Séminaire Lotharingien de Combinatoire **54A** (2006), Article B54Ab

COMBINATORICS OF PATIENCE SORTING PILES*

ALEXANDER BURSTEIN† AND ISAIAH LANKHAM‡

Abstract. Despite having been introduced in 1962 by C.L. Mallows, the combinatorial algorithm *Patience Sorting* is only now beginning to receive significant attention due to such recent deep results as the Baik-Deift-Johansson Theorem that connect it to fields including Probabilistic Combinatorics and Random Matrix Theory.

The aim of this work is to develop some of the more basic combinatorics of the Patience Sorting Algorithm. In particular, we exploit the similarities between Patience Sorting and the Schensted Insertion Algorithm in order to do things that include defining an analog of the Knuth relations and extending Patience Sorting to a bijection between permutations and certain pairs of set partitions. As an application of these constructions we characterize and enumerate the set $S_n(3\text{-}\bar{1}\text{-}42)$ of permutations that avoid the generalized permutation pattern 2-31 unless it is part of the generalized pattern 3-1-42.

Résumé. En dépit de l'introduction en 1962 par C.L. Mallows, l'algorithme combinatoire *Patience Sorting* commence seulement maintenant à susciter l'attention significative dû à des résultats profonds récents, tels que le théorème de Baik-Deift-Johansson, qui le relient à la combinatoire probabiliste et à la théorie des matrices aléatoires.

On développe une partie plus fondamentale de la combinatoire de l'algorithme de *Patience Sorting*. En particulier, on utilise les similarités entre *Patience Sorting* et la correspondance de Schensted pour définir un analogue des relations de Knuth et pour généraliser *Patience Sorting* à une bijection entre les permutations et certaines paires de partitions d'ensemble. Comme application de ces constructions on caractérise et énumère l'ensemble $S_n(3\text{-}\bar{1}\text{-}42)$ des permutations qui évitent le motif de permutation généralisé 2-31 à moins qu'il soit partie du motif généralisé 3-1-42.

Key words and phrases. Patience sorting, set partitions, Bell numbers, generalized permutation patterns, left-to-right minima subsequences, basic subsequences, shadow diagrams.

* This material is based in part upon work supported by the National Science Foundation under Grant No. DMS-0502858.

† Department of Mathematics, Iowa State University, Ames, IA 50011-2064, USA. (`burstein@math.iastate.edu`) Research supported in part by the U.S. NSA Young Investigator Grant H98230-06-1-0037.

‡ Department of Mathematics, University of California-Davis, Davis, CA 95616-8633, USA. (`issy@math.ucdavis.edu`) Research supported in part by the U.S. NSF Grants DMS-0135345 and DMS-0304414.

1. Introduction

The term *Patience Sorting* was introduced in 1962 by C. L. Mallows [15, 16] as the name of a card sorting algorithm invented by A. S. C. Ross. This algorithm works by first partitioning a shuffled deck of cards (which throughout this paper we take to be a permutation $\sigma \in \mathfrak{S}_n$) into sorted subsequences called *piles* using what Mallows referred to as a "patience sorting procedure":

Algorithm 1.1 (Mallows' Patience Sorting Procedure)**.** Given a shuffled deck of cards $\sigma = c_1 c_2 \cdots c_n$, inductively build the set of piles $R = R(\sigma) = \{r_1, r_2, \ldots, r_m\}$ as follows:

- Place the first card c_1 from the deck into a pile r_1 by itself.
- For each remaining card c_i $(i = 2, \ldots, n)$, consider the cards $d_1, d_2, \ldots, d_k$ atop the piles $r_1, r_2, \ldots, r_k$ that have already been formed.
  - If $c_i > \max\{d_1, d_2, \ldots, d_k\}$, then put c_i into a new right-most pile r_{k+1} by itself.
  - Otherwise, find the left-most card d_j that is larger than c_i and put the card c_i atop pile r_j.

We call the collection of piles $R(\sigma)$ the *pile configuration* associated to the deck of cards $\sigma \in \mathfrak{S}_n$ and illustrate their formation via an extended version of Algorithm 1.1 in Section 3.1 below.

Since each card c_i is either larger than the top card of every pile or is placed on top of the left-most top card d_j larger than it, the cards $d_1, d_2, \ldots, d_k$ atop the piles will be in increasing order from left to right at each step of the algorithm. Thus, Algorithm 1.1 resembles repeated application of the Schensted Insertion Algorithm (as discussed in [1]) for interposing a value into the increasing sequence $d_1, d_2, \ldots, d_k$ as if it were the top row of a Young tableau. The distinction is that cards remain in place and have other cards placed on top of them instead of being actively "bumped" from the row so that the Schensted Insertion Algorithm can then be recursively applied to the "bumped" value and the next lower row in the Young tableau. In this sense, Patience Sorting can be viewed as a non-recursive analog of the remarkable Robinson-Schensted-Knuth (or RSK) Algorithm due to G. Robinson [19] for permutations in 1938, C. Schensted [21] for words in 1961, and D. Knuth [12] for so-called $\mathbb{N}$-matrices in 1970. (See Fulton [10] for the appropriate definitions and for a detailed account of the differences between these algorithms.)

Recall that the RSK Algorithm bijectively associates an ordered pair of standard Young tableaux $(P(\sigma), Q(\sigma))$ to each permutation $\sigma = \sigma_1 \sigma_2 \cdots \sigma_n \in \mathfrak{S}_n$ by first building a so-called "insertion tableau" $P(\sigma)$ through repeated Schensted Insertion of the components $\sigma_1, \sigma_2, \ldots, \sigma_n$ into an initially empty tableau. It also simultaneously constructs the "recording tableau" $Q(\sigma)$ by literally *recording* how $P(\sigma)$ is formed. These tableaux have the same shape (a partition λ of n, denoted $\lambda \vdash n$), and this correspondence has many interesting properties. E.g., RSK applied to a permutation is symmetric in the sense that if $\sigma \in \mathfrak{S}_n$ corresponds to the ordered pair of tableaux $(P(\sigma), Q(\sigma))$, then $(Q(\sigma), P(\sigma))$ corresponds to the inverse permutation σ^{-1}. As a result, there is a bijection between the set of involutions $\mathfrak{I}_n \subset \mathfrak{S}_n$ and the set $\mathfrak{T}_n$ of all standard Young tableaux with entries $1, 2, \ldots, n$. (This is the famous Schützenberger symmetry property first proven in [22].)

In this paper we develop a bijective extension of Algorithm 1.1 and then study analogues for such properties of RSK. To facilitate this, we first characterize in Section 2 when two permutations have the same pile configurations under Algorithm 1.1. This yields an equivalence relation $\overset{PS}{\sim}$ on $\mathfrak{S}_n$ that is analogous to the Knuth relation $213 \overset{RSK}{\sim} 231$. (Recall that the Knuth relations describe when two permutations have the same "insertion tableau" P under RSK; see Sagan [20].)

In Section 3 we then explicitly describe a bijection between $\mathfrak{S}_n$ and certain pairs of pile configurations having the same shape (a composition γ of n, denoted $\gamma \models n$). Since there are many more possible pile configurations than standard Young tableaux (the former are enumerated by Bell numbers; see Theorem 4.1), it is necessary to specify which pairs are possible; this turns out to involve the same patterns as the other Knuth relation $312 \overset{RSK}{\sim} 132$ (see Definition 3.8). Moreover, this bijection shares the same Schützenberger symmetry property as RSK, and so we can immediately characterize a certain collection of pile configurations that are in bijection with the set of involutions $\mathfrak{I}_n$ (as well as with the set $\mathfrak{T}_n$ of standard Young tableaux).

In Section 4 we conclude by using the equivalence relation $\overset{PS}{\sim}$ to characterize and enumerate the set $S_n(3\text{-}\bar{1}\text{-}42)$ of permutations avoiding the barred (generalized) permutation pattern 3-$\bar{1}$-42. Such permutations avoid the pattern 2-31 unless it is contained in a 3-1-42 pattern. (See Sections 2.2 and 4 for the appropriate definitions.)

Another interesting property of RSK is that, given $\sigma \in \mathfrak{S}_n$, the number of boxes in the top row of the "insertion tableau" $P(\sigma)$ is exactly the length of the longest increasing subsequence in σ. (This was first proven by Schensted [21] but is now a special case of Greene's Theorem [11]). Due to the similarity between the Schensted Insertion Algorithm and Algorithm 1.1, it is clear that the cards atop the piles when Patience Sorting terminates will be exactly the elements in the top row of $P(\sigma)$. Thus, the number of piles formed under Patience Sorting is also equal to the length of the longest increasing subsequence in σ, and so one can apply the recent but now highly celebrated Baik-Deift-Johansson Theorem [3] in order to describe the asymptotic distribution for the number of piles (up to rescaling). Due to this deep connection between Patience Sorting and Probabilistic Combinatorics, it has been suggested (see, e.g., [13], [14] and [18]; cf. [7]) that studying generalizations of Patience Sorting might be the key to tackling certain open problems that can be viewed from the standpoint of Random Matrix Theory—the most notable being the Riemann Hypothesis.

At the same time, there is a lot more to Patience Sorting than just resembling the RSK Algorithm for permutations. E.g., after applying Algorithm 1.1 to a deck of cards, it is easy to recollect each card in ascending order from amongst the current top cards of the piles (and thus complete A. S. C. Ross' card sorting algorithm). While this is not necessarily the fastest sorting algorithm one can apply to a deck of cards, the *patience* in *Patience Sorting* is not intended to describe a prerequisite for its use. Instead it refers to how pile formation in Algorithm 1.1 resembles the way in which one places cards into piles when playing the popular single-person card game *Klondike Solitaire*, which is often called *Patience* in the UK. This is more than a coincidence, though, as Algorithm 1.1 also happens to be an optimal strategy (in the sense of forming as few piles as possible; see [1] for a proof) when playing an idealized model of Klondike Solitaire known as *Floyd's Game*:

Game 1.2 (Floyd's Game)**.** Given a shuffled deck of cards $c_1, c_2, \ldots, c_n$,

- Place the first card c_1 from the deck into a pile r_1 by itself.
- Then, for each card c_i ($i = 2, \ldots, n$), either
  - put c_i into a new pile by itself or
  - play c_i on top of any pile whose current top card is larger than c_i.
- The object of the game is to end with as few piles as possible.

In other words, the cards are played one at a time according to the order they appear in the deck so that piles are created in much the same way they are formed under Patience Sorting. According to [1], Floyd's Game was developed independently of Mallow's work and originated in unpublished correspondence between computer scientists Bob Floyd and Donald Knuth during 1964.

Note that unlike Klondike Solitaire, there is a known strategy (Algorithm 1.1) for Floyd's Game under which one will always win. In fact, Klondike Solitaire—though so popular that it has come pre-installed on the vast majority of personal computers shipped since 1989—is very poorly understood mathematically. (Recent progress, however, has been made in developing an optimal strategy for a version called *thoughtful solitaire* [26].) As such, Persi Diaconis ([1] and private communication with the second author) has suggested that a deeper understanding of Patience Sorting and its generalization would undoubtedly help in developing a better mathematical model for analyzing Klondike Solitaire.

2. Pile Configurations Coming from Patience Sorting

2.1. Pile Configurations, Shadow Diagrams, and Reverse Patience Words.

We begin by explicitly characterizing the pile configurations that result from applying Patience Sorting (Algorithm 1.1) to a permutation:

Lemma 2.1. *Let $\sigma \in \mathfrak{S}_n$ be a permutation and $R(\sigma) = \{r_1, r_2, \ldots, r_k\}$ be the pile configuration associated to σ under Algorithm 1.1. Then $R(\sigma)$ is a partition of the* *set* $[n] = \{1, 2, \ldots, n\}$ *such that denoting* $r_j = \{r_{j1} > r_{j2} > \cdots > r_{js_j}\}$,

$$r_{js_j} < r_{is_i} \quad \text{if} \quad j < i \tag{2.1}$$

Moreover, for every set partition $T = \{t_1, t_2, \ldots, t_k\}$ *satisfying Equation (2.1), there is a permutation* $\tau \in \mathfrak{S}_n$ *such that* $R(\tau) = T$.

Proof. Given a pile configuration $R(\sigma) = \{r_1, r_2, \ldots, r_k\}$, suppose that for some pair of indices $i, j \in [k]$ we have $j < i$ but $r_{js_j} > r_{is_i}$. Then r_{is_i} was put atop pile r_i when pile r_j had top card $d_j \geq r_{js_j}$ so that $d_j > r_{is_i}$. However, it then follows that, under Algorithm 1.1, the card r_{is_i} would actually then have been placed atop either pile r_j or a pile to the left of r_j instead of atop pile r_i. The resulting contradiction implies that $r_{js_j} < r_{is_i}$ for each $j < i$.

Conversely, let $T = \{t_1, t_2, \ldots, t_k\}$ be any set partition of $[n]$ with $t_j = \{t_{j1} > t_{j2} > \cdots > t_{js_j}\}$ for every $j \in [k]$ and $t_{js_j} < t_{is_i}$ for all pairs $j < i$. Consider then the permutation

$$\tau = t_{11}t_{12}\ldots t_{1s_1}\, t_{21}t_{22}\ldots t_{2s_2}\, \ldots\, t_{k1}t_{k2}\ldots t_{ks_k} \in \mathfrak{S}_n.$$

We show that $R(\tau) = T$: Given any $j \in [k]$ and any $i > j$, each $m \in [s_i]$ satisfies $t_{js_j} < t_{is_i} < t_{im}$, and so no value to the right of t_{js_j} in τ will be placed on top of t_{js_j} in $R(\tau)$. Since $t_{11} < t_{12} < \cdots < t_{1s_1}$, all these entries must be placed in the leftmost (a.k.a. first) pile $r_1 \in R(\tau)$, so the first pile of $R(\tau)$ is $r_1 = t_1$. Now suppose that the entries t_1 through t_{js_j} of τ have been placed in piles $t_1, t_2, \ldots, t_j$. Then the entries $t_{j+1,1} > t_{j+1,2} > \cdots > t_{j+1,s_{j+1}}$ cannot be placed atop any t_{is_i} for $i \leq j$, so we must form at least one new pile, say t_{j+1}. However, since these entries occur in decreasing order in τ, these cards will then all be placed in the $(j+1)^{\text{st}}$ pile t_{j+1}. Moreover, no entry to the right of $t_{j+1,s_{j+1}}$ can be placed in pile t_{j+1}, so the $(j+1)^{\text{st}}$ pile of $R(\tau)$ is $r_{j+1} = t_{j+1}$. Therefore, we obtain $R(\sigma) = T$ by induction. $\square$

We will often express a pile configuration R with its constituent piles $r_1, r_2, \ldots, r_k$ written vertically and bottom-justified with respect to the largest value r_{j1} in each pile r_j. This motivates the following definition (which reverses the so-called "far-eastern reading"):

Definition 2.2. The *reverse patience word* $RPW(R)$ for a pile configuration R is the permutation formed by concatenating the piles $r_1, r_2, \ldots, r_k$ together with each pile r_j written in decreasing order (i.e., read from bottom to top in order from left to right). In the notation of Lemma 2.1,

$$RPW(R) = r_{11}r_{12}\ldots r_{1s_1}\; r_{21}r_{22}\ldots r_{2s_2} \quad \ldots \quad r_{k1}r_{k2}\ldots r_{ks_k}.$$

Example 2.3. The pile configuration $R = \{\{6 > 4 > 1\}, \{5 > 2\}, \{8 > 7 > 3\}\}$ is represented by the piles

$$\begin{matrix} 1 & & 3 \\ 4 & 2 & 7 \\ 6 & 5 & 8 \end{matrix}$$

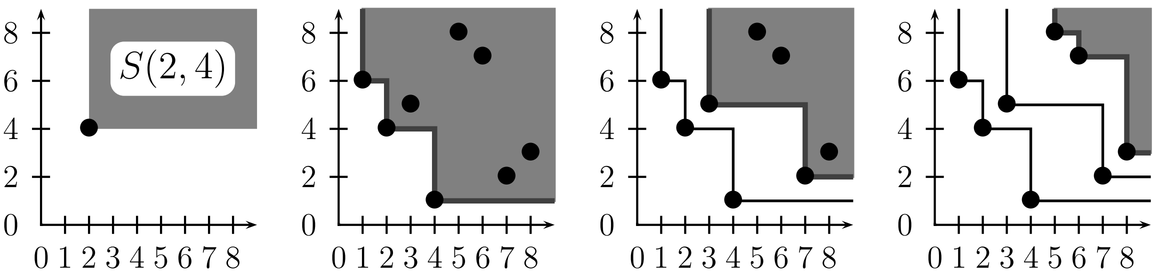

(a) The Shadow $S_{NE}(2,4)$.

(b) Shadowline $L_1(64518723)$.

(c) Shadowline $L_2(64518723)$.

(d) Shadowline $L_3(64518723)$.

FIGURE 2.1. Examples of Northeast Shadow and Shadowline Constructions

and has the reverse patience word $RPW(R) = 64152873$. Moreover, note that as in the proof of Lemma 2.1, $R(RPW(R)) = R(64152873) = R$.

The following Lemma should now be clear from the above definitions and example:

Lemma 2.4. *Given a permutation $\sigma \in \mathfrak{S}_n$, $R(RPW(R(\sigma))) = R(\sigma)$.*

Proof. From the proof of Lemma 2.1, we have that if $\tau = RPW(T)$, where T is any set partition of $[n]$ satisfying condition (2.1), then $T = R(\tau)$. Thus, $R(RPW(T)) = T$ for any partition T of $[n]$ satisfying (2.1), in particular for any $T = R(\sigma)$ where $\sigma \in \mathfrak{S}_n$. □

At the same time, it is also clear that in general there will be many permutations $\sigma, \tau \in \mathfrak{S}_n$ for which $R(\sigma) = R(\tau)$. In Section 2.2 below we characterize when two permutations have the same pile configuration, and we will denote this equivalence relation by $\sigma \overset{PS}{\sim} \tau$. Moreover, we will also see that the reverse patience word $RPW(R(\sigma))$ is the most natural representative for the equivalence class generated by a given permutation σ.

We close this section by giving an alternate characterization for pile configurations in terms of the *shadow diagram* construction that Viennot [24] introduced in the context of studying the RSK Algorithm for permutations.

Definition 2.5. Given a lattice point $(m,n) \in \mathbb{Z}^2$, we define the *(northeast) shadow* of (m,n) to be the quarter space $S(m,n) = \{(x,y) \in \mathbb{R}^2 \mid x \geq m,\ y \geq n\}$.

See Figure 2.1(a) for an example of a point's shadow.

The most important use of shadows is in building shadowlines:

Definition 2.6. Given lattice points $(m_1,n_1),(m_2,n_2),\ldots,(m_k,n_k) \in \mathbb{Z}^2$, we define their *(northeast) shadowline* to be the boundary of the region formed from the union of the shadows $S(m_1,n_1), S(m_2,n_2),\ldots,S(m_k,n_k)$.

In particular, we wish to associate to each permutation a certain collection of shadowlines (as illustrated in Figure 2.1(b)–(d)):

Definition 2.7. Given a permutation $\sigma = \sigma_1\sigma_2\cdots\sigma_n \in \mathfrak{S}_n$, the *(northeast) shadow diagram* $D(\sigma)$ of σ consists of the shadowlines $L_1(\sigma), L_2(\sigma), \ldots, L_k(\sigma)$ formed as follows:

- $L_1(\sigma)$ is the shadowline for the lattice points $\{(1,\sigma_1),(2,\sigma_2),\ldots,(n,\sigma_n)\}$.
- While at least one of the points $(1,\sigma_1),(2,\sigma_2),\ldots,(n,\sigma_n)$ is not contained in the shadowlines $L_1(\sigma), L_2(\sigma), \ldots, L_j(\sigma)$, define $L_{j+1}(\sigma)$ to be the shadowline for the points
$$\{(i,\sigma_i) \mid (i,\sigma_i) \notin \bigcup_{k=1}^{j} L_k(\sigma)\}.$$

In other words, we define the shadow diagram $D(\sigma) = \{L_1(\sigma), L_2(\sigma), \ldots, L_k(\sigma)\}$ inductively with $L_1(\sigma)$ the shadowline for the diagram $\{(1,\sigma_1),(2,\sigma_2),\ldots,(n,\sigma_n)\}$ of the permutation $\sigma \in \mathfrak{S}_n$. Then we ignore the points whose shadows were actually used in building $L_1(\sigma)$ and define $L_2(\sigma)$ to be the shadowline of the resulting subset of the permutation diagram. We then build $L_3(\sigma)$ as the shadowline for the points not yet used in constructing both $L_1(\sigma)$ and $L_2(\sigma)$, and this process continues until all points in the permutation diagram are exhausted.

One of the most basic properties of the shadow diagram for a permutation σ is that it encodes the top row of the insertion tableau $P(\sigma)$ (resp. recording tableau $Q(\sigma)$) as the smallest ordinates (resp. smallest abscissae) of all points belonging to the constituent shadowlines $L_1(\sigma), L_2(\sigma), \ldots, L_k(\sigma)$. (A proof of this can be found in Sagan [20].) In particular, this means that if σ has pile configuration $R(\sigma) = \{r_1, r_2, \ldots, r_m\}$, then $m = k$ since the number of piles is equal to the length of the top row of $P(\sigma)$ (as both are the length of the longest increasing subsequence of σ). We can say even more about the relationship between $D(\sigma)$ and $R(\sigma)$ when both are viewed in terms of left-to-right minima subsequences (a.k.a. basic subsequences or records):

Definition 2.8. Let $\pi = \pi_1\pi_2\cdots\pi_l$ be a partial permutation on $[n] = \{1, 2, \ldots, n\}$. Then the *left-to-right minima subsequence* of π consists of those components π_j of π such that $\pi_j = \min\{\pi_i \mid 1 \le i \le j\}$.

We then inductively define the left-to-right minima subsequences $s_1, s_2, \ldots, s_k$ of a permutation σ by taking s_1 to be the left-to-right minima subsequence for σ itself and then each subsequent subsequence s_i to be the left-to-right minima subsequence for the partial permutation obtained by removing the elements of $s_1, s_2, \ldots, s_{i-1}$ from σ. The subsequence s_j is called the j^{th} *left-to-right minima subsequence* of σ.

Lemma 2.9. *Suppose $\sigma \in \mathfrak{S}_n$ has shadow diagram $D(\sigma) = \{L_1(\sigma), L_2(\sigma), \ldots, L_k(\sigma)\}$. Then the ordinates of the southwest corners of L_j are exactly the cards in the j^{th} pile*

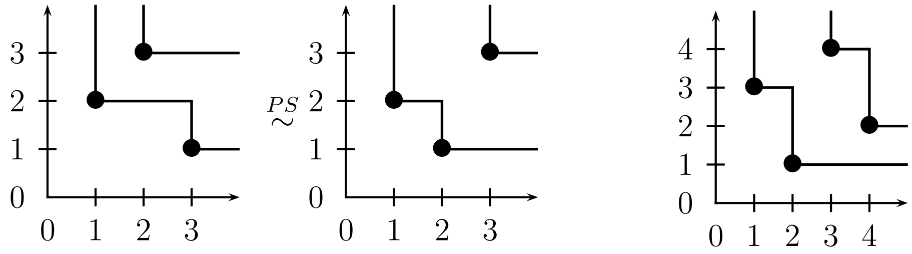

(a) "Stretching" shadowlines effects $231 \overset{PS}{\sim} 213$. Thus, $\widetilde{231} = \{231, 213\}$.

(b) No "stretching" can interchange "4" and "2".

FIGURE 2.2. Examples of patience sorting equivalence and non-equivalence

$r_j \in R(\sigma)$ formed by applying Patience Sorting (Algorithm 1.1) to σ. In other words, the j^{th} pile r_j contains exactly the elements of the j^{th} left-to-right minima subsequence of σ.

Proof. The i^{th} basic subsequence s_i of σ consists of those elements σ_t that appear at the end of an increasing subsequence of length i but not at the end of an increasing subsequence of length $i+1$. Thus, since each element added to a pile must be smaller than all other elements already in the pile, $s_1 = r_1$. It then follows similarly by induction that $s_i = r_i$ for $i = 2, \ldots, k$.

The proof that the ordinates of the southwest corners of the shadowlines L_i are also exactly the elements of the left-to-right minima subsequences s_i is similar. $\square$

Lemma 2.9 gives a particularly nice correspondence between the piles formed under Patience Sorting and the shadowlines that constitute the shadow diagram of a permutation. In particular, we have that forming $RPW(R(\sigma))$ essentially amounts to sorting $\sigma \in \mathfrak{S}_n$ into its left-to-right minima subsequences.

We will rely heavily upon this correspondence in the sections below.

2.2. **Permutations Having Equivalent Pile Configurations.** In this section we characterize the following equivalence relation:

Definition 2.10. Two permutations $\sigma, \tau \in \mathfrak{S}_n$ are said to be *patience sorting equivalent*, written $\sigma \overset{PS}{\sim} \tau$, if they have the same pile configuration $R(\sigma) = R(\tau)$ under Algorithm 1.1. We denote the equivalence class generated by σ as $\widetilde{\sigma}$.

By Lemma 2.9 in Section 2.1 above, the pile configurations $R(\sigma)$ and $R(\tau)$ correspond to certain shadow diagrams. Thus, it should be intuitively clear that preserving a given pile configuration is equivalent to preserving the ordinates for the southwest corners of the shadowlines. In particular, this means that we are limited to horizontally "stretching" shadowlines up to the point of not allowing them to cross as is illustrated in Figure 2.2 and the following examples.

Example 2.11. The only non-singleton patience sorting equivalence class for $\mathfrak{S}_3$ consists of $\widetilde{231} = \{231, 213\}$. We illustrate $231 \overset{PS}{\sim} 213$ in Figure 2.2(a).

Notice that the actual values of the elements interchanged in Example 2.11 are immaterial so long as they have the same relative magnitudes as the literal values in the word 231. (I.e., they have to be order-isomorphic.) Moreover, it should also be clear that any value greater than the element playing the role of "1" can be inserted between the elements playing the roles of "2" and "3" without affecting the ability to interchange the "1" and "3" elements. Problems with this interchange only start to arise when a value smaller than the element playing the role of "1" is inserted between the elements playing the roles of "2" and "3". We can formally describe this idea using the language of generalized permutation patterns (as was recently defined in [2]; cf. [4]).

Definition 2.12. Let $\sigma = \sigma_1\sigma_2 \cdots \sigma_n \in \mathfrak{S}_n$ and $\tau \in \mathfrak{S}_m$ for $m \leq n$. Then we say that σ *contains the (classical) pattern* τ if there exists a subsequence $\sigma_{i_1}, \sigma_{i_2}, \ldots, \sigma_{i_m}$ of σ (meaning $i_1 < i_2 < \cdots < i_m$) such that the word $\sigma_{i_1}\sigma_{i_2} \ldots \sigma_{i_m}$ is order-isomorphic to τ.

If σ does not contain τ, then we say that σ *avoids* the pattern τ, and we denote by $S_n(\tau)$ the subset of the symmetric group $\mathfrak{S}_n$ that avoids τ.

Note that the elements in the subsequence $\sigma_{i_1}, \sigma_{i_2}, \ldots, \sigma_{i_m}$ are not required to be contiguous in σ. In a *generalized pattern* one assumes that every element in the subsequence must be taken contiguously unless a dash is inserted in the pattern τ between elements that are not required to be contiguous in σ. (A generalized patterns with no dashes is sometimes called a *segment* or a *consecutive pattern.*)

Example 2.13.

(1) Notice that 2431 contains exactly one instance of a 2-31 pattern as the bold underlined subsequence $\underline{\mathbf{2}}4\underline{\mathbf{31}}$. (Conversely, $\underline{\mathbf{24}}3\underline{\mathbf{1}}$ is an instance of 23-1 but not of 2-31.) Moreover, it is clear that $2431 \overset{PS}{\sim} 2413$.
(2) Even though 3142 contains a 2-31 pattern (as the subsequence $\underline{\mathbf{3}}1\underline{\mathbf{42}}$), we cannot interchange "4" and "2", and so $R(3142) \neq R(3124)$. As illustrated in Figure 2.2(b), this is because "4" and "2" are on the same shadowline.

We can now state our main result on patience sorting equivalence:

Theorem 2.14. *Let $\sigma, \tau \in \mathfrak{S}_n$. Then σ and τ have the same pile configuration $R(\sigma) = R(\tau)$ under Algorithm 1.1 (so that $\sigma \overset{PS}{\sim} \tau$) if and only if there exists a sequence of* 2-31 *to* 2-13 *interchanges (with no* 2-31 *pattern contained in a* 3-1-42 *pattern) that transform σ into τ.*

In other words, $\overset{PS}{\sim}$ is the transitive closure of such interchanges.

Proof. By Lemma 2.9 it suffices to show that 2-31 to 2-13 interchanges (with no 2-31 pattern contained in a 3-1-42 pattern), preserve the left-to-right minima subsequences $s_1, s_2, \ldots, s_k$ of σ. This amounts to showing by induction that such interchanges suffice to transform σ into $RPW(R(\sigma))$ via the sequence of pattern interchanges

$$\sigma = \sigma^{(0)} \rightsquigarrow \sigma^{(1)} \rightsquigarrow \sigma^{(2)} \rightsquigarrow \cdots \rightsquigarrow \sigma^{(\ell)} = RPW(R(\sigma))$$

where each $\sigma^{(i)} \overset{PS}{\sim} \sigma^{(i+1)}$.

Let (a, b, c) be a subsequence of σ_i that is an instance of 2-31 not contained in an instance of 3-1-42. Then $c < a < b$, and there is no d between a and b in σ such that $d < c$. Clearly, we do not lose any increasing subsequences of $\sigma^{(i)}$ by interchanging b and c. Moreover, the only new increasing subsequences $\sigma^{(i+1)}$ created by this interchange are those that end with the subsequence (c, b). If such an increasing subsequence contains other terms, say $d_1 < d_2 < \cdots < d_m < c < b$, then d_m must be to the left of a in σ. But then $\sigma^{(i)}$ contains an increasing sequence $d_1 < d_2 < \cdots < d_m < a < b$ with the same number of terms. Hence, interchanging b and c does not create any longer increasing subsequences with the same final term, so $\sigma^{(i)} \overset{PS}{\sim} \sigma^{(i+1)}$.

Let $r_{11} > r_{12} > \cdots > r_{1s_1}$ be the subsequence of left-to-right minima of σ; from the proof of Lemma 2.9, these are the entries that form the leftmost pile r_1 of $R(\sigma)$. Now suppose that, for some $j < s_1$, there is an entry between r_{1j} and $r_{1,j+1}$, and let b be the entry immediately preceding $r_{1,j+1}$. Then $b > r_{1j}$ so that $(r_{1j}, b, r_{1,j+1})$ is an instance of 2-31. On the other hand, $r_{1,j+1}$ is the leftmost entry of σ that is less than r_{1j}, so no entry $d < r_{1,j+1}$ may occur between r_{1j} and b. Hence, $(r_{1j}, b, r_{1,j+1})$ is not an instance of 3-1-42, and so interchanging $r_{1,j+1}$ and b will not change the pile configuration. We may repeat this until there are no entries between consecutive left-to-right minima of σ. We thus obtain $\sigma' = r_{11}r_{12}\ldots r_{1s_1}\sigma''$, where σ'' is obtained by deleting $r_{11}, r_{12}, \ldots, r_{1s_1}$ from σ. Since no instance of 2-31 may start with r_{1s_1} and since any instance (r_{1i}, a, b) of 2-31 is part of an instance (r_{1i}, r_{1s_1}, a, b) of 3-1-42, no further interchanges will involve any r_{1i}. Thus, by induction, we can now apply the same procedure to σ'', etc., to ultimately obtain $r_{11}r_{12}\ldots r_{1s_1}r_{21}r_{22}\ldots r_{2s_2} \ \ldots \ r_{k1}r_{k2}\ldots r_{ks_k} = RPW(R(\sigma))$. □

Remark 2.15. It follows from Theorem 2.14 that Examples 2.11 and 2.13(2) sufficiently characterize when two permutations yield the same pile configurations under Patience Sorting. However, it is worth pointing out that these examples also begin to illustrate how one can build an infinite sequence of generalized permutation patterns (all of them containing either 2-13 or 2-31) with the following property: an interchange of the pattern 2-13 with the pattern 2-31 is allowed within an odd-length pattern in this sequence unless the elements used to form the odd-length pattern can also be used as part of a longer even-length pattern in this sequence.

Example 2.16. Even though the permutation 34152 contains a 3-1-42 pattern in the suffix "4152", one can still directly interchange the "5" and the "2" because of the

"3" prefix (or via the following sequence of interchanges: $34152 \rightsquigarrow 31452 \rightsquigarrow 31425 \rightsquigarrow 34125$).

3. Bijectively Extending Patience Sorting to "Stable Pairs" of Pile Configurations

3.1. **The Extended Patience Sorting Algorithm.** Recall from Section 1 that Patience Sorting (Algorithm 1.1) can be viewed as an iterated, non-recursive form of the Schensted Insertion Algorithm for inserting a value into the top row of a Young tableau. In this section we extend the Patience Sorting construction so that it becomes a full non-recursive analog of the RSK Algorithm for permutations. In particular, we mimic the RSK recording tableau construction so that "recording piles" are formed while assembling the usual pile configuration under Patience Sorting (which by analogy to RSK we will similarly now call "insertion piles").

Algorithm 3.1 (Extended Patience Sorting Algorithm)**.** Given a shuffled deck of cards $\sigma = c_1 c_2 \cdots c_n$, inductively build *insertion piles* $R = R(\sigma) = \{r_1, r_2, \ldots, r_m\}$ and *recording piles* $S = S(\sigma) = \{s_1, s_2, \ldots, s_m\}$ as follows:

- Place the first card c_1 from the deck into a pile r_1 by itself and set $s_1 = \{1\}$.
- For each remaining card c_i ($i = 2, \ldots, n$), consider the cards $d_1, d_2, \ldots, d_k$ atop the piles $r_1, r_2, \ldots, r_k$ that have already been formed.
  - If $c_i > \max\{d_1, d_2, \ldots, d_k\}$, then put c_i into a new pile r_{k+1} by itself and set $s_{k+1} = \{i\}$.
  - Otherwise, find the left-most card d_j that is larger than c_i and put the card c_i atop pile r_j while simultaneously putting i at the bottom of pile s_j.

We call the pile configuration pairs that result from Algorithm 3.1 *stable pairs* and give a characterization for them in Section 3.2 below. Note that the pile configurations that comprise a resulting stable pair must have the same "shape", which we define as follows:

Definition 3.2. Given a pile configuration $R = \{r_1, r_2, \ldots, r_m\}$ on n cards, we call the composition $\gamma = (|r_1|, |r_2|, \ldots, |r_m|)$ of n the *shape* of R and denote this by $\mathrm{sh}(R) = \gamma \models n$.

Example 3.3. Let $\sigma = 64518723 \in \mathfrak{S}_8$. Then according to Algorithm 3.1 we simultaneously form the following pile configurations with shape $\mathrm{sh}(R(\sigma)) = \mathrm{sh}(S(\sigma)) = (3, 2, 3) \models 8$.

	insertion piles	recording piles		insertion piles	recording piles
Form a new pile with **6**:	$\begin{array}{c} \\ \\ \mathbf{6} \end{array}$	$\begin{array}{c} \\ \\ \mathbf{1} \end{array}$	Then play the **4** on it:	$\begin{array}{c} \\ \mathbf{4} \\ 6 \end{array}$	$\begin{array}{c} \\ 1 \\ \mathbf{2} \end{array}$
Form a new pile with **5**:	$\begin{array}{cc} & \\ 4 & \\ 6 & \mathbf{5} \end{array}$	$\begin{array}{cc} & \\ 1 & \\ 2 & \mathbf{3} \end{array}$	Add the **1** to left pile:	$\begin{array}{cc} \mathbf{1} & \\ 4 & \\ 6 & 5 \end{array}$	$\begin{array}{cc} 1 & \\ 2 & \\ \mathbf{4} & 3 \end{array}$
Form a new pile with **8**:	$\begin{array}{ccc} 1 & & \\ 4 & & \\ 6 & 5 & \mathbf{8} \end{array}$	$\begin{array}{ccc} 1 & & \\ 2 & & \\ 4 & 3 & \mathbf{5} \end{array}$	Then play the **7** on it:	$\begin{array}{ccc} 1 & & \\ 4 & & \mathbf{7} \\ 6 & 5 & 8 \end{array}$	$\begin{array}{ccc} 1 & & \\ 2 & & 5 \\ 4 & 3 & \mathbf{6} \end{array}$
Add the **2** to middle pile:	$\begin{array}{ccc} 1 & & \\ 4 & \mathbf{2} & 7 \\ 6 & 5 & 8 \end{array}$	$\begin{array}{ccc} 1 & & \\ 2 & 3 & 5 \\ 4 & \mathbf{7} & 6 \end{array}$	Add the **3** to right pile:	$\begin{array}{ccc} 1 & & \mathbf{3} \\ 4 & 2 & 7 \\ 6 & 5 & 8 \end{array}$	$\begin{array}{ccc} 1 & & 5 \\ 2 & 3 & 6 \\ 4 & 7 & \mathbf{8} \end{array}$

The idea behind Algorithm 3.1 is that we are using the recording piles $S(\sigma)$ to implicitly label the order in which the elements of the permutation σ are added to the insertion piles $R(\sigma)$. It is clear that this information then allows us to uniquely reconstruct σ by reversing the order in which the cards were played. However, even though reversing the Extended Patience Sorting Algorithm is much easier than reversing the RSK Algorithm through recursive "reverse row bumping", the trade-off is that the stable pairs that result from the former are not independent whereas the tableau pairs generated by RSK are completely independent (up to shape).

That $S(\sigma) = \{s_1, s_2, \ldots, s_m\}$ records the order of the cards being added to the insertion piles is made clear if we instead add cards to the tops of new piles s'_j in Algorithm 3.1 rather than to the bottoms of the piles s_j. This yields modified recording piles $S'(\sigma)$ from which each original recording pile $s_j \in S(\sigma)$ can be recovered by simply reflecting the corresponding pile s'_j vertically.

Example 3.4. As in Example 3.3 above, let $\sigma = 64518723 \in \mathfrak{S}_8$. Then $R(\sigma)$ is formed as before and

$$S'(\sigma) = \begin{array}{ccc} 4 & & 8 \\ 2 & 7 & 6 \\ 1 & 3 & 5 \end{array} \overset{\text{reflect}}{\dashrightarrow} \begin{array}{ccc} 1 & & 5 \\ 2 & 3 & 6 \\ 4 & 7 & 8 \end{array} = S(\sigma)$$

We are now in a position to prove that the Extended Patience Sorting Algorithm has the same form of symmetry as the RSK Algorithm has for permutations.

Proposition 3.5. *Let $(R(\sigma), S(\sigma))$ be the insertion and recording piles, respectively, formed by applying Algorithm 3.1 to $\sigma \in \mathfrak{S}_n$. Then reversing Algorithm 3.1 for the pair $(S(\sigma), R(\sigma))$ yields the inverse permutation σ^{-1}.*

Proof. Construct $S'(\sigma)$ from $S(\sigma)$ as discussed above, and form the n ordered pairs (r_{ij}, s'_{ij}) where i indexes the individual piles and j the cards in the i^{th} piles. Then these n points correspond to the diagram of a permutation $\tau \in \mathfrak{S}_n$. However, since reflecting these points through the line $y = x$ yields the diagram for σ, it follows that $\tau = \sigma^{-1}$. $\square$

Proposition 3.5 suggests that Algorithm 3.1 is the right generalization of Algorithm 1.1 since we obtain the same symmetry property as for RSK. At the same time, though, since there are many more possible pile configurations than standard Young tableau (as we'll show in Section 4 below), not every ordered pair of pile configurations with the same shape will result from Algorithm 3.1. Thus, it is necessary to first characterize the "stable pairs" that result from applying Extended Patience Sorting to a permutation. We do this in Section 3.2.

3.2. **Characterizing "Stable Pairs" of Pile Configurations and Pile Configurations for Involutions.** Based upon Proposition 3.5 above, there is a bijection between involutions and certain pile configurations. We will describe this bijection as a corollary to a more general characterization for the "stable pairs" of pile configurations that can result from applying the Extended Patience Sorting Algorithm to a permutation.

The following example, though very small, illustrates the most generic behavior that must be avoided in constructing stable pairs. As in Section 3.1 above, we denote by S' the "reverse pile configuration" of S (which has all piles listed in reverse order).

Example 3.6. Even though the pile configuration $R = \{\{3 > 1\}, \{2\}\}$ cannot result as the insertion piles given by an involution under the Extended Patience Sorting Algorithm, we can still try to look at the shadow diagram for the pre-image of the pair (R, R) under Algorithm 3.1:

$$R \;=\; \begin{array}{cc} 1 & \\ 3 & 2 \end{array} \quad \text{and} \quad S' \;=\; \begin{array}{cc} 3 & \\ 1 & 2 \end{array} \quad \Longrightarrow$$

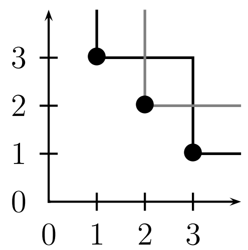

Note that there are two competing constructions here. On the one hand, we have the diagram $\{(1, 3), (2, 2), (3, 1)\}$ of a permutation given by the entries in the pile configurations. (In particular, the values in R specify the ordinates and the values in the corresponding boxes of S' the abscissae.) On the other hand, the piles in R also specify shadowlines with respect to this permutation diagram. Here the pair (R, S) of pile configurations is "unstable" because their combination yields crossing shadowlines—which is clearly not allowed.

Similar considerations lead to avoiding the crossings of the form

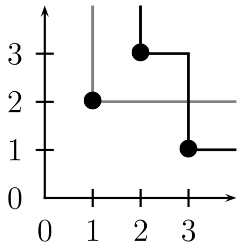

and

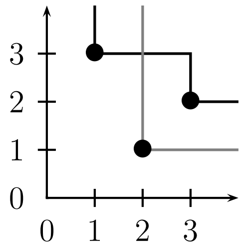

Note also that these latter two crossings can also be used together to build something like the first crossing but with "extra" elements on the boundary of the polygon formed:

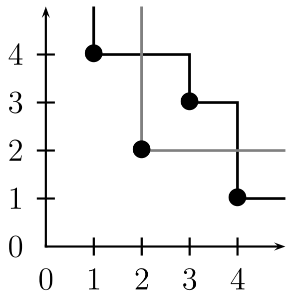

We can now make the following important definitions:

Definition 3.7. Given a composition γ of n (denoted $\gamma \models n$), we define $\mathfrak{P}_\gamma(n)$ to be the set of all pile configurations R having shape $\mathrm{sh}(R) = \gamma$ and put

$$\mathfrak{P}(n) = \bigcup_{\gamma \models n} \mathfrak{P}_\gamma(n).$$

Definition 3.8. Define the set $\Sigma(n) \subset \mathfrak{P}(n) \times \mathfrak{P}(n)$ to consist of all ordered pairs (R, S) with $\mathrm{sh}(R) = \mathrm{sh}(S)$ such that the pair $(RPW(R), RPW(S'))$ avoids simultaneous occurrences of pairs of patterns $(31\text{-}2, 13\text{-}2)$, $(31\text{-}2, 32\text{-}1)$ and $(32\text{-}1, 13\text{-}2)$ at the same positions in $RPW(R)$ and $RPW(S')$.

In other words, if $RPW(R)$ contains an occurrence of the first pattern in any of the above pairs, then $RPW(S')$ cannot contain an occurrence at the same positions of the second pattern in the same pair, and vice versa. Note that Definition 3.8 characterizes "stable pairs" of pile configurations (R, S) by forcing R and S to avoid certain sub-pile pattern pairs. As in Example 3.6, we are characterizing when the induced shadowlines cross.

Theorem 3.9. *Extended Patience Sorting (Algorithm 3.1) gives a bijection between the symmetric group $\mathfrak{S}_n$ and the "stable pairs" set $\Sigma(n)$ given in Definition 3.8 above.*

Proof. We will prove that for any stable pair $(R, S) \in \Sigma(n)$ and any permutation $\sigma \in \mathfrak{S}_n$,

$$\sigma = \begin{pmatrix} RPW(S') \\ RPW(R) \end{pmatrix} \text{ (in the two-line notation) } \iff (R, S) = (R(\sigma), S(\sigma)).$$

Clearly, if $(R, S) = (R(\sigma), S(\sigma))$ for some $\sigma \in \mathfrak{S}_n$, then $\sigma = \begin{pmatrix} RPW(S') \\ RPW(R) \end{pmatrix}$, so we only need to prove that $(R, S) \in \Sigma(n)$. Suppose that $(R, S) \notin \Sigma(n)$; then $RPW(R)$ and $RPW(S')$ contain instances of one of the three forbidden pairs.

Suppose $RPW(R)$ contains an occurrence (r_3, r_1, r_2) of 31-2, and $RPW(S')$ contains an occurrence (s'_1, s'_3, s'_2) of 13-2 at the same positions. Since $r_3 > r_1$ and r_3, r_1 are consecutive entries in $RPW(R)$, it follows that r_3 and r_1 must be in the same column $c_i(R)$ of R (in fact, r_1 is immediately on top of r_3). Since $r_1 < r_2$ and r_2 is to the right of r_1 in R, it follows that the column $c_j(R)$ of R containing r_2 must be to the right of the column containing r_1 atop r_3. Therefore, s'_2 must also be in a column $c_i(S')$ of S' to the right of the column $c_j(S')$ containing s'_3 on top of s'_1.

Consider the subpermutation τ of σ formed by deleting entries of $RPW(R)$ and $RPW(S')$ that are not in these two columns. Alternatively, let R_* and S'_* consist only of the columns $(c_i(R), c_j(R))$ of R and $(c_i(S'), c_j(S'))$ of S', respectively. Then

$$\tau = \begin{pmatrix} RPW(S'_*) \\ RPW(R_*) \end{pmatrix}.$$

Note that the values r_3 and r_1 in $c_i(S')$ are consecutive left-to-right minima of τ, whereas r_2 is not a left-to-right minimum of τ. Since $r_1 < r_2 < r_3$, it follows that r_2 cannot occur between r_1 and r_3 in τ. However, since $\begin{pmatrix} s'_1 & s'_3 & s'_2 \\ r_3 & r_1 & r_2 \end{pmatrix}$ is a subpermutation of τ and $s'_1 < s'_2 < s'_3$, it follows that r_2 does occur between r_1 and r_3, a contradiction.

A similar argument applies to the other two forbidden pairs, $(31\text{-}2, 32\text{-}1)$ and $(32\text{-}1, 13\text{-}2)$. The resulting contradiction in each case implies that we must have $(R, S) \in \Sigma(n)$.

Conversely, if $(R, S) \in \Sigma(n)$, then set $\sigma = \begin{pmatrix} RPW(S') \\ RPW(R) \end{pmatrix}$. We must show that $(R, S) = (R(\sigma), S(\sigma))$. This can be done by showing that $c_1(R)$ is the sequence of left-to-right minima of σ and $c_1(S')$ is the sequence of their positions, and then proceeding by induction on the length of σ.

We know that $c_1(R)$ is decreasing and $c_1(S')$ is increasing, with both columns of the same length. Moreover, $c_1(R)$ is the sequence of left-to-right minima of $RPW(R)$, and $c_1(S')$ is increasing, so the values of $c_1(R)$ are in decreasing order in σ. Note that the first term of $c_1(R)$ is the leftmost value in σ, and the last term of $c_1(R)$ is the least value of σ.

If every value in the first column of R is less than every value not in the first column of R, and every value in the first column of S' is less than every value not in the first column of S', then we are done. Now suppose this is not so. Then there exists a point (s', r) with $s' \in (RPW(S') \setminus c_1(S')$ and $r \in RPW(R) \setminus c_1(R))$ such that there are points (s'_1, r_1) and (s'_2, r_2), where $r_1, r_2 \in c_1(R)$ and $s'_1, s'_2 \in c_1(S'))$ are consecutive terms occupying the same positions in the first columns of R and S', for which either $r_1 > r > r_2$ or $s'_1 < s' < s'_2$. None of the three forbidden pairs may occur at positions $(s'_1, r_1), (s'_2, r_2), (s', r)$, in other words, given that one of $r_1 > r > r_2$ or $s'_1 < s' < s'_2$ must hold, we must avoid the following cases:

$$\begin{aligned} r_1 > r > r_2 \text{ and } s' < s'_1, &\qquad \text{i.e. an occurrence of } (31\text{-}2, 32\text{-}1), \\ r_1 > r > r_2 \text{ and } s'_1 < s' < s'_2, &\qquad \text{i.e. an occurrence of } (31\text{-}2, 13\text{-}2), \\ r < r_1 \text{ and } s'_1 < s' < s'_2, &\qquad \text{i.e. an occurrence of } (32\text{-}1, 13\text{-}2). \end{aligned}$$

Hence, $r_1 > r > r_2$ implies $s' > s'_2$, while $s'_1 < s' < s'_2$ implies $r > r_1$. Thus, r_2 is the first value of σ to the right of r_1 that is less than r_1, and all values between r_1 and r_2 are to the right of r_2. Therefore, if r_1 is a left-to-right minimum of σ, then r_2 is the next left-to-right minimum of σ. Since the first term of $c_1(R)$ is a left-to-right minimum of σ, it follows by induction that $c_1(R)$ is the sequence of left-to-right minima of σ, which occur at positions in $c_1(S')$. This finishes the proof. $\square$

Example 3.10. The pair of piles

$$(R, S) = \begin{pmatrix} 1 & & 3 & & 1 & & 5 \\ 4 & 2 & 7 & , & 2 & 3 & 6 \\ 6 & 5 & 8 & & 4 & 7 & 8 \end{pmatrix} \in \Sigma(8)$$

corresponds to the permutation

$$\sigma = \begin{pmatrix} RPW(S') \\ RPW(R) \end{pmatrix} = \begin{pmatrix} 1 & 2 & 4 & 3 & 7 & 5 & 6 & 8 \\ 6 & 4 & 1 & 5 & 2 & 8 & 7 & 3 \end{pmatrix} = 64518723 \in \mathfrak{S}_8.$$

Based upon the characterization of stable pairs given in Theorem 3.9 and the Schützenberger-type Symmetry Property proven in Proposition 3.5, we can immediately describe a bijection between involutions and certain pile configurations. In particular, these pile configurations must avoid simultaneously containing the symmetric sub-pile patterns corresponding to the patterns given in Definition 3.8.

This corresponds to the reverse patience word for a pile configuration simultaneously avoiding a symmetric pair of the generalized patterns 31-2 and 32-1. As such it is interesting to compare this construction to two results recently obtained by Claesson and Mansour [6]:

(1) The size of $S_n(3\text{-}12, 3\text{-}21)$ is equal to the number of involutions $|\mathfrak{I}_n|$ in $\mathfrak{S}_n$.
(2) The size of $S_n(31\text{-}2, 32\text{-}1)$ is 2^{n-1}.

The first result suggests that there should be a way to relate the result in Theorem 3.9 to simultaneous avoidance of the very similar patterns 3-12 and 3-21. The second result suggests that restricting to complete avoidance of all simultaneous occurrences of 31-2 and 32-1 will yield a natural bijection between $S_n(31\text{-}2, 32\text{-}1)$ and a subset $\mathfrak{N} \subset \mathfrak{P}(n)$ such that $\mathfrak{N} \cap \mathfrak{P}_\gamma(n)$ contains exactly one pile configuration of each shape γ. A natural family for this collection of pile configurations consists of what we call *non-crossing pile configurations*; namely, for the composition $\gamma = (\gamma_1, \gamma_2, \ldots, \gamma_k) \models n$,

$$\mathfrak{N} \cap \mathfrak{P}_\gamma(n) = \{\{\{\gamma_1 > \cdots > 1\}, \{\gamma_1 + \gamma_2 > \cdots > \gamma_1 + 1\}, \ldots, \{n > \cdots > n - \gamma_{k-1}\}\}\}$$

so that there are exactly 2^{n-1} such pile configurations. One can also show that $\mathfrak{N}$ is the image $R(S_n(3\text{-}1\text{-}2))$ of all permutations avoiding the classical pattern 3-1-2 under the Patience Sorting Algorithm.

4. Enumerating $S_n(3\text{-}\bar{1}\text{-}42)$

In this section we use the results from Section 2 to both enumerate and characterize the permutations that avoid the generalized permutation pattern 2-31 unless it is part of the generalized pattern 3-1-42. We call this restricted form of the generalized pattern 2-31 a *barred (generalized) permutation pattern* and denote it by $3\text{-}\bar{1}\text{-}42$. (This notation is due to J. West, et al., and first appeared in the study of two-stack sortable permutations [8, 9, 25].)

Theorem 4.1.

(1) *The set of permutations* $S_n(3\text{-}\bar{1}\text{-}42)$ *that avoids the pattern* $3\text{-}\bar{1}\text{-}42$ *is exactly the set* $RPW(R(\mathfrak{S}_n))$ *of reverse patience words obtainable from the symmetric group* $\mathfrak{S}_n$.

(2) *The size of* $S_n(3\text{-}\bar{1}\text{-}42)$ *is given by the* n^{th} *Bell number* B_n.

Proof.

(1) Let $\sigma \in S_n(3\text{-}\bar{1}\text{-}42)$. Then, for $i = 1, 2, \ldots, n-1$, define

$$\sigma_{m_i} = \min\{\sigma_j \mid i \leq j \leq n\}.$$

Since σ avoids $3\text{-}\bar{1}\text{-}42$, the subpermutation $\sigma_i \sigma_{i+1} \cdots \sigma_{m_i}$ must be a decreasing subsequence of σ. (Otherwise σ would necessarily contain a 2-31 pattern that is not part of a 3-1-42 pattern.) It follows that the left-to-right minima subsequences $s_1, s_2, \ldots, s_k$ of σ must be disjoint and satisfy Equation (2.1) so that the result follows by Lemmas 2.1 and 2.9.

(2) Recall that the Bell number B_n enumerates the set partitions of the set $[n] = \{1, 2, \ldots, n\}$. From Part (1), the elements of $S_n(3\text{-}\bar{1}\text{-}42)$ are in bijection with pile configurations. Thus, since pile configurations are themselves set partitions by Lemma 2.1, we need only show that every set partition is also a pile configuration. But this follows by ordering the components of a given set partition by their smallest element so that Equation (2.1) is satisfied. □

Remark 4.2. We conclude by remarking that even though the set $S_n(3\text{-}\bar{1}\text{-}42)$ is enumerated by the very well known Bell numbers, it cannot be described in a simpler way using classical pattern avoidance. This means that there does not exist a countable set of non-generalized (a.k.a. classical) permutation patterns $\tau_1, \tau_2, \ldots$ such that

$$S_n(3\text{-}\bar{1}\text{-}42) = S_n(\tau_1, \tau_2, \ldots) = \bigcap_{i \geq 1} S_n(\tau_i).$$

There are two very important reasons that this cannot happen:

First of all, the Bell numbers satisfy $\log B_n = n(\log n - \log\log n + O(1))$ and so exhibit superexponential growth. However, in light of the Stanley-Wilf ex-Conjecture (which was recently proven by Marcus and Tardos [17]), the set of permutations $S_n(\tau)$ avoiding any classical pattern τ can only grow at most exponentially in n.

On the other hand, the so-called *avoidance class* of permutations

$$Av(3\text{-}\bar{1}\text{-}42) = \bigcup_{n \geq 0} S_n(3\text{-}\bar{1}\text{-}42)$$

is not closed under taking order-isomorphic subpermutations, whereas it is easy to see that classes of permutations defined by classical pattern avoidance must be closed. (See Bóna [4], Chap. 5.) In particular, $3142 \in Av(3\text{-}\bar{1}\text{-}42)$ but $231 \notin Av(3\text{-}\bar{1}\text{-}42)$.

At the same time, Theorem 4.1(2) implies that 3-$\bar{1}$-42 belongs to the so-called *Wilf Equivalence class* for the generalized pattern 1-23. That is, if

$$\tau \in \{1\text{-}23, 3\text{-}21, 12\text{-}3, 32\text{-}1, 1\text{-}32, 3\text{-}12, 21\text{-}3, 23\text{-}1\}$$

then the size of the avoidance class $S_n(\tau)$ is also given by the n^{th} Bell number B_n. In particular, Claesson [5] showed that $|S_n(23\text{-}1)| = B_n$ via a direct bijection between permutations avoiding 23-1 and set partitions. Furthermore, given $\sigma \in S_n(3\text{-}\bar{1}\text{-}42)$, each segment between consecutive right-to-left minima must be a single decreasing run (when from read left to right), so it is easy to see that $S_n(3\text{-}\bar{1}\text{-}42) = S_n(23\text{-}1)$. Thus, the barred pattern 3-$\bar{1}$-42 and the generalized pattern 23-1 are not just in the same Wilf equivalence class but also have identical avoidance classes.

Still, even though $S_n(3\text{-}\bar{1}\text{-}42) = S_n(23\text{-}1)$, it is more natural to use avoidance of 3-$\bar{1}$-42 when studying Patience Sorting. Fundamentally, this lets us look at $S_n(3\text{-}\bar{1}\text{-}42)$ as the set of equivalence classes in $\mathfrak{S}_n$ modulo $3\text{-}\bar{1}\text{-}42 \overset{PS}{\sim} 3\text{-}\bar{1}\text{-}24$, where each equivalence class corresponds to a unique pile configuration. The same equivalence relation is not easy to describe when starting with an occurrence of 23-1. (Note that $23\text{-}1 \sim 2\text{-}13$ or $23\text{-}1 \sim 21\text{-}3$ is wrong since we would incorrectly get $2431 \sim 2314$ or $2431 \sim 2134$ instead of the correct $2431 \sim 2413$).

This suggests that there is even more information about pattern avoidance to be gotten from such a simple algorithm as Patience Sorting.

References

[1] D. Aldous and P. Diaconis. "Longest Increasing Subsequences: From Patience Sorting to the Baik-Deift-Johansson Theorem." *Bull. Amer. Math. Soc.* **36** (1999), 413–432.

[2] E. Babson and E. Steingrímsson. "Generalized Permutation Patterns and a Classification of the Mahonian Statistics." *Séminaire Lotharingien de Combinatoire* **B44b** (2000), 18 pp.

[3] J. Baik, P. Deift, and K. Johansson. "On the Distribution of the Length of the Longest Increasing Subsequence in a Random Permutation." *J. Amer. Math. Soc.* **12** (1999), 1119–1178.

[4] M. Bóna. *Combinatorics of Permutations.* Chapman & Hall/CRC Press, 2004.

[5] A. Claesson. "Generalized Pattern Avoidance." *Europ. J. Combin.* **22** (2001), 961–971.

[6] A. Claesson and T. Mansour. "Enumerating Permutations Avoiding a Pair of Babson-Steingrímsson Patterns." *Ars Combinatoria* **77** (2005), 17–31.

[7] J. B. Conrey. "The Riemann Hypothesis." *Not. Amer. Math. Soc.* **50** (2003), 341–353.

[8] S. Dulucq, S. Gire, and O. Guibert. "A Combinatorial Proof of J. West's Conjecture." *Discrete Math.* **187** (1998), 71–96.

[9] S. Dulucq, S. Gire, and J. West. "Permutations with Forbidden Subsequences and Nonseparable Maps." *Discrete Math.* **153** (1996), 85–103.

[10] W. Fulton. *Young Tableaux.* LMS Student Texts 35. Cambridge University Press, 1997.

[11] C. Greene. "An Extension of Schensted's Theorem." *Adv. in Math.* **14** (1974), 254–265.

[12] D. E. Knuth. "Permutations, Matrices, and Generalized Young Tableaux." *Pac. J. Math.* **34** (1970), 709–727.

[13] C. Kuykendall. "Analyzing Solitaire." *Science* **283** (1999), 794–795.

[14] D. Mackenzie. "From Solitaire, a Clue to the World of Prime Numbers." *Science* **282** (1998), 1631–1633.

[15] C. L. Mallows. "Problem 62-2, Patience Sorting." *SIAM Review* **4** (1962), 148–149.

[16] C. L. Mallows. "Problem 62-2'.' *SIAM Review* **5** (1963), 375–376.

[17] A. Marcus and G. Tardos. "Excluded Permutation Matrices and the Stanley-Wilf Conjecture." *J. Combin. Th. Ser. A* **107** (2004), 153–160.

[18] I. Peterson. "Solitaire-y Sequences." *MathTrek.* Available online at `http://www.maa.org/mathland/mathtrek_7_5_99.html`

[19] G. de B. Robinson. "On the Representations of the Symmetric Group." *Amer. J. Math.* **60** (1938), 745–760.

[20] B. Sagan. *The Symmetric Group, Second Edition.* Graduate Texts in Mathematics 203. Springer-Verlag, 2000.

[21] C. Schensted. "Longest Increasing and Decreasing Subsequences." *Can. J. Math.* **13** (1961), 179–191.

[22] M. P. Schützenberger. "Quelques remarques sur une construction de Schensted." *Math. Scand.* **12** (1963), 117–128.

[23] R. Stanley. *Enumerative Combinatorics, Volume 2.* Studies in Advanced Mathematics 62. Cambridge University Press, 1999.

[24] G. Viennot. "Une forme géométrique de la correspondance de Robinson-Schensted." In *Combinatoire et Représenatation du Groupe Symétrique.* D. Foata, ed. Lecture Notes in Mathematics 579. Springer-Verlag, 1977, pp. 29–58.

[25] J. West. Permutations with Forbidden Subsequences and Stack-sortable Permutations. Ph.D. thesis, M.I.T., 1990.

[26] X. Yan, P. Diaconis, P. Rusmevichientong, and B. Van Roy. "Solitaire: Man Versus Machine." In *Advances in Neural Information Processing Systems 17.* L.K. Saul, Y. Weiss, and L. Bottou, eds. MIT Press, 2005, to appear.