly, a *bottom column* is one that touches the bottom row. (It is possible for a column to be both top and bottom.) Finally, a *middle column* is one that touches neither the top nor bottom row of the pre-cluster.

We prove our claim about (‡), and thus Theorem 1.1, with a series of five propositions.

Proposition 4.1 *In the expansion of* (‡)*, none of the remaining terms after cancellation correspond to top or bottom columns of height $k - 2$ or less.*

Proof: We prove the proposition by constructing a bijection Φ which maps a k-pre-cluster of length $2k + i - 2$ with a single marked non-middle column of height at most $k - 2$ to a shorter $(k-1)$-pre-cluster with the same marked non-middle column. Let C be a k-pre-cluster of length $2k + i - 2$ with marked non-middle column c. If c is a top column, then $\Phi(C, c)$ is obtained by deleting the bottom row of C while leaving c marked. Because the height of c is at most $k - 2$, it cannot be both a top and a bottom column in C, and thus c will be unchanged in $\Phi(C, c)$. If c is a bottom column, $\Phi(C, c)$ is obtained similarly by deleting the top row of C while again leaving c marked.

The map is inverted as follows. Let C' be a $(k-1)$-pre-cluster of length at most $2k + i - 3$ and with marked non-middle column c' of height at most $k - 2$. Again, by the height restriction on c', it cannot be both a top and a bottom column. If c' is a top column, then there is a unique k-pre-cluster of length $2k + i - 2$ which has c' as a top column, which is obtained by adding a new bottom row to C' in the unique position to obtain a k-pre-cluster of length $2k + i - 2$. The resulting k-pre-cluster C (with marked column c corresponding to c') clearly has the property that $\Phi(C, c) = (C', c')$. Analogously, if c' is a bottom column, then there is exactly one way to build a k-pre-cluster C with marked bottom column c' such that $\Phi(C, c) = (C', c')$. □

Our next step is to study the contribution of columns of height k or $k - 1$ to (‡). There are two cases, $i = 1$ and $i \geq 2$. We first consider the special case $i = 1$. In this case we see that there is a unique k-pre-cluster of length $2k - 1$:

$$
\begin{array}{ccccccc}
u_1 & \cdots & u_{k-1} & u_k & & & \\
& \cdots & u_{k-2} & u_{k-1} & u_k & & \\
& & \vdots & \vdots & \vdots & & \\
& & u_1 & u_2 & u_3 & \cdots & \\
& & & u_1 & u_2 & u_3 & \cdots & u_k
\end{array}
$$

The columns of height k or $k - 1$ in this pre-cluster correspond to the entries $u_{1,\ldots,k}$, $u_{1,\ldots,k-1}$ and $u_{2,\ldots,k}$ in the associated symbolic cluster. There is also a unique $(k-1)$-pre-cluster of length $2k - 2$:

$$
\begin{array}{ccccccc}
u_1 & \cdots & u_{k-2} & u_{k-1} & u_k & & \\
& \cdots & u_{k-1} & u_{k-2} & u_{k-1} & u_k & \\
& & \vdots & \vdots & \vdots & & \\
& & u_1 & u_2 & u_3 & \cdots & u_k
\end{array}
$$

This pre-cluster contains two columns of height $k - 1$, corresponding to the entries $u_{1,\ldots,k-1}$ and $u_{2,\ldots,k}$ in the associated symbolic cluster. These two entries cancel with two of the three entries from the symbolic

k-cluster. Because all other terms of (‡) cancel by Proposition 4.1 in this case, (‡) reduces to $u_{1,\ldots,k} = \lambda_1$ when $i = 1$, as desired. Notice that when $i = 1$ all columns which contribute to (‡) are top or bottom columns, and thus we are completely done with this case and will not consider it again in our proof.

Our next propositions give the $i \geq 2$ case.

Proposition 4.2 *In the expansion of* (‡) *for $i \geq 2$, the total contribution of top and bottom columns is* $(k-1)\lambda_1 + \lambda_2$.

Proof: We begin by looking at such columns in $(k-1)$-pre-clusters. Trivially, these pre-clusters cannot contain columns of height k. If a $(k-1)$-pre-cluster contains a column of height $k-1$ then it can have length at most $2k-1$. As in the $i = 1$ case, there is a unique $(k-1)$-pre-cluster of length $2k-2$ with a column of height $k-1$. This pre-cluster actually has two columns of height $k-1$, corresponding to the terms $u_{1,\ldots,k-1}$ and $u_{2,\ldots,k}$. Additionally, for every other subset $I \subseteq \{1,\ldots k\}$ of size $k-1$, there is a unique $(k-1)$-pre-cluster of length $2k-1$ which has a column corresponding to u_I.

Next we consider columns of height at least $k-1$ in k-pre-clusters. There is a unique k-pre-cluster with a column of height k, but this pre-cluster has length $2k-1$ so will not contribute to (‡) when $i \geq 2$. Now choose any subset $I \subseteq \{1,\ldots,k\}$ of size $k-1$. There are precisely two k-pre-clusters with a marked top or bottom column u_I of each length between $2k$ and $3k-3$ (i.e., for i between 2 and $k-1$), because u_I can arise from either a top column or a bottom column in such pre-clusters.

Therefore the total contribution of these columns is

$$\sum_{\substack{I \subseteq \{1,\ldots,k\}, \\ |I| = k-1}} u_I = (k-1)\lambda_1 + \lambda_2,$$

as desired. □

In the $i = 2$ case, the terms of (‡) correspond to k-pre-clusters of length $2k$ and $(k-1)$-pre-clusters of lengths $2k-1$ and $2k-2$. As no column of such a pre-cluster can be a middle column, our claim follows from Proposition 4.2 in this case. We move on to consider the contribution of middle columns when $i \geq 3$.

Proposition 4.3 *For all $i \geq 3$, middle columns of height less than $k-i+1$ do not contribute to the expansion of* (‡).

Proof: Let C be a k-pre-cluster of minimal length which contains a middle column c of height $k-i$. Since c is a middle column, C must have rows both above and below it. Together, these two rows and the column c contribute at least $2k+1$ to the length of C. Moreover, C has $(k-(k-i+2)) = i-2$ more rows. Because each such row increases the length of C by at least 1, C has total length at least $2k+i-1$. The expression (‡) does not involve any k-clusters of this length. Obviously, the same argument holds for shorter middle columns as well (those of length less than $k-i$).

Analogously, any $(k-1)$-pre-cluster which contains a middle column of height at most $k-i$ has length at least $2k+i-2$. Again, the expression (‡) does not involve any $(k-1)$-clusters of this length. □

Proposition 4.4 *Fix $3 \leq i \leq k-1$ and a column c of height at least $k-i+1$. There is a bijection between k-pre-clusters of length $2k+i-2$ which contain c as a marked middle column and $(k-1)$-pre-clusters of length between $2k-2$ and $2k+i-3$ which contain c as a marked middle column.*

Proof: Let h denote the height of c. To construct our bijection we first consider the k-pre-clusters. Let c and c' be columns of the same height $h \geq k - i + 1$. We define a bijection, $\Psi_c^{c'}$, which maps from a k-pre-cluster of length $2k + i - 2$ with a marked middle column c to a k-pre-cluster of the same length with marked middle column c'. Showing that $\Psi_c^{c'}$ is indeed a bijection will obviously prove our claim about middle columns in k-pre-clusters.

Given such a k-pre-cluster C with marked middle column c, the map $\Psi_c^{c'}(C, c)$ is defined by replacing the rows which involve c in C by the rows which involve c' (which are the same in any pre-cluster it is a column of) in such a way that c' (which becomes our new marked column) is in the same column of $\Psi_c^{c'}(C, c)$ as c was in C. Note that, because c was a middle column, the length of $\Psi_c^{c'}(C, c)$ will be the same as the length of C.

However, we need to show that $\Psi_c^{c'}(C, c)$ is indeed a pre-cluster. The concern we need to address is that it is a priori possible for the rows of $\Psi_c^{c'}(C, c)$ to fail to overlap. As we will show, however, this is prevented by our length condition. Consider, for the moment, only the rows which involve c'. The least entry that can lie on the top of the column c' is u_h, and thus this set of rows contains at least $h - 1$ columns to the left of c'. Similarly, the greatest entry than can lie on the bottom of c' is u_{k-h+1}, so this set of rows contains at least $k - (k - h + 1) = h - 1$ columns to the right of c'. Were the rows of $\Psi_c^{c'}(C, c)$ to fail to overlap, there would be one row completely to the left (and above) the rest of the array, or a row completely to the right (and below) the array. Let us suppose for the sake of contradiction that there is a row of $\Psi_c^{c'}(C, c)$ completely to the left of the rest of the array. This row therefore contributes its total length, k, to the length of $\Psi_c^{c'}(C, c)$. Moreover, there must be at least $h - 1$ columns separating it from the column c' in $\Psi_c^{c'}(C, c)$. Thus the column c' contributes 1 to the length of $\Psi_c^{c'}(C, c)$, and because c' is a middle column in $\Psi_c^{c'}(C, c)$, there must be a row which contributes its total length to the length of $\Psi_c^{c'}(C, c)$. Finally, we have only accounted for $h + 2$ rows, and each additional row must contribute at least 1 to the length of $\Psi_c^{c'}(C, c)$. This shows that the length of $\Psi_c^{c'}(C, c)$ is at least

$$k + h - 1 + 1 + k + (k - (h + 2)) = 3k - 2.$$

However, the length of $\Psi_c^{c'}(C, c)$ is $2k + i - 2$ (the length of C), which is at most $3k - 3$ because $i \leq k - 1$.

This contradiction completes the proof of the proposition for k-pre-clusters. The case of $(k - 1)$-pre-clusters is completely analogous. $\square$

Our claim about ($\ddagger$) in the $i \geq 3$ case, and thus also our proof of Theorem 1.1, follows from the following result.

Proposition 4.5 *For $3 \leq i \leq k - 1$, the contribution of middle columns to the expansion of ($\ddagger$) is a linear combination of $\lambda_1, \ldots, \lambda_i$ with strictly positive coefficients.*

Proof: Consider the contribution of a particular middle column c to the expansion of ($\ddagger$). The k-pre-clusters of length $2k + i - 2$ with this middle column will contribute positively, while the $(k - 1)$-pre-clusters of lengths between $2k - 2$ and $2k + i - 3$ contribute negatively.

Suppose that (C, c) is a $(k - 1)$-pre-cluster of length between $2k - 2$ and $2k + i - 3$ with marked middle column c. Because c is a middle column of C, C must contain at least k columns both to the left and the right of c. Therefore by adding a new bottom column in the appropriate position we can create a k-pre-cluster of length $2k + i - 2$ containing marked middle column c. This mapping defines an injection from $(k-1)$-pre-clusters of lengths between $2k - 2$ and $2k + i - 3$ and k-pre-clusters of length $2k + i - 2$, both with marked middle column c. The corresponding terms of ($\ddagger$) therefore cancel.

However, there are also terms of (‡) corresponding to middle columns of k-pre-clusters where the middle column begins in the second row from the bottom of the k-pre-cluster, and these terms are not canceled in (‡). Therefore each middle column c of height between $k - i + 1$ and $k - 2$ contributes positively to (‡), and this contribution is the same for each column c by the previous proposition. Finally, note that the contribution of one copy of every column of height h for $k - i + 1 \leq h \leq k - 2$ is

$$\sum_{\substack{I \subseteq \{1,\dots,k\}, \\ |I|=h}} u_I = \binom{k-1}{h-1}\lambda_1 + \binom{k-2}{h-1}\lambda_2 + \cdots + \binom{h-1}{h-1}\lambda_{k-h+1},$$

and thus the contribution of these columns to (‡) is a linear combination of $\lambda_1, \dots, \lambda_i$ with positive coefficients, completing the proof.   □

5 Further applications of the cluster method

In our final section, we provide short new proofs using the cluster method for a result and a two conjectures of Kitaev, Liese, Remmel, and Sagan [4]. For our first result we need to introduce some new notation. Given $u \in \mathbb{P}^*$, we write $1u$ for the word obtained by prepending the letter 1 to u, and we write u^+ for the word obtained by adding 1 to every letter of u.

Proposition 5.1 ([4, Lemma 4.1]) *For $u, v \in \mathbb{P}^*$, we have the following Wilf-equivalences,*

(a) u is Wilf-equivalent to its reverse, u^r,

(b) u and v are Wilf-equivalent if and only if $1u$ and $1v$ are Wilf-equivalent,

(c) if u and v are Wilf-equivalent then u^+ and v^+ are Wilf-equivalent.

In fact, the proposition also holds with Wilf-equivalence replaced by strong Wilf-equivalence.

Proof: For part (a), note that the reversal of a minimal m-cluster of u is a minimal m cluster of u^r with the same length and sum of its entries. Thus $M_u(x, y, z) = M_{u^r}(x, y, z)$, so u is strongly Wilf-equivalent to its reverse (so the two words are trivially Wilf-equivalent).

Next, we consider how to turn minimal m-clusters of u into minimal m-clusters of $1u$. Suppose that we have d minimal clusters of u, $c^{(1)}, \dots, c^{(d)}$, where $c^{(i)}$ is an m_i-cluster of u. Then the concatenation $1c^{(1)} \cdots c^{(d)}$ is a minimal $(m_1 + \cdots + m_d)$-cluster of $1u$. Moreover, because every minimal cluster of $1u$ is of this form, we see that

$$M_{1u}(x, y, z) = \frac{xyM_u(x, y, z)}{1 - M_u(x, y, z)}.$$

Therefore u and v are strongly Wilf-equivalent (resp., Wilf-equivalent) if and only if $1u$ and $1v$ are strongly Wilf-equivalent (resp., Wilf-equivalent).

Finally, part (c) follows immediately from the observation that if c is a minimal m-cluster of u then c^+ is a minimal m-cluster of u^+, so $M_{u^+}(x, y, z) = M_u(xy, y, z)$.   □

Next we establish a conjecture of [4] proving the converse of Proposition 5.1 (c).

Proposition 5.2 ([4, Item (3) of Subsection 8.4]) *If u^+ and v^+ are Wilf-equivalent (resp., strongly Wilf-equivalent) then u and v are Wilf-equivalent (resp., strongly Wilf-equivalent).*

Proof: From the proof of part (c) of the previous proposition, we see that

$$M_u(x, y, z) = M_{u^+}\left(\frac{x}{y}, y, z\right).$$

Therefore if $M_{u^+}(x, y, -1) = M_{v^+}(x, y, -1)$ then $M_u(x, y, -1) = M_v(x, y, -1)$, and of course the same holds in the context of strong Wilf-equivalence (when -1 is not substituted for z). $\qquad\square$

In order to explain many of the Wilf-equivalences that were found in [4], the authors made the following conjecture.

Theorem 5.3 ([4, Conjecture 8.3]) *The words $a1b2c$ and $a2b1c$ are (strongly) Wilf-equivalent for any choice of positive integers $a, b, c \geq 2$.*

Proof: We prove the result by constructing a bijection Π between m-pre-clusters of $a1b2c$ and m-pre-clusters of $c1b2a$, which preserves the length and sum of entries of the corresponding minimal m-clusters. This will prove the result because it shows that

$$M_{a1b2c}(x, y, z) = M_{c1b2a}(x, y, z),$$

and by Proposition 5.1 (a), $c1b2a$ is strongly Wilf-equivalent to $a2b1c$.

Let C be an m-pre-cluster of $a1b2c$. To construct $\Pi(C)$ we simply replace every row of C by the word $c1b2a$ (without moving the rows). For example, we have the mapping of 3-pre-clusters

a	1	b	2	c					c	1	b	2	a		
		a	1	b	2	c					c	1	b	2	a
				a	1	b	2	c					c	1	b	2	a
a	1	$\max\{a,b\}$	a	$\max\{b,c\}$	b	c	c		c	1	$\max\{b,c\}$	c	$\max\{a,b\}$	b	a	a

$\longrightarrow$

under Π. Clearly Π is a bijection, but we must show that the cluster corresponding to $\Pi(C)$ has the same length and sum of its entries as the cluster corresponding to C. Instead, we prove the stronger claim that these two clusters are rearrangements of each other.

Since Π essentially swaps the locations of each a with the locations of each c, it is clear that any entry which does not involve a or c or which involves both a and c occurs equally frequently in the clusters corresponding to C and $\Pi(C)$. To complete the proof, we show that in any cluster of $a1b2c$, the entries a and c occur equally often and the entries $\max\{a, b\}$ and $\max\{b, c\}$ occur equally often, from which it follows immediately that the clusters corresponding to C and $\Pi(C)$ are indeed rearrangements of each other.

We prove this by induction, for which the base case (a 1-pre-cluster of $a1b2c$) is trivial. Let C be an $a1b2c$ pre-cluster such that the cluster corresponding to C has an equal number of a and c entries and an equal number of $\max\{a, b\}$ and $\max\{b, c\}$ entries. Consider the effect of adding another row to the bottom of C. There are fifteen possible alignments of the columns involving this new row; in each case the new cluster resulting from adding a row to C preserves equality of the number of occurrences of a and c entries and of $\max\{a, b\}$ and $\max\{b, c\}$ entries. The fifteen total cases (in which we only show the rightmost five columns) are shown below.

c a 1 b 2 c	2 c a 1 b 2 c	b 2 c a 1 b 2 c	1 b 2 c a 1 b 2 c	c 2 c a 1 b 2 c
c b 2 c a 1 b 2 c	c 1 b 2 c a 1 b 2 c	2 c b 2 c a 1 b 2 c	2 c 1 b 2 c a 1 b 2 c	b 2 c 1 b 2 c a 1 b 2 c
c 2 c b 2 c a 1 b 2 c	c 2 c 1 b 2 c a 1 b 2 c	c b 2 c 1 b 2 c a 1 b 2 c	2 c b 2 c 1 b 2 c a 1 b 2 c	c 2 c b 2 c 1 b 2 c a 1 b 2 c

For example, in the pre-cluster in the top-right corner, adding the row $a1b2c$ destroys one c entry from the corresponding cluster by turning it into a $\max\{a, c\}$ entry, preserves one c entry, and creates a new c entry. It does not create or destroy any occurrences of a, $\max\{a, b\}$, or $\max\{b, c\}$.

In the pre-cluster in the bottom-right corner, adding the last row destroys an occurrence of c, preserves an occurrence of c, and creates an occurrence of c. Similarly, it both creates and destroys one occurrence of $\max\{b, c\}$, while preserving one other existing occurrence of $\max\{b, c\}$.

Since we have shown that the clusters corresponding to C and $\Pi(C)$ are rearrangements of each other, this completes the proof.   □

Finally, note that our proof of Theorem 5.3 extends to prove a more general result. For positive integers a, b, c, x, y with $a, b, c \geq x, y$, the same proof shows that $axbyz$ is Wilf-equivalent to $aybxc$.

References

[1] S. Elizalde. The most and the least avoided consecutive patterns. *Proc. Lond. Math. Soc.*, 106(5):957–979, 2013.

[2] S. Elizalde and M. Noy. Clusters, generating functions and asymptotics for consecutive patterns in permutations. *Adv. in Appl. Math.*, 49(3-5):351–374, 2012.

[3] I. P. Goulden and D. M. Jackson. *Combinatorial Enumeration*. A Wiley-Interscience Publication. John Wiley & Sons Inc., New York, 1983.

[4] S. Kitaev, J. Liese, J. Remmel, and B. E. Sagan. Rationality, irrationality, and Wilf equivalence in generalized factor order. *Electron. J. Combin.*, 16(2):Paper 22, 26, 2009.

[5] B. Nakamura. Computational approaches to consecutive pattern avoidance in permutations. *Pure Math. Appl. (PU.M.A.)*, 22(2):253–268, 2011.

[6] H. S. Wilf. *generatingfunctionology*. A K Peters Ltd., Wellesley, MA, third edition, 2006.