

k -pop stack sortable permutations and 2-avoidance

Murray Elder* Yoong Kuan Goh†

School of Mathematical and Physical Sciences
University of Technology Sydney
Ultimo NSW 2007, Australia

{murrayelder, gohyoongkuan}@gmail.com

Submitted: May 24, 2020; Accepted: Mar 7, 2021; Published: Mar 26, 2021

© M. Elder and Y. K. Goh. Released under the CC BY-ND license (International 4.0).

Abstract

We consider permutations sortable by k passes through a deterministic pop stack. We show that for any $k \in \mathbb{N}$ the set is characterised by finitely many patterns, answering a question of Claesson and Guðmundsson. Moreover, these sets of patterns are algorithmically constructible.

Our characterisation demands a more precise definition than in previous literature of what it means for a permutation to avoid a set of barred and unbarred patterns. We propose a new notion called *2-avoidance*.

Mathematics Subject Classifications: 05A05

1 Introduction

A *pop stack* is a sorting device which operates as follows: at each step it can either push one entry from the input stream onto the top of the stack, or else pop the *entire stack contents* to the output stream. We consider the entries to be distinct real numbers with the usual ordering. A *deterministic* pop stack always performs the push move unless the entry on the top of the stack is smaller in value than the entry to be pushed from the input, or if there is no further input. See for example Figure 1. For convenience from now on we assume entries are integers and write a sequence 2, 1, 3 as 213 when entries are single digits.

Observe that by definition the stack remains ordered from smallest on top to largest on the bottom during the operation of a deterministic pop stack.

A permutation (ordered sequence of distinct real numbers) can be sorted by k passes through a deterministic pop stack if after repeating the procedure k times, the sequence

*Both authors supported by Australian Research Council grant DP160100486.

†Supported by the Australian Government Research Training Program.






Figure 1: Sorting 213 using a deterministic pop stack.

is ordered from smallest to largest. For example, 41352 can be sorted by two passes (Figure 2).



(a) First pass.



(b) Second pass.

Figure 2: Sorting 41352 with a 2-pass pop stack.

Let $p_1(\alpha)$ denote the sequence obtained by passing a sequence α through a deterministic pop stack once, and define $p_k(\alpha) = p_{k-1}(p_1(\alpha))$. For example $p_1(41352) = 14325$ and $p_2(41352) = 12345$. We say α is k -pass deterministic pop stack sortable, or k -pop stack sortable for short, if $p_k(\alpha)$ is an increasing sequence.

In 1981 Avis and Newborn characterised permutations sorted by a single pass through a pop stack [1], initiating the study of pop stack sorting. Specifically they showed a permutation can be sorted in one pass if and only if it avoids 231 and 312 in the usual sense of pattern avoidance (see Section 2 for precise definitions). Pudwell and Smith characterised permutations sorted by 2 passes, in terms of avoiding a set of six usual patterns and two special *barred* patterns (defined in Subsection 2.1 below), and computed a rational generating function for the number of such permutations [11]. Claesson and Guðmundsson then computed a rational generating function for permutations sorted by any finite number of passes [3], and asked whether a “useful permutation pattern characterization of the k -pop stack-sortable permutations” exists for $k \geq 3$.

In this paper we provide such a characterisation. To do so, we realised the current notions of what it means to avoid a set of barred and unbarred patterns would not suffice (see Subsection 2.1). We therefore introduce a new notion called *2-avoidance*, which could

be of independent interest. In Appendix B we explain how 2-avoidance differs from the notion of mesh pattern avoidance.

We make the following observations. First, our present result is in contrast to the usual (nondeterministic) stacks-in-series model where in many cases no finite pattern-avoidance characterisation is possible due to the existence of infinite antichains [4, 10]. Second, the operation of a pop stack is related to classical sorting: “bubble-sort” is exactly sorting by arbitrarily many passes through a pop stack of depth 2. Third, pop stacks are a natural model for genome rearrangement [11].

2 Preliminaries

Define a *permutation* to be a sequence of distinct real numbers, written as $\alpha = a_1 a_2 \cdots a_r$ (we continue the convention to write sequences without commas). The *reduced form* of a permutation α , denoted $\text{red}(\alpha)$, is the permutation obtained by replacing the i th smallest entry of α by the integer i . A permutation in reduced form is called *reduced*. We denote the set of all reduced permutations by S^∞ . Two permutations $\alpha = a_1 \cdots a_r$ and $\beta = b_1 \cdots b_s$ are *order-isomorphic*, denoted $\alpha \sim \beta$, if they have the same reduced form. For example 253 and 132 are order-isomorphic. In general, we will write permutations in their reduced form.


A *subpermutation* of $\alpha = a_1 \cdots a_r$ is a sequence $a_{i_1} \cdots a_{i_s}$ where $1 \leq i_1 < \dots < i_s \leq r$, while a *factor* is a sequence $a_{i_1} \cdots a_{i_s}$ where $i_{j+1} = i_j + 1$ for $1 \leq j \leq s - 1$. For two permutations α and β , we say β *contains* α if some subpermutation of β is order-isomorphic to α . We use the notation $\alpha \leq \beta$ if β contains α , and the notation $\alpha <_{\text{subperm}} \beta$ if α is a subpermutation of β . For example, $423 <_{\text{subperm}} 14253$, and $312 \leq 14253$ since $312 \sim 423$, but $312 \not<_{\text{subperm}} 14253$.

We say β *avoids* α if no subpermutation of β is order-isomorphic to α . For example 54123 contains 312 and avoids 231. For any set of permutations F , let $\text{Av}(F) \subseteq S^\infty$ denote the set of all permutations that simultaneously avoid every $\alpha \in F$.


Knuth famously observed that a permutation can be sorted by passing it through a single stack if and only if it avoids 231 [8]. Avis and Newborn proved a permutation can be sorted by passing it through a pop stack once if and only if it avoids both 231 and 312. However, for multiple passes through a pop stack, the situation arises that some permutation cannot be sorted, while a longer permutation containing it can, so the usual notion of pattern avoidance is not useful for characterising permutations in this context (in other words, for $k \geq 2$, k -pass pop stack sortable permutations are not a closed class with respect to usual pattern avoidance). As a concrete example, Figure 3 shows that 3241 is not 2-pass pop stack sortable, whereas 41352 is (as demonstrated by Figure 2), and contains 3241.

This leads us to the following notion.

Definition 1 (2-containment). Let σ be a permutation and $F, G \subseteq S^\infty$ be two sets of reduced permutations. We say that σ *2-contains* (F, G) if there exists a subpermutation γ of σ such that



(a) First pass.



(b) Second pass.

Figure 3: 3241 is not a 2-pass pop stack sortable.

- $\text{red}(\gamma) \in F$ and
- there is no $\delta <_{\text{subperm}} \sigma$ such that $\gamma <_{\text{subperm}} \delta$ and $\text{red}(\delta) \in G$.

Informally, we think of the set G as patterns which can potentially “save” a permutation from being forbidden by F .

A permutation *2-avoids* (F, G) if it does not 2-contain (F, G) . Using propositional logic we can express this as follows.

Definition 2 (2-avoidance). Let σ be a permutation and $F, G \subseteq S^\infty$ be two sets of reduced permutations. We say that σ *2-avoids* (F, G) if for all $\gamma <_{\text{subperm}} \sigma$, if $\text{red}(\gamma) \in F$ then there exists $\delta <_{\text{subperm}} \sigma$ such that $\gamma <_{\text{subperm}} \delta$ and $\text{red}(\delta) \in G$.

We denote the set of all permutations in S^∞ which 2-avoid (F, G) by $\text{Av}_2(F, G)$.

Example 3. Let $F = \{3241\}, G = \{41352\}$. Then $\sigma_1 = 143562$ has the subpermutation $4352 \sim 3241$ which is not part of a longer subpermutation of σ_1 order-isomorphic to 41352 , so σ_1 2-contains (F, G) . Now consider $\sigma_2 = 152463$ which has subpermutation $5463 \sim 3241$. However 5463 is a subpermutation of $52463 \sim 41352$, so 5463 is *saved* by G . Since there are no other subpermutations of σ_2 that are order-isomorphic to 3241 , we have that σ_2 2-avoids (F, G) .

Example 4. Let $F = \{1\}, G = \{12, 21\}$. Then $\text{Av}_2(F, G)$ consists of *all* permutations in S^∞ except the permutation of length 1. This is a rather extreme example, but shows that the growth of 2-avoidance sets can be factorial.

Using this notion, we can express the result of Pudwell and Smith as follows.

Theorem 5 (Pudwell and Smith [11]). *The set of 2-pass pop stack sortable permutations is equal to*

$$\text{Av}_2(\{2341, 3412, 3421, 4123, 4231, 4312, 3241, 4132\}, \{41352\}).$$

At this point the reader might object and say that Pudwell and Smith’s use of *barred patterns* (see the next subsection) is more efficient to describe this set, which it is. However for 3-pop stack sortable permutations, the usual definition of avoiding sets of barred and unbarred patterns will fail.

2.1 The problem with barred patterns

A *barred pattern* is a permutation where certain entries (possibly none) are marked with a bar. For example, $\beta = 4\bar{3}\bar{5}\bar{1}2$. A barred pattern encodes two permutations, one called $\text{removebar}(\beta)$ obtained by deleting entries marked with a bar, and the other called $\text{unbar}(\beta)$ obtained by removing bars. For example, $\text{removebar}(4\bar{3}\bar{5}\bar{1}2) = 452 \sim 231$, and $\text{unbar}(4\bar{3}\bar{5}\bar{1}2) = 43512$.

In [7, Definition 1.2.3] a permutation σ is said to *avoid a barred pattern* β if each occurrence of $\text{removebar}(\beta)$ in σ (if any) is a part of an occurrence of $\text{unbar}(\beta)$ in σ . There are two issues with this definition. Firstly, as written, this does not agree with the usual pattern avoidance when β itself has no bar entries. For example, $\sigma = 21$ obviously does not avoid itself in the usual sense of pattern avoidance, but if $\beta = 21$ is considered as a barred pattern then σ avoids β since it has a subpermutation, $\kappa = 21 = \text{removebar}(\beta)$, and κ is part of an occurrence of $\text{unbar}(\beta) = 21$ in σ .

Secondly and more seriously, in applications such as [11, 13] some set of permutations S is characterised by being those permutations avoiding some list of barred and unbarred patterns, where this means that each permutation in S must avoid every pattern individually. Explicitly, Tenner [12] defines, for P a collection of barred and unbarred patterns, $\text{Av}(P)$ to be the set of permutations simultaneously avoiding all patterns in P :

$$\text{Av}(P) = \bigcap_{p \in P} \text{Av}(p)$$

However, this notion does not suffice to describe 3-pop stack sortable permutations: one may verify that 32451 is not 3-pop stack sortable while both 4631572 and 4731562 are. If we were to characterise 3-pop stack sortable permutations as those avoiding a list containing $4\bar{6}\bar{3}\bar{1}572$ and $4\bar{7}\bar{3}\bar{1}562$, then we would be mistaken since 4731562 does not avoid this list since it fails to avoid the barred pattern $4\bar{6}\bar{3}\bar{1}572$. Theorem 1 says that even though some permutation may contain a subpermutation which is forbidden, it can be saved if it extends to another subpermutation which appears somewhere on the list G . For the applications in [11, 13] the sets of barred and unbarred patterns to be avoided have no “overlap”: in the case of [11], the two barred patterns have removebar equal to 2341 and 4312 which are both different from the unbarred patterns in their list; and in the case of [13] there is only one barred pattern whose removebar is different from the unbarred pattern.

Note that our definition of 2-avoidance also differs from the notion of avoiding *mesh patterns* as defined in [2]. See Appendix B for a brief discussion.

2.2 Removing redundant patterns

In general, one cannot simply delete patterns from F if they contain shorter elements of F , or delete patterns from G if they are contained in longer elements of G , since 2-containment involves a subtle interplay between the two sets. The following two examples demonstrate this.

Example 6. If $F = \{43251, 3241\}$ and $G = \{41352\}$ then we claim that the pattern $43251 \in F$ is not redundant, even though it contains a shorter element in F . Let $F_1 = \{3241\}$. Then $\sigma = 6251473$ 2-avoids (F_1, G) , but 2-contains (F, G) . Thus $\text{Av}_2(F, G) \neq \text{Av}_2(F_1, G)$.

Example 7. If $F = \{4123, 4231, 43251, 3241\}$ and $G = \{41352\}$, then we claim that the pattern $43251 \in F$ is redundant. Suppose σ 2-contains (F, G) . Either this is because of $\beta \in F \setminus \{43251\}$, or not. If we suppose not, then σ must 2-contain (F, G) because of 43251 . If σ contains 43251 then it contains 3241 , so each of the subpermutations $\gamma \in \{4351, 4251, 3251\}$ must be saved by $41352 \in G$ (otherwise we would say σ 2-contains (F, G) because of some $\beta \in F \setminus \{43251\}$). Thus σ must have a subpermutation $4a3b251$ such that $4a351 \sim 4a251 \sim 3b251 \sim (41352)$, so $a, b < 1$. Therefore, $4a3b251$ is either order isomorphic to $403(-1)251$ (so σ contains 4231) or $4(-1)30251$ (so σ contains 4123). Thus σ 2-contains F because of 4123 or 4231 which contradicts that it is because of 43251 only. Thus σ 2-contains $(F \setminus \{43251\}, G)$. Conversely if σ 2-contains $(F \setminus \{43251\}, G)$ then it clearly 2-contains (F, G) . Thus $\text{Av}_2(F, G) = \text{Av}_2(F \setminus \{43251\}, G)$.

However, we can state some general rules for removing redundant elements of F or G .

Lemma 8. If $\alpha \in G$ and for all $\gamma \in F$, $\gamma \not\leq \alpha$, then $\text{Av}_2(F, G) = \text{Av}_2(F, G \setminus \{\alpha\})$.

Lemma 9. If $\alpha, \beta \in G$, $\beta \leq \alpha$, $\beta \neq \alpha$, and for all $\kappa \in F$, $\kappa \leq \alpha$ implies $\kappa \leq \beta$, then $\text{Av}_2(F, G) = \text{Av}_2(F, G \setminus \{\alpha\})$.

Lemma 10. If $\kappa, \lambda \in F$, $\kappa \leq \lambda$ and for all $\alpha \in G$, $\kappa \not\leq \alpha$, then $\text{Av}_2(F, G) = \text{Av}_2(F \setminus \{\lambda\}, G)$.

Proofs can be found in Appendix A (no other results depend on them). Certainly we believe that further lemmas could be stated and proved to remove more redundant elements. For the purpose of this paper we do not pursue this; we content ourselves to find a characterisation for k -pop stack sortable permutations in terms of finite sets only.

3 Blocks

Let σ be a permutation. Call a factor $B_i = a_{i,1}a_{i,2} \cdots a_{i,n_i}$ of σ a *block* if $n_i > 0$ and $a_{i,j} > a_{i,j+1}$ for all $1 \leq j < n_i$. (Recall that factor means the entries are contiguous in σ .) A (*maximal*) *block decomposition* of σ is an expression of the form $\sigma = B_1B_2B_3 \cdots B_m$

where each B_i is a block and for any two adjacent blocks $B_i = a_{i,1}a_{i,2} \cdots a_{i,n_i}$ and $B_{i+1} = a_{i+1,1}a_{i+1,2} \cdots a_{i+1,n_{i+1}}$ we have $a_{i,n_i} < a_{i+1,1}$. For example $\sigma = 87634521$ has block decomposition $B_1 = 8763, B_2 = 4, B_3 = 521$. For convenience we indicate the block decomposition of σ by inserting $|$ symbols to separate blocks, so for our example we write $8763 | 4 | 521$.

If $B_i = a_{i,1}a_{i,2} \cdots a_{i,n_i}$ is a block, let $\widetilde{B}_i = a_{i,n_i} \cdots a_{i,2}a_{i,1}$. We have the following.

Lemma 11 ([11]). *If σ has block decomposition $B_1B_2B_3 \cdots B_m$ then*

$$p_1(\sigma) = \widetilde{B}_1\widetilde{B}_2\widetilde{B}_3 \cdots \widetilde{B}_m.$$

For example $\sigma = 987354621 = 9873 | 54 | 621$, so $p_1(\sigma) = 3789 45 126$.

Lemma 12 ([3]). *Let σ be a permutation. Then each block in the block decomposition of $p_1(\sigma)$ contains at most 3 entries.*

For example, if $\sigma = 52341 = 52 | 3 | 41$ then $p_1(\sigma) = 25314 = 2 | 531 | 4$.

Lemma 13. *Let σ be a permutation with block decomposition $B_1B_2B_3 \cdots B_m$, and let $a \in B_{i+1}, b \in B_{i+n}$ be two entries of σ with $a > b, n \geq 1$. (See for example Figure 4). If $n \geq 3^k$ then σ is not k -pass pop stack sortable.*

Proof. We proceed by induction. For $k = 0$, σ is not sortable by 0 passes since $a > b$. Assume the statement is true for $k \in \mathbb{N}$, and consider $\sigma = B_1 \cdots B_{i+1} \cdots B_{i+n} \cdots B_m$, $a \in B_{i+1}, b \in B_{i+n}$ with $n \geq 3^{k+1}$. Colour the entries a, b and all entries in $B_{i+2}, \dots, B_{i+n-1}$ **bold**. Then the number of bold entries in σ is at least $n \geq 3^{k+1}$ since each block contains at least one entry. After one pass, each block of $p_1(\sigma)$ can contain at most 3 entries, so the number of blocks with bold entries is at least $\frac{3^{k+1}}{3} = 3^k$. Since $p_1(\sigma) = \widetilde{B}_1 \cdots \widetilde{B}_{i+1} \widetilde{B}_{i+2} \cdots \widetilde{B}_{i+n-1} \widetilde{B}_{i+n} \cdots \widetilde{B}_m$ and $\widetilde{B}_{i+1}, \widetilde{B}_{i+n}$ contain just one bold entry each (a, b respectively), we have that a must be in the first block with bold entries in the block decomposition of $p_1(\sigma)$, and b in the last, so by inductive hypothesis $p_1(\sigma)$ cannot be sorted by k passes. Thus σ cannot be sorted by $k + 1$ passes. \square

We remark that the bound of 3^k in the preceding statement is most likely an extreme overestimate. However for the purpose of this paper any bound suffices.

4 General characterisation of k -pass pop stack sortable permutations

Theorem 14. *Let $k \in \mathbb{N}_+$. There exists a pair of finite sets (F_k, G_k) such that the set of all k -pass pop stack sortable permutations is equal to $\text{Av}_2(F_k, G_k)$. Moreover, the sets F_k, G_k can be algorithmically constructed.*

Proof. Let S_k denote the set of all k -pass pop stack sortable permutations. We proceed by induction, with the base case $k = 1$ established by Avis and Newborn [1] (specifically, $F_1 = \{231, 312\}, G_1 = \emptyset$). Assume F_{k-1}, G_{k-1} have been constructed, are finite, and $S_{k-1} = \text{Av}_2(F_{k-1}, G_{k-1})$.

Let $f_{\max} = \max\{|\beta| \mid \beta \in F_{k-1}\}$ and $C = 3^{k+2}f_{\max}$. Then define the following two sets

$$\begin{aligned}\Omega_1 &= \{\tau \in S^\infty \mid |\tau| \leq 3f_{\max} \text{ and } \tau \notin S_k\}, \\ \Omega_2 &= \{\kappa \in S^\infty \mid |\kappa| \leq C, \kappa \in S_k \text{ and there exists } \tau \in \Omega_1 \text{ such that } \tau \leq \kappa\}.\end{aligned}$$

Claim 1: Ω_1 and Ω_2 are both finite: both are subsets of the set of all permutations of length at most C .

Claim 2: Ω_1 and Ω_2 are algorithmically constructible: we only have finitely many τ, κ of length at most $3f_{\max}, C$ respectively. For each τ, κ we can check $\tau \notin S_k$ or $\kappa \in S_k$ in linear time by passing them according to the deterministic procedure, and we can check all subpermutations of κ of length at most $3f_{\max}$ to see whether or not at least one has reduced form in Ω_1 .

Claim 3: $\sigma \notin S_k$ if and only if σ 2-contains (Ω_1, Ω_2) : we have $\sigma \notin S_k$ if and only if $p_1(\sigma) \notin S_{k-1}$ if and only if $p_1(\sigma)$ 2-contains (F_{k-1}, G_{k-1}) .

To prove the forwards direction of Claim 3, suppose $p_1(\sigma)$ 2-contains (F_{k-1}, G_{k-1}) , and further assume this is because of some $\zeta <_{\text{subperm}} p_1(\sigma)$ with $\text{red}(\zeta) \in F_{k-1}$, and there is no $\delta <_{\text{subperm}} \sigma$ with $\zeta <_{\text{subperm}} \delta$ and $\text{red}(\delta) \in G_{k-1}$. Note, there may be many choices of ζ to take, but fix one choice.

1. Mark entries corresponding to ζ in $p_1(\sigma)$ **bold**. Let $\zeta' <_{\text{subperm}} \sigma$ be such that after one pass, the entries belonging to ζ' are the bold entries corresponding to $\zeta <_{\text{subperm}} p_1(\sigma)$. Mark the ζ' entries bold as well. Note that $|\zeta'| \leq f_{\max}$.

For example, if $\sigma = 987354621$ then $p_1(\sigma) = 378945126$ which 2-contains $(\{231, 312\}, \emptyset)$ because of (for instance) the subpermutation $\zeta = 956 \sim 312$ of $p_1(\sigma)$. We write $p_1(\sigma)$ as **378945126**, and thus σ as **987354621**.

2. Next, write σ in block decomposition $\sigma = B_1B_2B_3 \cdots B_m$ with $B_i = a_{i,1} \cdots a_{i,n_i}$. Say that B_i is **bold** if it contains at least one bold entry (from ζ'). We wish to delete non-bold entries of σ but we do not want to *merge* bold blocks, so we apply the following subroutine.

- Set $\tau = \sigma$.
- While $a_{i,j}$ is a non-bold letter,
 - if removing $a_{i,j}$ from τ does not cause two or more bold blocks to merge, delete $a_{i,j}$ from τ .

We claim that at the end of this process $|\tau| \leq 3|\zeta'| \leq 3f_{\max}$. Let $a_{i,j} \in B_i$ with $1 < i < m$ be a non-bold entry in τ . If at most one of B_{i-1}, B_i, B_{i+1} is bold, then removing $a_{i,j}$ cannot merge bold blocks. Else assume at least two of B_{i-1}, B_i, B_{i+1} are bold. If $a_{i,j}$ is not the first or last entry in B_i , it can be deleted without merging blocks. This leaves at most two unbold entries in B_i . For B_1 (resp. B_m) we can delete all except the last (resp. first) entry without merging bold blocks. This leaves at most two unbold entries in each block. Then in the worst case each block contains just one bold entry, with an unbold entry on either side. For example, if we get to $\tau = \dots | 12, \mathbf{10}, 8 | \mathbf{975} | \mathbf{642} | \mathbf{31}$ then we cannot delete 8, 9, 5, 6, 2, 3 without merging blocks.

3. After this, we obtain a permutation $\tau <_{\text{subperm}} \sigma$ such that the bold letters $\zeta' <_{\text{subperm}} \tau$ and $|\tau| \leq 3f_{\max}$.

We now claim that $p_1(\tau)$ 2-contains (F_{k-1}, G_{k-1}) because of the same subpermutation ζ of $p_1(\sigma)$. Since bold blocks are preserved in τ , we know that $p_1(\tau)$ will also contain ζ . Now suppose there is some $\delta <_{\text{subperm}} p_1(\tau)$ ($<_{\text{subperm}} p_1(\sigma)$) with $\text{red}(\delta) \in G_{k-1}$ and $\zeta <_{\text{subperm}} \delta$. This means that the same δ saves $(\zeta, p_1(\sigma))$, which contradicts our original assumption. Thus $p_1(\tau)$ 2-contains (F_{k-1}, G_{k-1}) , which implies $p_1(\tau) \notin S_{k-1}$, which implies $\tau \notin S_k$.

Thus since $|\tau| \leq 3f_{\max}$ and $\tau \notin S_k$, we have $\tau \in \Omega_1$ by definition. To finish this direction, we will show that τ is not saved by any subpermutation order-isomorphic to something in Ω_2 .

Suppose (for contradiction) that there is some $\delta <_{\text{subperm}} \sigma$ with $\tau <_{\text{subperm}} \delta$ and $\text{red}(\delta) \in \Omega_2$. This means $\delta \in S_k$, so $p_1(\delta) \in S_{k-1}$, and so $p_1(\delta)$ 2-avoids (F_{k-1}, G_{k-1}) .

Now $p_1(\delta)$ will contain ζ since blocks containing $\zeta' <_{\text{subperm}} \tau <_{\text{subperm}} \delta$ will not merge after one pass. Since $p_1(\delta)$ 2-avoids (F_{k-1}, G_{k-1}) and contains ζ , there must be some $\alpha \in G_{k-1}$ which saves $\zeta <_{\text{subperm}} p_1(\delta)$. This means there is some $\alpha' <_{\text{subperm}} p_1(\delta)$ with $\zeta <_{\text{subperm}} \alpha'$ and $\alpha' \sim \alpha$.

We claim $\alpha \in G_{k-1}$ also saves $\zeta <_{\text{subperm}} p_1(\sigma)$, since there exists

$$\alpha' <_{\text{subperm}} p_1(\delta) <_{\text{subperm}} p_1(\sigma)$$

with $\zeta <_{\text{subperm}} \alpha'$ and $\alpha' \sim \alpha$. This contradicts that $p_1(\sigma)$ 2-contains (F_{k-1}, G_{k-1}) because of ζ . Thus we have shown $\sigma \notin S_k$ implies σ 2-contains (Ω_1, Ω_2) .

Now for the converse direction of Claim 3, suppose that σ 2-contains (Ω_1, Ω_2) . Then we can assume that this is because of $\gamma <_{\text{subperm}} \sigma$ and $\tau \in \Omega_1$ with $\gamma \sim \tau$, which is not saved by any $\alpha \in \Omega_2$, and so by definition $\gamma \notin S_k$. Thus $p_1(\gamma)$ 2-contains (F_{k-1}, G_{k-1}) .

Assume (for contradiction) that $\sigma \in S_k$. We will show that this implies we can construct some $\kappa <_{\text{subperm}} \sigma$ such that $\kappa \in S_k$, $\gamma <_{\text{subperm}} \kappa$ and $|\kappa| \leq C$. If so, then we can construct $\alpha \in \Omega_2$ with $\alpha \sim \kappa$ and $\text{red}(\gamma) \leq \alpha$, which means α saves $\gamma <_{\text{subperm}} \sigma$, and this gives a contradiction that σ 2-contains (Ω_1, Ω_2) because of γ .

Here is how we construct κ . In σ , mark the entries corresponding to γ **bold**.

Call a block of σ bold if it contains at least one bold entry, and otherwise call it non-bold. Starting with $\kappa = \sigma$, we delete non-bold entries using the following procedure,

which is more careful than the subroutine used in the proof of the forward direction above. The goal is to delete non-bold entries to obtain a permutation κ with subpermutation γ such that for *every* block B in κ there is a block B' in σ so that the entries in B are entries in B' . That is, we do not allow *any* (bold or unbold) blocks to merge, only to be deleted entirely.

- Set $\kappa = \sigma$.
- While $a_{i,j} \in B_i$ is a non-bold letter,
 - if removing $a_{i,j}$ from κ does not cause two or more blocks of *any kind* (bold or non-bold) in κ to merge, delete $a_{i,j}$ from κ ,
 - if B_i is non-bold and removing the entire block B_i at once does not cause any of the remaining blocks to merge, then delete B_i .

We claim that at the end of this process each block contains at most two non-bold entries, which will be the first and last entries of the block. However, since we have not deleted non-bold blocks if their removal would cause other blocks to merge, we could have arbitrarily long factors of non-bold blocks, as in the example in Figure 4.




Figure 4: $\kappa = \mathbf{16}, 14, 15, 12, 13, 10, 11, 896745231$.

Suppose κ has block decomposition $T_1 \cdots T_m$, and κ contains two bold entries a, b with $a \in T_{i+1}$ and $b \in T_{i+n}$, where $T_{i+2}, \dots, T_{i+n-1}$ have no bold entries. (Note that each T_j is a subset of some B_k in the block decomposition of σ , by construction.) If $a < b$ then the subroutine is not complete: we could delete $T_{i+2}, \dots, T_{i+n-1}$ completely without

merging T_{i+1} with T_{i+n} since the last entry of T_{i+1} is smaller than the first entry of T_{i+n} . Therefore since we assume the subroutine is complete, we know that $a > b$. If $n \geq 3^k$ then, by Lemma 13, σ cannot be sorted by k passes, which is a contradiction. Thus we have $|\kappa|$ is bounded above by $3|\gamma| + 2 \cdot 3^k |\gamma|$, where $3|\gamma|$ counts the bold blocks with a non-bold first and last entry, and $2 \cdot 3^k |\gamma|$ counts the 3^k factors each containing 2 entries, in the worst case, occurring between every pair of bold entries from γ , as in Figure 4.

Thus

$$|\kappa| \leq (3 + 2 \cdot 3^k) |\gamma| \leq (3 + 2 \cdot 3^k) 3f_{\max} \leq 3^{k+2} f_{\max} = C.$$

If κ cannot be sorted, then $p_1(\kappa) \notin S_{k-1}$ because of some subpermutation τ with $\text{red}(\tau) \in F_{k-1}$, but since no block has merged in obtaining $p_1(\kappa)$, $p_1(\sigma)$ also contains τ which cannot be saved since blocks containing the entries forming τ are fixed. Thus $p_1(\sigma) \notin S_{k-1}$, so $\sigma \notin S_k$, contradiction.

This proves Claim 3.

Claim 3 implies that we could take $F_k = \Omega_1, G_k = \Omega_2$ and the theorem is done. However, we can first apply Theorems 9 and 10¹ to obtain smaller sets with the same 2-avoidance set, so we will call the result of applying these (in some order) (F_k, G_k) . As remarked in Subsection 2.2, the result of applying the lemmas is not guaranteed to give a set that is minimal or unique. Note that each of these lemmas needs to check a finite set, so each is algorithmic. \square

Applying the construction in the proof above to the case $k = 2$ yields a much larger pair of sets than those appearing in Theorem 5 given by Pudwell and Smith. Explicitly, we have $f_{\max} = 3$ and $k = 2$, so $C = 3 \cdot 3^4 = 243$. The point of our argument is that it is general; our upper bounds can certainly be lowered, and more lemmas to reduce the size of the sets F, G could be proved along the lines of Lemmas 8–10 to sharpen the result and render it more tractable for computation. See [5] for a discussion of an attempt to compute explicit avoidance sets for the case $k = 3$, where if we start with the sets F_2, G_2 in Theorem 5 given by Pudwell and Smith rather than the larger sets produced by our algorithm, we have $f_{\max} = 4$ and $C = 4 \cdot 3^5 = 972$ (so a brute force calculation to find Ω_1, Ω_2 would need to examine $C! = 972!$ permutations).

5 Outlook

Our notion of 2-containment opens up some interesting possibilities. Recall that by Kaiser-Klazar [6, Thm. 3.4] and Marcus-Tardos [9] the function counting the number of permutations of length n in any $\text{Av}(F)$ for F non-empty is either polynomial or exponential. It is conceivable that some pair of (finite or infinite) sets (F, G) could have a 2-avoidance set with a growth function strictly between polynomial and exponential or between exponential and factorial. Example 4 shows non-trivial 2-avoidance sets with factorial growth.

¹Note that by construction there is no need to apply Lemma 8.

Generating functions for 2-avoidance sets might also exhibit interesting behaviour. For the sets (F_k, G_k) in Theorem 14 we know by [3] the generating functions are rational for all k , but for general sets F, G , the set $\text{Av}_2(F, G)$ could have interesting enumerations.

References

- [1] David Avis and Monroe Newborn. On pop-stacks in series. *Utilitas Math.*, 19:129–140, 1981.
- [2] Petter Brändén and Anders Claesson. Mesh patterns and the expansion of permutation statistics as sums of permutation patterns. *Electron. J. Combin.*, 18(2):Paper 5, 14, 2011.
- [3] Anders Claesson and Bjarki Ágúst Guðmundsson. Enumerating permutations sortable by k passes through a pop-stack. *Sém. Lothar. Combin.*, 80B:Art. 43, 12, 2018.
- [4] Murray Elder and Yoong Kuan Goh. Permutations sorted by a finite and an infinite stack in series. In *Language and automata theory and applications*, volume 10792 of *Lecture Notes in Comput. Sci.*, pages 220–231. Springer, Cham, 2018.
- [5] Yoong Kuan Goh. *Variations on stack sorting and pattern avoidance*. PhD thesis, University of Technology Sydney, 2020.
- [6] Tomáš Kaiser and Martin Klazar. On growth rates of closed permutation classes. *Electron. J. Combin.*, 9(2):Research paper 10, 20, 2002/03. Permutation patterns (Otago, 2003).
- [7] Sergey Kitaev. *Patterns in Permutations and Words*. Springer Verlag Berlin and Heidelberg GmbH and Co. KG, 2011. Monographs in Theoretical Computer Science. an Eatsc.
- [8] Donald E. Knuth. *The art of computer programming. Volume 3*. Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1973. Sorting and searching, Addison-Wesley Series in Computer Science and Information Processing.
- [9] Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Combin. Theory Ser. A*, 107(1):153–160, 2004.
- [10] Maximilian M. Murphy. *Restricted permutations, antichains, atomic classes, stack sorting*. PhD thesis, University of St Andrews, 2003.
- [11] Lara Pudwell and Rebecca Smith. Two-stack-sorting with pop stacks. *Australasian Journal of Combinatorics*, 74(1):179 – 195, 2019.
- [12] Bridget Eileen Tenner. Coincidental pattern avoidance. *J. Comb.*, 4(3):311–326, 2013.
- [13] Julian West. Sorting twice through a stack. *Theoret. Comput. Sci.*, 117(1-2):303–313, 1993. Conference on Formal Power Series and Algebraic Combinatorics (Bordeaux, 1991).

A Removing redundant permutations

We start with the following observation.

Lemma 15. *Let $F_1, F, G_1, G \subseteq S^\infty$ with $F_1 \subset F, G_1 \subset G$. Then*

- $\text{Av}_2(F, G_1) \subseteq \text{Av}_2(F, G)$ and
- $\text{Av}_2(F, G) \subseteq \text{Av}_2(F_1, G)$.

Proof. If σ 2-contains (F, G) , then there exists a subpermutation γ of σ with $\text{red}(\gamma) \in F$, such that there is no δ with $\gamma <_{\text{subperm}} \delta <_{\text{subperm}} \sigma$ and $\text{red}(\delta) \in G$. In particular, there is no δ with $\gamma <_{\text{subperm}} \delta <_{\text{subperm}} \sigma$ and $\text{red}(\delta) \in G_1$, so σ 2-contains (F, G_1) . Thus $\text{Av}_2(F, G_1) \subseteq \text{Av}_2(F, G)$.

On the other hand, if σ 2-contains (F_1, G) then there exists a subpermutation γ of σ with $\text{red}(\gamma) \in F_1$, such that there is no δ with $\gamma <_{\text{subperm}} \delta <_{\text{subperm}} \sigma$ and $\text{red}(\delta) \in G$. In particular $\text{red}(\gamma) \in F$, so σ 2-contains (F, G) . Thus $\text{Av}_2(F, G) \subseteq \text{Av}_2(F_1, G)$. \square

Proof of Lemma 8. Recall that we wish to show that if $\alpha \in G$ and for all $\gamma \in F, \gamma \not\leq \alpha$, then $\text{Av}_2(F, G) = \text{Av}_2(F, G \setminus \{\alpha\})$.

By Theorem 15 it suffices to show $\text{Av}_2(F, G) \subseteq \text{Av}_2(F, G \setminus \{\alpha\})$. If σ 2-contains $(F, G \setminus \{\alpha\})$ then there exists a subpermutation γ of σ such that $\text{red}(\gamma) \in F$ and there is no δ satisfying $\gamma <_{\text{subperm}} \delta <_{\text{subperm}} \sigma$ and $\text{red}(\delta) \in G \setminus \{\alpha\}$. Since $\gamma \not\leq \alpha$, then $\text{red}(\delta)$ is not equal to α , so we have that there can be no δ with $\gamma <_{\text{subperm}} \delta <_{\text{subperm}} \sigma$ and $\text{red}(\delta) \in G$. Thus σ 2-contains (F, G) , and so $\text{Av}_2(F, G) \subseteq \text{Av}_2(F, G \setminus \{\alpha\})$. \square

Proof of Lemma 9. Recall that we wish to show that if $\alpha, \beta \in G, \beta \leq \alpha, \beta \neq \alpha$, and for all $\kappa \in F, \kappa \leq \alpha$ implies $\kappa \leq \beta$, then $\text{Av}_2(F, G) = \text{Av}_2(F, G \setminus \{\alpha\})$.

By Theorem 15 it suffices to show $\text{Av}_2(F, G) \subseteq \text{Av}_2(F, G \setminus \{\alpha\})$. If σ 2-contains $(F, G \setminus \{\alpha\})$ then there exists a subpermutation γ of σ such that $\text{red}(\gamma) \in F$, and there is no δ satisfying $\gamma <_{\text{subperm}} \delta <_{\text{subperm}} \sigma$ and $\text{red}(\delta) \in G \setminus \{\alpha\}$. If there is a δ' with $\gamma <_{\text{subperm}} \delta' <_{\text{subperm}} \sigma$ and $\text{red}(\delta') = \alpha$, then $\gamma \leq \alpha$ which implies $\gamma \leq \beta$. This contradicts that no δ with $\text{red}(\delta) = \beta \in G \setminus \{\alpha\}$ can exist with $\gamma <_{\text{subperm}} \delta <_{\text{subperm}} \delta' <_{\text{subperm}} \sigma$. Thus σ 2-contains (F, G) , and so $\text{Av}_2(F, G) \subseteq \text{Av}_2(F, G \setminus \{\alpha\})$. \square

Proof of Lemma 10. Recall that we wish to show that if $\kappa, \lambda \in F, \kappa \leq \lambda$ and for all $\alpha \in G, \kappa \not\leq \alpha$, then $\text{Av}_2(F, G) = \text{Av}_2(F \setminus \{\lambda\}, G)$.


By Theorem 15 it suffices to show $\text{Av}_2(F \setminus \{\lambda\}, G) \subseteq \text{Av}_2(F, G)$. If σ 2-contains (F, G) then there exists a subpermutation γ of σ such that $\text{red}(\gamma) \in F$, and there is no δ satisfying $\gamma <_{\text{subperm}} \delta <_{\text{subperm}} \sigma$ and $\text{red}(\delta) \in G$. If $\text{red}(\gamma) = \lambda$, then there exists a subpermutation γ' of σ with $\text{red}(\gamma') = \kappa$, since $\kappa \leq \lambda$. Since $\kappa \not\leq \alpha$ for all $\alpha \in G$, then there is no $\delta <_{\text{subperm}} \sigma$ with $\gamma' <_{\text{subperm}} \delta$ and $\text{red}(\delta) \in G$. Thus σ 2-contains $(F \setminus \{\lambda\}, G)$, and so $\text{Av}_2(F \setminus \{\lambda\}, G) \subseteq \text{Av}_2(F, G)$. \square

B Mesh patterns


For the sake of disambiguation we show here that 2-avoidance is in general a different notion from that of avoiding (a set of) mesh patterns.

For $n \in \mathbb{N}$ let S^n denote the set of all reduced permutations of length n , and let $[a, b]$ be the integer interval $\{i \in \mathbb{Z} \mid a \leq i \leq b\}$.

Following [2], a *mesh pattern* is a pair (π, R) where $\pi \in S^n$ and $R \subseteq [0, n] \times [0, n]$. We visualise it by drawing the points $\{(i, \pi(i)) \mid 1 \leq i \leq n\}$ in the $[0, n] \times [0, n]$ lattice, then shading all squares whose bottom left corner lies in R . See Figure 5.



(a) Mesh pattern $(3241, \{(1, 0)\})$.



(b) Mesh pattern $(4132, \{(3, 4)\})$.

Figure 5: Two mesh patterns.

A permutation $\tau \in S^\infty$ *contains* the mesh pattern (π, R) if $\pi \leq \tau$ (so $\pi \sim \pi' <_{\text{subperm}} \tau$) and when τ is mapped to the $[0, n] \times [0, n]$ lattice sending the entries of π' to the points $\{(i, \pi(i)) \mid 1 \leq i \leq n\}$, no entry from τ lies inside a shaded region. (See [2] for a formal definition.) A permutation *avoids* a mesh pattern if it does not contain it. For example, 41352 avoids both mesh patterns $(3241, \{(1, 0)\})$ and $(4132, \{(3, 4)\})$ shown in Figure 5.

Barred patterns with a single barred entry can be viewed as mesh patterns, where R consists of a single point. For example, the two mesh patterns in Figure 5 represent the two barred patterns $4\bar{1}352$ and $413\bar{5}2$ respectively. Explicitly, from [2], if $\pi(i)$ is the only barred letter of a given barred pattern π , then the corresponding mesh pattern is $(\text{red}(\text{removebar}(\pi)), \{(i-1, \pi(i)-1)\})$. However, avoiding a barred pattern β with more than one barred entry is not the same as avoiding the mesh pattern obtained in the analogous way (that is, taking $(\text{red}(\text{removebar}(\beta)), R)$, with R containing a point for each barred entry of β in the appropriate location). To see this, let $\pi = \text{red}(\text{removebar}(\beta))$ and note that by definition τ avoids (π, R) if for every subpermutation π' of τ with $\pi' \sim \pi$, there is *at least one* entry from τ inside a shaded square. By contrast, τ avoids the barred pattern β if for every subpermutation π' of τ with $\pi' \sim \pi$, there is a subpermutation δ of τ with $\pi' <_{\text{subperm}} \delta$ and $\delta \sim \text{unbar}(\beta)$. In other words, we would require one entry of τ in *every* shaded square.

One could potentially fix this by modifying the definition of avoiding a meshed pattern to encompass barred patterns with multiple barred entries. However, recall that in Subsection 2.1 we observed that 32451 is not 3-pop stack sortable while both 4631572 and 4731562 are. It is not clear how avoiding a set of mesh patterns could be defined to describe this situation. One might try to define mesh patterns $(32451, R_1)$ and $(32451, R_2)$ where R_1, R_2 encode positions which must be present in any permutation containing 32451, and modify the definition of avoiding a set of mesh patterns to say that a permutation τ avoids

a set of mesh patterns if for every subpermutation π' of τ with $\pi' \sim \pi$ and π appears in the first coordinate of a mesh pattern in the set, there is some R so that (π, R) is in the set, and τ contains an entry in *every* region coded by R . Following this kind of reasoning, we came to the definition of 2-avoidance rather than the existing notions.