

Analysis of Algorithms for Permutations Biased by Their Number of Records

Nicolas Auger¹, Mathilde Bouvel^{2†}, Cyril Nicaud¹, Carine Pivoteau¹

¹ Université Paris-Est, LIGM (UMR 8049), F77454 Marne-la-Vallée, France

² CNRS and Institut für Mathematik, Universität Zürich, Zurich, Switzerland

The topic of the article is the parametric study of the complexity of algorithms on arrays of pairwise distinct integers. We introduce a model that takes into account the non-uniformness of data, which we call the Ewens-like distribution of parameter θ for records on permutations: the weight θ^r of a permutation depends on its number r of records. We show that this model is meaningful for the notion of presortedness, while still being mathematically tractable. Our results describe the expected value of several classical permutation statistics in this model, and give the expected running time of three algorithms: the Insertion Sort, and two variants of the Min-Max search.

Keywords: permutation, Ewens distribution, random generation, analysis of algorithms

1 Introduction

A classical framework for analyzing the average running time of algorithms is to consider uniformly distributed inputs. Studying the complexity of an algorithm under this uniform model usually gives a quite good understanding of the algorithm. However, it is not always easy to argue that the uniform model is relevant, when the algorithm is used on a specific data set. Observe that, in some situations, the uniform distribution arises by construction, from the randomization of a deterministic algorithm. This is the case with Quick Sort for instance, when the pivot is chosen uniformly at random. In other situations, the uniformity assumption may not fit the data very well, but still is a reasonable first step in modeling it, which makes the analysis mathematically tractable.

In practical applications where the data is a sequence of values, it is not unusual that the input is already partially sorted, depending on its origin. Consequently, assuming that the input is uniformly distributed, or shuffling the input as in the case of Quick Sort, may not be a good idea. Indeed, in the last decade, standard libraries of well-established languages have switched to sorting algorithms that take advantage of the “almost-sortedness” of the input. A noticeable example is Tim Sort Algorithm, used in Python (since 2002) and Java (since Java 7): it is particularly efficient to process data consisting of long increasing (or decreasing) subsequences.

In the case of sorting algorithms, the idea of taking advantage of some bias in the data towards sorted sequences dates back to Knuth [9, p. 336]. It has been embodied by the notion of *presortedness*, which

[†]Supported by a Marie Heim-Vögtlin grant of the Swiss National Science Foundation.

quantifies how far from sorted a sequence is. There are many ways of defining measures of presortedness, and it has been axiomatized by Mannila [10] (see Section 2.2 for a brief overview). For a given measure of presortedness m , the classical question is to find a sorting algorithm that is optimal for m , meaning that it minimizes the number of comparisons as a function of both the size of the input and the value of m . For instance, Knuth’s Natural Merge Sort [9] is optimal for the measure r = “number of runs”, with a worst case running time of $\mathcal{O}(n \log r)$ for an array of length n .

Most measures of presortedness studied in the literature are directly related to basic statistics on permutations. Consequently, it is natural to define biased distributions on permutations that depend on such statistics, and to analyze classical algorithms under these non-uniform models. One such distribution is very popular in the field of discrete probability: the Ewens distribution. It gives to each permutation σ a probability that is proportional to $\theta^{\text{cycle}(\sigma)}$, where $\theta > 0$ is a parameter and $\text{cycle}(\sigma)$ is the number of cycles in σ . Similarly, for any classical permutation statistics χ , a non-uniform distribution on permutations may be defined by giving to any σ a probability proportional to $\theta^{\chi(\sigma)}$. We call such distributions *Ewens-like distributions*. Note that the Ewens-like distribution for the number of inversions is quite popular, under the name of *Mallows distribution* [7, and references therein].

In this article, we focus on the Ewens-like distribution according to χ = number of records (a.k.a. *left to right maxima*). The motivation for this choice is twofold. First, the number of records is directly linked to the number of cycles by the fundamental bijection (see Section 2.1). So, we are able to exploit the nice properties of the classical Ewens distribution, and have a non-uniform model that remains mathematically tractable. Second, we observe that the number of non-records is a measure of presortedness. Therefore, our distribution provides a model for analyzing algorithms which is meaningful for the notion of presortedness, and consequently which may be more realistic than the uniform distribution. We first study how this distribution impacts the expected value of some classical permutation statistics, depending on the choice of θ . Letting θ depend on n , we can reach different kinds of behavior. Then, we analyze the expected complexity of Insertion Sort under this biased distribution, as well as the effect of branch prediction on two variants of the simultaneous minimum and maximum search in an array.

2 Permutations and Ewens-like distributions

2.1 Permutations as words or sets of cycles

For any integers a and b , let $[a, b] = \{a, \dots, b\}$ and for every integer $n \geq 1$, let $[n] = [1, n]$. By convention $[0] = \emptyset$. If E is a finite set, let $\mathfrak{S}(E)$ denote the set of all permutations on E , i.e., of bijective maps from E to itself. For convenience, $\mathfrak{S}([n])$ is written \mathfrak{S}_n in the sequel. Permutations of \mathfrak{S}_n can be seen in several ways (reviewed for instance in [3]). Here, we use both their representations as words and as sets of cycles.

A permutation σ of \mathfrak{S}_n can be represented as a word $w_1 w_2 \dots w_n$ containing exactly once each symbol in $[n]$: by simply setting $w_i = \sigma(i)$ for all $i \in [n]$. Conversely, any sequence (or word) of n distinct integers can be interpreted as representing a permutation of \mathfrak{S}_n . For any sequence $s = s_1 s_2 \dots s_n$ of n distinct integers, the rank $\text{rank}_s(s_i)$ of s_i is defined as the number of integers appearing in s that are smaller than or equal to s_i . Then, for any sequence s of n distinct integers, the *normalization* $\text{norm}(s)$ of s is the unique permutation σ of \mathfrak{S}_n such that $\sigma(i) = \text{rank}_s(s_i)$. For instance, $\text{norm}(8254) = 4132$.

Many permutation statistics are naturally expressed on their representation as words. One will be of particular interest for us: the number of records. If σ is a permutation of \mathfrak{S}_n and $i \in [n]$, there is a *record* at position i in σ (and subsequently, $\sigma(i)$ is a record) if $\sigma(i) > \sigma(j)$ for every $j \in [i - 1]$. In the word representation of permutations, records are therefore elements that have no larger elements to their

left. This equivalent definition of records naturally extends to sequences of distinct integers, and for any sequence s of distinct integers, the positions of the records in s and in $\text{norm}(s)$ are the same. A position that is not a record is called a *non-record*.

A cycle of size k in a permutation $\sigma \in \mathfrak{S}_n$ is a subset $\{i_1, \dots, i_k\}$ of $[n]$ such that $i_1 \xrightarrow{\sigma} i_2 \dots \xrightarrow{\sigma} i_k \xrightarrow{\sigma} i_1$. It is written (i_1, i_2, \dots, i_k) . Any permutation can be decomposed as the set of its cycles. For instance, the cycle decomposition of τ represented by the word 6321745 is $(32)(641)(75)$.

These two ways of looking at permutations (as words or as set of cycles) are rather orthogonal, but there is still a link between them, provided by the so-called *fundamental bijection* or *transition lemma*. The fundamental bijection, denoted F , is the following transformation:

1. Given σ a permutation of size n , consider the cycle decomposition of σ .
2. Write every cycle starting with its maximal element, and write the cycles in increasing order of their maximal (i.e., first) element.
3. Erasing the parenthesis gives $F(\sigma)$.

Continuing our previous example gives $F(\tau) = 3264175$. This transformation is a bijection, and transforms a permutation as set of cycles into a permutation as word. Moreover, it maps the number of cycles to the number of records. For references and details about this bijection, see for example [3, p. 109–110].

2.2 The number of non-records as a measure of presortedness

The concept of presortedness, formalized by Mannila [10], naturally arises when studying sorting algorithms which efficiently sort sequences already almost sorted. Let E be a totally ordered set. We denote by E^* the set of all nonempty sequences of distinct elements of E , and by \cdot the concatenation on E^* . A mapping m from E^* to \mathbb{N} is a *measure of presortedness* if it satisfies:

1. if $X \in E^*$ is sorted then $m(X) = 0$;
2. if $X = (x_1, \dots, x_\ell)$ and $Y = (y_1, \dots, y_\ell)$ are two elements of E^* having same length, and such that for every $i, j \in [\ell]$, $x_i < x_j \Leftrightarrow y_i < y_j$ then $m(X) = m(Y)$;
3. if X is a subsequence of Y then $m(X) \leq m(Y)$;
4. if every element of X is smaller than every element of Y then $m(X \cdot Y) \leq m(X) + m(Y)$;
5. for every symbol $a \in E$ that does not occur in X , $m(a \cdot X) \leq |X| + m(X)$.

Classical measures of presortedness [10] are the number of inversions, the number of swaps, ... One can easily see, checking conditions 1 to 5, that $m_{\text{rec}}(s) = \text{number of non-records in } s = |s| - \text{number of records in } s$ defines a measure of presortedness on sequences of distinct integers. Note that because of condition 2, studying a measure of presortedness on \mathfrak{S}_n is not a restriction with respect to studying it on sequences of distinct integers.

Given a measure of presortedness m , we are interested in optimal sorting algorithms with respect to m . Let $\text{below}_m(n, k) = \{\sigma : \sigma \in \mathfrak{S}_n, m(\sigma) \leq k\}$. A sorting algorithm is *m-optimal* (see [10] and [11] for more details) if it performs in the worst case $\mathcal{O}(n + \log |\text{below}_m(n, k)|)$ comparisons when applied to $\sigma \in \mathfrak{S}_n$ such that $m(\sigma) = k$, uniformly in k . There is a straightforward algorithm that is m_{rec} -optimal. First scan σ from left to right and put the records in one (sorted) list L_R and the non-records in another list L_N . Sort L_N using a $\mathcal{O}(|L_N| \log |L_N|)$ algorithm, then merge it with L_R . The worst case running time of this algorithm is $\mathcal{O}(n + k \log k)$ for permutations σ of \mathfrak{S}_n such that $m_{\text{rec}}(\sigma) = k$. Moreover, $|\text{below}_{m_{\text{rec}}}(n, k)| \geq k!$ for any $k \geq n$, since it contains the $k!$ permutations of the form

$(k+1)(k+2)\dots n \cdot \tau$ for $\tau \in \mathfrak{S}_k$. Consequently, $\mathcal{O}(n + k \log k) = \mathcal{O}(n + \log |\text{below}_{m_{\text{rec}}}(n, k)|)$, proving m_{rec} -optimality.

2.3 Ewens and Ewens-like distribution

The Ewens distribution on permutations (see for instance [1, Ch. 4 & 5]) is a generalization of the uniform distribution on \mathfrak{S}_n : the probability of a permutation depends on its number of cycles. Denoting $\text{cycle}(\sigma)$ the number of cycles of any permutation σ , the Ewens distribution of parameter θ (where θ is any fixed positive real number) gives to any σ the probability $\frac{\theta^{\text{cycle}(\sigma)}}{\sum_{\rho \in \mathfrak{S}_n} \theta^{\text{cycle}(\rho)}}$. As seen in [1, Ch. 5], the normalization constant $\sum_{\rho \in \mathfrak{S}_n} \theta^{\text{cycle}(\rho)}$ is $\theta^{(n)}$, where the notation $x^{(n)}$ (for any real x) denotes the *rising factorial* defined by $x^{(n)} = x(x+1) \cdots (x+n-1)$ (with the convention that $x^{(0)} = 1$).

Mimicking the Ewens distribution, it is natural (and has appeared on several occasions in the literature, see for instance [4, Example 12]) to define other non-uniform distributions on \mathfrak{S}_n , where we introduce a bias according to some statistics χ . The Ewens-like distribution of parameter θ (again θ is any fixed positive real number) for statistics χ is then the one that gives to any $\sigma \in \mathfrak{S}_n$ the probability $\frac{\theta^{\chi(\sigma)}}{\sum_{\rho \in \mathfrak{S}_n} \theta^{\chi(\rho)}}$. The classical Ewens distribution corresponds to $\chi = \text{number of cycles}$. Ewens-like distributions can be considered for many permutations statistics, like the number of inversions, of fixed points, of runs, ... In this article, we focus on the distribution associated with $\chi = \text{number of records}$. We refer to it as the Ewens-like distribution for records (with parameter θ). For any σ , we let $\text{record}(\sigma)$ denote the number of records of σ , and define the *weight* of σ as $w(\sigma) = \theta^{\text{record}(\sigma)}$. The Ewens-like distribution for records on \mathfrak{S}_n gives probability $\frac{w(\sigma)}{W_n}$ to any $\sigma \in \mathfrak{S}_n$, where $W_n = \sum_{\rho \in \mathfrak{S}_n} w(\rho)$. Note that the normalization constant is $W_n = \theta^{(n)}$, like in the classical Ewens distribution: indeed, the fundamental bijection reviewed above shows that there are as many permutations with c cycles as permutations with c records. Fig. 1 shows random permutations under the Ewens-like distribution for records, for various values of θ .




Fig. 1: Random permutations under the Ewens-like distribution on \mathfrak{S}_{100} with, from left to right, $\theta = 1$ (corresponding to the uniform distribution), 50, 100, and 500. For each diagram, the darkness of a point (i, j) is proportional to the number of generated permutations σ such that $\sigma(i) = j$, for a sampling of 10000 random permutations.

2.4 Linear random samplers

Efficient random samplers have several uses for the analysis of algorithms in general. They allow to estimate quantities of interest (even when their computation with a theoretical approach is not feasible), and can be used to double-check theoretical results. They are also a precious tool to visualize the objects under study (the diagrams in Fig. 4 were obtained in this way), allowing to define new problems on these objects (for example: can we describe the limit shape of the diagrams shown in Fig. 4?).

As mentioned in [6, §2.1], one can easily obtain a linear time and space algorithm to generate a random permutation according to the Ewens distribution (for cycles), using a variant of the Chinese restaurant process reviewed in what follows. To generate a permutation of size n , we start with an empty array⁽ⁱ⁾ σ of length n that is used to store the values of the $\sigma(i)$'s. For i from 1 to n , we choose to either create a new cycle containing only i with probability $\frac{\theta}{\theta+i-1}$ or to insert i in one of the existing cycles with probability $\frac{i-1}{\theta+i-1}$. To create a new cycle, we set $\sigma[i] = i$. To insert i in an existing cycle, we choose uniformly at random an element j in $[i-1]$ to be the element following i in its cycle, and we set $\sigma[i] = j$ and $\sigma[\sigma^{-1}[j]] = i$. To avoid searching for $\sigma^{-1}[j]$ in the array σ , we only need to keep σ^{-1} in a second array while adding the elements in σ .

Starting from this algorithm, we can easily design a linear random sampler for permutations according to the Ewens-like distribution for records, using the fundamental bijection. The first step is to generate a permutation σ in \mathfrak{S}_n with the above algorithm. Then, we write the cycles of σ in reverse order of their maximum, as sequences, starting from the last element and up to exhaustion of the cycle: $n, \sigma[n], \sigma[\sigma[n]], \dots, \sigma^{-1}[n]$. Each time we write an element i , we set $\sigma[i] = 0$ and each time a cycle is finished, we search the next value of i such that $\sigma[i] \neq 0$ to start the next cycle. This new cycle will be written before the one that has just been written. Note that all these operations can be performed in time complexity $O(1)$ using doubly linked lists for the resulting permutation. In the end, the cycles will be written as sequences starting by their maximum, sorted in increasing order of their maximum, which is the fundamental bijection.

Note that there exists another branching process, known as the *Feller coupling*, to generate permutations according to the Ewens distribution (see for instance [1, p.16]). Although it is less natural than with the Chinese restaurant process, it is also possible to infer linear random samplers from it. Details will be provided in an extended version of this work.

3 Average value of statistics in biased random permutations

Let θ be any fixed positive real number. In this section, we study the behavior of several statistics on permutations, when they follow the Ewens-like distribution for records with parameter θ . Our purpose is mostly to illustrate methods to obtain precise descriptions of the behavior of such statistics. Such results allow a fine analysis of algorithms whose complexity depends on the studied statistics.

Recall that, for any $\sigma \in \mathfrak{S}_n$, $w(\sigma) = \theta^{\text{record}(\sigma)}$ and the probability of σ is $\frac{w(\sigma)}{W_n}$, with $W_n = \theta^{(n)}$. Recall also that the records of any sequence of distinct integers are well-defined. For any such sequence s we subsequently set $\text{record}(s)$ to be the number of records of s and $w(s) = \theta^{\text{record}(s)}$. Note that for any such sequence s , $w(s) = w(\text{norm}(s))$, because the positions (and hence the number) of records do not change when normalizing.

3.1 Technical lemmas

Seeing a permutation of \mathfrak{S}_n as a word, it can be split (in many ways) into two words as $\sigma = \pi \cdot \tau$ for the usual concatenation on words. Note that here π and τ are not normalized permutations: τ belongs to the set $\mathfrak{S}_{k \text{ in } n}$ of all sequences of k distinct integers in $[n]$ where $k = |\tau|$, and π belongs to $\mathfrak{S}_{n-k \text{ in } n}$. The weight function w behaves well with respect to this decomposition, as shown in the following lemmas.

⁽ⁱ⁾ Note that our array starts at index 1.

Lemma 1 *Let n be an integer, and τ be a sequence of $k \leq n$ distinct integers in $[n]$. Denote by m the number of records in τ whose value is larger than the largest element of $[n]$ which does not appear in τ , and define $w'_n(\tau)$ as θ^m . For all $\sigma \in \mathfrak{S}_n$, if $\sigma = \pi \cdot \tau$, then $w(\sigma) = w(\pi) \cdot w'_n(\tau)$.*

For instance, the definition of $w'_n(\tau)$ gives $w'_9(6489) = \theta^2$ (8 and 9 are records of τ larger than 7) and $w'_{10}(6489) = 1$ (there are no records in τ larger than 10).

We extend the weight function w to subsets X of \mathfrak{S}_n as $w(X) = \sum_{\sigma \in X} w(\sigma)$. For any sequence τ of $k \leq n$ distinct integers in $[n]$, the right-quotient of X with τ is $X/\tau = \{\pi : \pi \cdot \tau \in X\}$. Since $w(\pi) = w(\text{norm}(\pi))$ for all sequences π of distinct integers, we have $w(X/\tau) = w(\text{norm}(X/\tau))$ for all X and τ as above (As expected, $\text{norm}(Y)$ means $\{\text{norm}(\pi) : \pi \in Y\}$).

For $k \in [n]$, we say that $X \subseteq \mathfrak{S}_n$ is *quotient-stable for k* if $w(X/\tau)$ is constant when τ runs over $\mathfrak{S}_{k \text{ in } n}$. When X is quotient-stable for k , we denote $w_k^q(X)$ the common value of $w(X/\tau)$ for τ as above. For instance, $X = (4321, 3421, 4132, 3142, 4123, 2143, 3124, 1324)$ is quotient-stable for $k = 1$. Indeed,

$$\begin{aligned} w_1^q(X) &= w(X/1) = w(\{432, 342\}) = w(X/2) = w(\{413, 314\}) = \\ &= w(X/3) = w(\{412, 214\}) = w(X/4) = w(\{312, 132\}) = \theta + \theta^2. \end{aligned}$$

Note that \mathfrak{S}_n is quotient-stable for all $k \in [n]$: indeed, for any τ of size k , $\text{norm}(\mathfrak{S}_n/\tau) = \mathfrak{S}_{n-k}$ so that $w(\mathfrak{S}_n/\tau) = w(\mathfrak{S}_{n-k})$ for all τ of size k . It follows that $w_k^q(\mathfrak{S}_n) = w(\mathfrak{S}_{n-k}) = \theta^{(n-k)}$.

Lemma 2 *Let $X \subseteq \mathfrak{S}_n$ be quotient-stable for $k \in [n]$. Then $w(X) = \frac{\theta^{(n)}}{\theta^{(n-k)}} w_k^q(X)$.*

A typical example of use of Lemma 2 is given in the proof of Theorem 3.

Remark: Lemma 2 is a combinatorial version of a simple probabilistic property: Let E_τ be the set of elements of \mathfrak{S}_n that end with τ . If A is an event on \mathfrak{S}_n and if the probability of A given E_τ is the same for every $\tau \in \mathfrak{S}_{k \text{ in } n}$, then it is equal to the probability of A , by the law of total probabilities. \diamond

3.2 Summary of asymptotic results

The rest of this section is devoted to studying the expected behavior of some permutation statistics, under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ . We are especially interested in the asymptotics in n when θ is constant or is a function of n . The studied statistics are: number of records, number of descents, first value, and number of inversions. A summary of our results is presented in Table 3.2. The asymptotics reported in Table 3.2 follow from Corollaries 4, 6, 9, 11 either immediately or using the so-called *digamma* function. The digamma⁽ⁱⁱ⁾ function is defined by $\Psi(x) = \Gamma'(x)/\Gamma(x)$. It satisfies the identity $\sum_{i=0}^{n-1} \frac{1}{x+i} = \Psi(x+n) - \Psi(x)$, and its asymptotic behavior as $x \rightarrow \infty$ is $\Psi(x) = \log(x) - \frac{1}{2x} - \frac{1}{12x^2} + o\left(\frac{1}{x^2}\right)$. We also define $\Delta(x, y) = \Psi(x+y) - \Psi(x)$, so that $\Delta(x, n) = \sum_{i=0}^{n-1} \frac{1}{x+i}$ for any positive integer n . In Table 3.2 and in the sequel, we use the notations $\mathbb{P}_n(E)$ (resp. $\mathbb{E}_n[\chi]$) to denote the probability of an event E (resp. the expected value of a statistics χ) under the Ewens-like distribution on \mathfrak{S}_n for records.

Remark: To some extent, our results may also be interpreted on the classical Ewens distribution, via the fundamental bijection. Indeed the number of records (resp. the number of descents, resp. the first value) of σ corresponds to the number of cycles (resp. the number of anti-excedances⁽ⁱⁱⁱ⁾, resp. the minimum over

⁽ⁱⁱ⁾ For details, see https://en.wikipedia.org/wiki/Digamma_function (accessed on April 27, 2016).

⁽ⁱⁱⁱ⁾ An anti-excedance of $\sigma \in \mathfrak{S}_n$ is $i \in [n]$ such that $\sigma(i) < i$. The proof that descents of σ are equinumerous with anti-excedances of $F^{-1}(\sigma)$ is a simple adaptation of the proof of Theorem 1.36 in [3], p. 110–111.

	$\theta = 1$ (uniform)	fixed $\theta > 0$	$\theta := n^\epsilon,$ $0 < \epsilon < 1$	$\theta := \lambda n,$ $\lambda > 0$	$\theta := n^\delta,$ $\delta > 1$	See Cor.
$\mathbb{E}_n[\text{record}]$	$\log n$	$\theta \cdot \log n$	$(1 - \epsilon) \cdot n^\epsilon \log n$	$\lambda \log(1 + 1/\lambda) \cdot n$	n	4
$\mathbb{E}_n[\text{desc}]$	$n/2$	$n/2$	$n/2$	$n/2(\lambda + 1)$	$n^{2-\delta}/2$	6
$\mathbb{E}_n[\sigma(1)]$	$n/2$	$n/(\theta + 1)$	$n^{1-\epsilon}$	$(\lambda + 1)/\lambda$	1	9
$\mathbb{E}_n[\text{inv}]$	$n^2/4$	$n^2/4$	$n^2/4$	$n^2/4 \cdot f(\lambda)$	$n^{3-\delta}/6$	11

Tab. 1: Asymptotic behavior of some permutation statistics under the Ewens-like distribution on \mathfrak{S}_n for records. We use the shorthand $f(\lambda) = 1 - 2\lambda + 2\lambda^2 \log(1 + 1/\lambda)$. All the results in this table are asymptotic equivalents.

all cycles of the maximum value in a cycle) of $F^{-1}(\sigma)$. Consequently, Corollary 4 is just a consequence of the well-known expectation of the number of cycles under the Ewens distribution (see for instance [1, §5.2]). Similarly, the expected number of anti-excedances (Corollary 6) can be derived easily from the results of [6]. Those results on the Ewens distribution do not however give access to results as precise as those stated in Theorems 3 and 5, which are needed to prove our results of Section 4. Finally, to the best of our knowledge, the behavior of the third statistics (minimum over all cycles of the maximum value in a cycle) has not been previously studied, and we are not aware of any natural interpretation of the number of inversions of σ in $F^{-1}(\sigma)$. \diamond

3.3 Expected values of some permutation statistics

We start our study by computing how the value of parameter θ influences the expected number of records.

Theorem 3 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , for any $i \in [n]$, the probability that there is a record at position i is: $\mathbb{P}_n(\text{record at } i) = \frac{\theta}{\theta+i-1}$.*

Proof: We prove this theorem by splitting permutations seen as words after their i -th element, as shown in Fig. 2. Let $\mathcal{R}_{n,i}$ denote the set of permutations of \mathfrak{S}_n having a record at position i . We claim that the set $\mathcal{R}_{n,i}$ is quotient-stable for $n-i$, and that $w(\mathcal{R}_{n,i}) = \frac{\theta^{(n)}}{\theta^{(i)}} \cdot \theta^{(i-1)} \cdot \theta$. It will immediately follow that $\mathbb{P}_n(\text{record at } i) = \frac{w(\mathcal{R}_{n,i})}{\theta^{(n)}} = \frac{\theta^{(i-1)} \cdot \theta}{\theta^{(i)}} = \frac{\theta}{\theta+i-1}$. We now prove the claim. Let τ be any sequence in \mathfrak{S}_{n-i} . Observe that $\text{norm}(\mathcal{R}_{n,i}/\tau) = \mathcal{R}_{i,i}$. Since the number of records is stable by normalization, it follows that $w(\mathcal{R}_{n,i}/\tau) = w(\mathcal{R}_{i,i})$. By definition, $\pi \in \mathfrak{S}_i$ is in $\mathcal{R}_{i,i}$ if and only if $\pi(i) = i$. Thus $\mathcal{R}_{i,i} = \mathfrak{S}_{i-1} \cdot i$ in the word representation of permutations. Hence, $w(\mathcal{R}_{i,i}) = \theta^{(i-1)}\theta$, since the last element is a record by definition. This yields $w(\mathcal{R}_{n,i}/\tau) = \theta^{(i-1)}\theta$ for any $\tau \in \mathfrak{S}_{n-i}$, proving that $\mathcal{R}_{n,i}$ is quotient-stable for $n-i$, and that $w_{n-i}^q(\mathcal{R}_{n,i}) = \theta^{(i-1)}\theta$. By Lemma 2, it follows that $w(\mathcal{R}_{n,i}) = \frac{\theta^{(n)}}{\theta^{(n-(n-i))}} \cdot w_{n-i}^q(\mathcal{R}_{n,i}) = \frac{\theta^{(n)}}{\theta^{(i)}} \cdot \theta^{(i-1)} \cdot \theta$. \square




Fig. 2: The decomposition used to compute the probability of having a record at i . This record has weight θ and thus, for any fixed τ , the weights of all possible π sum to $w(\mathfrak{S}_{i-1}) \cdot \theta = \theta^{(i-1)} \cdot \theta$.

Corollary 4 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , the expected value of the number of records is: $\mathbb{E}_n[\text{record}] = \sum_{i=1}^n \frac{\theta}{\theta+i-1} = \theta \cdot \Delta(\theta, n)$.*

Next, we study the expected number of descents. Recall that a permutation σ of \mathfrak{S}_n has a *descent* at position $i \in \{2, \dots, n\}$ if $\sigma(i-1) > \sigma(i)$. We denote by $\text{desc}(\sigma)$ the number of descents in σ . We are interested in descents as they are directly related to the number of increasing runs in a permutation (each such run but the last one is immediately followed by a descent, and conversely). Some sorting algorithms, like Knuth's Natural Merge Sort, use the decomposition into runs.

The following theorem is proved using Lemmas 1 and 2 and the decomposition of Fig. 3.

Theorem 5 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , for any $i \in \{2, \dots, n\}$, the probability that there is a descent at position i is: $\mathbb{P}_n(\sigma(i-1) > \sigma(i)) = \frac{(i-1)(2\theta+i-2)}{2(\theta+i-1)(\theta+i-2)}$.*




Fig. 3: The two cases for the probability of having a descent at i . We decompose σ as $\pi \cdot \sigma(i-1) \cdot \sigma(i) \cdot \tau$, and we let $\rho = \text{norm}(\pi \cdot \sigma(i-1) \cdot \sigma(i))$. On the left, the case where $\sigma(i-1)$ is a record, that is, $\rho(i-1) = i$: there are $i-1$ possibilities for $\rho(i)$. On the right, the case where $\sigma(i-1)$ is not a record: there are $\binom{i-1}{2}$ possibilities for the values of $\rho(i)$ and $\rho(i-1)$. In both cases, once the images of $j \in \{i-1, \dots, n\}$ by σ have been chosen, the weight of all possible beginnings sum to $w(\mathfrak{S}_{i-2}) = \theta^{(i-2)}$.

Corollary 6 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , the expected value of the number of descents is: $\mathbb{E}_n[\text{desc}] = \frac{n(n-1)}{2(\theta+n-1)}$.*

In the second row of Table 3.2, remark that the only way of obtaining a sublinear number of descents is to take very large values for θ .

Finally, we study the expected value of $\sigma(1)$. We are interested in this statistic to show a proof that differs from the ones for the numbers of records and descents: the expected value of the first element of a permutation is not obtained using Lemma 2.

Lemma 7 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , for any $k \in [0, n-1]$, the probability that a permutation starts with a value larger than k is: $\mathbb{P}_n(\sigma(1) > k) = \frac{(n-1)! \theta^{(n-k)}}{(n-k-1)! \theta^{(n)}}$.*

Proof: Let $\mathcal{F}_{n,k}$ denote the set of permutations of \mathfrak{S}_n such that $\sigma(1) > k$. Such a permutation can uniquely be obtained by choosing the preimages of the elements in $[k]$ in $\{2, \dots, n\}$, then by mapping bijectively the remaining elements to $[k+1, n]$. Since none of the elements in $[k]$ is a record and since the elements of $[k+1, n]$ can be ordered in all possible ways, we get that $w(\mathcal{F}_{n,k}) = \binom{n-1}{k} k! \theta^{(n-k)}$. Indeed, there are $\binom{n-1}{k} k!$ ways to position and order the elements of $[k]$, and the total weight of the elements larger than k is $\theta^{(n-k)}$. Hence, $\mathbb{P}_n(\sigma(1) > k) = \frac{w(\mathcal{F}_{n,k})}{w(\mathfrak{S}_n)} = \frac{\binom{n-1}{k} k! \theta^{(n-k)}}{\theta^{(n)}} = \frac{(n-1)! \theta^{(n-k)}}{(n-k-1)! \theta^{(n)}}$. \square

Theorem 8 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , for any $k \in [n]$, the probability that a permutation starts with k is: $\mathbb{P}_n(\sigma(1) = k) = \frac{(n-1)! \theta^{(n-k)} \theta}{(n-k)! \theta^{(n)}}$.*

Corollary 9 Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , the expected value of the first element of a permutation is: $\mathbb{E}_n[\sigma(1)] = \frac{\theta+n}{\theta+1}$.

Remark: Our proof of Corollary 9 relies on calculus, but gives a very simple expression for $\mathbb{E}_n[\sigma(1)]$. We could therefore hope for a more combinatorial proof of Corollary 9, but we were not able to find it. \diamond

3.4 Number of inversions and expected running time of INSERTIONSORT

Recall that an *inversion* in a permutation $\sigma \in \mathfrak{S}_n$ is a pair $(i, j) \in [n] \times [n]$ such that $i < j$ and $\sigma(i) > \sigma(j)$. In the word representation of permutations, this corresponds to a pair of elements in which the largest is to the left of the smallest. This equivalent definition of inversions naturally generalizes to sequences of distinct integers. For any $\sigma \in \mathfrak{S}_n$, we denote by $\text{inv}(\sigma)$ the number of inversions of σ , and by $\text{inv}_j(\sigma)$ the number inversions of the form (i, j) in σ , for any $j \in [n]$. More formally, $\text{inv}_j(\sigma) = |\{i \in [j-1] : (i, j) \text{ is an inversion of } \sigma\}|$.

Theorem 10 Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , for any $j \in [n]$ and $k \in [0, j-1]$, the probability that there are k inversions of the form (i, j) is: $\mathbb{P}_n(\text{inv}_j(\sigma) = k) = \frac{1}{\theta+j-1}$ if $k \neq 0$ and $\mathbb{P}_n(\text{inv}_j(\sigma) = k) = \frac{\theta}{\theta+j-1}$ if $k = 0$.

Corollary 11 Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , the expected value of the number of inversions is: $\mathbb{E}_n[\text{inv}] = \frac{n(n+1-2\theta)}{4} + \frac{\theta(\theta-1)}{2} \Delta(\theta, n)$.

Recall that the INSERTIONSORT algorithm works as follows: at each step $i \in \{2, \dots, n\}$, the first $i-1$ elements are already sorted, and the i -th element is then inserted at its correct place, by swapping the needed elements.

It is well known that the number of swaps performed by INSERTIONSORT when applied to σ is equal to the number of inversions $\text{inv}(\sigma)$ of σ . Moreover, the number of comparisons $C(\sigma)$ performed by the algorithm satisfies $\text{inv}(\sigma) \leq C(\sigma) \leq \text{inv}(\sigma) + n - 1$ (see [5] for more information on INSERTIONSORT).

As a direct consequence of Corollary 11 and the asymptotic estimates of the fourth row of Table 3.2, we get the expected running time of INSERTIONSORT:

Corollary 12 Under the Ewens-like distribution for records with parameter $\theta = \mathcal{O}(n)$, the expected running time of INSERTIONSORT is $\Theta(n^2)$, like under the uniform distribution. If $\theta = n^\delta$ with $1 < \delta < 2$, it is $\Theta(n^{3-\delta})$. If $\theta = \Omega(n^2)$, it is $\Theta(n)$.

4 Expected Number of Mispredictions for the Min/Max Search

4.1 Presentation

In this section, we turn our attention to a simple and classical problem: computing both the minimum and the maximum of an array of size n . The straightforward approach (called *naive* in the sequel) is to compare all the elements of the array to the current minimum and to the current maximum, updating them when it is relevant. This is done^(iv) in Algorithm 1 and uses exactly $2n - 2$ comparisons. A classical optimization is to look at the elements in pairs, and to compare the smallest to the current minimum and the largest to the current maximum (see Algorithm 2). This uses only $3n/2$ comparisons, which is optimal. However, as reported in [2], with an implementation in C of these two algorithms, the naive algorithm proves to be

^(iv) Note that, for consistency, our arrays start at index 1, as stated at the beginning of this paper.

the fastest on uniform permutations as input. The explanation for this is a trade-off between the number of comparisons involved and an other inherent but less obvious factor that influences the running time of these algorithms: the behavior of the branch predictor.

Algorithm 1: NAIVEMINMAX(T, n)

```

1  $min \leftarrow T[1]$ 
2  $max \leftarrow T[1]$ 
3 for  $i \leftarrow 2$  to  $n$  do
4   if  $T[i] < min$  then
5      $min \leftarrow T[i]$ 
6   if  $T[i] > max$  then
7      $max \leftarrow T[i]$ 
8 return  $min, max$ 

```

Algorithm 2: 3/2-MINMAX(T, n)

```

1  $min, max \leftarrow T[n], T[n]$ 
2 for  $i \leftarrow 2$  to  $n$  by 2 do
3   if  $T[i-1] < T[i]$  then
4      $pMin, pMax \leftarrow T[i-1], T[i]$ 
5   else  $pMin, pMax \leftarrow T[i], T[i-1]$  if
6      $pMin < min$  then  $min \leftarrow pMin$  if
7      $pMax > max$  then  $max \leftarrow pMax$ 
8 return  $min, max$ 

```

In a nutshell, when running on a modern processor, the instructions that constitute a program are not executed strictly sequentially but instead, they usually overlap one another since most of the instructions can start before the previous one is finished. This mechanism is commonly described as a *pipeline* (see [8] for a comprehensive introduction on this subject). However, not all instructions are well-suited for a pipelined architecture: this is specifically the case for branching instructions such as an *if* statement. When arriving at a branch, the execution of the next instruction should be delayed until the outcome of the test is known, which stalls the pipeline. To avoid this, the processor tries to predict the result of the test, in order to decide which instruction will enter the pipeline next. If the prediction is right, the execution goes on normally, but in case of a *misprediction*, the pipeline needs to be flushed, which can significantly slow down the execution of a program.

There is a large variety of branch predictors, but nowadays, most processors use *dynamic* branch prediction: they remember partial information on the results of the previous tests at a given *if* statement, and their prediction for the current test is based on those previous results. These predictors can be quite intricate, but in the sequel, we will only consider local *1-bit predictors* which are state buffers associated to each *if* statement: they store the last outcome of the test and guess that the next outcome will be the same.

Let us come back to the problem of simultaneously finding the minimum and the maximum in an array. We can easily see that, for Algorithm 1, the behavior of a 1-bit predictor when updating the maximum (resp. minimum) is directly linked to the succession of records (resp. min-records^(v)) in the array. As we explain later on, for Algorithm 2, this behavior depends on the “pattern” seen in four consecutive elements of the array, this “pattern” indicating not only which elements are records (resp. min-records), but also where we find descents between those elements. As shown in [2], for uniform permutations, Algorithm 1 outperforms Algorithm 2, because the latter makes more mispredictions than the former, compensating for the fewer comparisons made by Algorithm 2. This corresponds to our Ewens-like distribution for $\theta = 1$. But when θ varies, the way records are distributed also changes, influencing the performances of both Algorithms 1 and 2. Specifically, when $\theta = \lambda n$, we have a linear number of records (as opposed to a logarithmic number when $\theta = 1$). The next subsections provide a detailed analysis of the number of mispredictions in Algorithms 1 and 2, under the Ewens-like distribution for records, with a particular emphasis on $\theta = \lambda n$ (which exhibits a very different behavior w.r.t. the uniform distribution – see Fig. 4).

^(v) A *min-record* (a.k.a. left to right minimum) is an element of the array such that no smaller element appears to its left.

4.2 Expected Number of Mispredictions in NaiveMinMax

Theorem 13 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , the expected numbers of mispredictions at lines 4 and 6 of Algorithm 1 satisfy respectively $\mathbb{E}_n[\mu_4] \leq \frac{2}{\theta} \mathbb{E}_n[\text{record}]$ and $\mathbb{E}_n[\mu_6] = 2\theta \Delta(\theta, n-1) - \frac{(2\theta+1)(n-1)}{\theta+n-1}$.*

Consequently, with our previous results on $\mathbb{E}_n[\text{record}]$, the expected number of mispredictions at line 4 is $\mathcal{O}(\log n)$ when $\theta = \Omega(1)$ (i.e., when $\theta = \theta(n)$ is constant or larger). Moreover, using the asymptotic estimates of the digamma function, the asymptotics of the expected number of mispredictions at line 6 is such that (again, for $\lambda > 0$, $0 < \epsilon < 1$ and $\delta > 1$):

$$\frac{\mathbb{E}_n[\mu_6]}{\mathbb{E}_n[\text{record}]} \left| \begin{array}{c|c|c|c} \text{fixed } \theta > 0 & \theta := n^\epsilon & \theta := \lambda n & \theta := n^\delta \\ \hline \sim 2\theta \cdot \log n & \sim 2(1-\epsilon) \cdot n^\epsilon \log n & \sim 2\lambda(\log(1+1/\lambda) - 1/(\lambda+1)) \cdot n & o(n) \end{array} \right.$$

In particular, asymptotically, the expected total number of mispredictions of Algorithm 1 is given by $\mathbb{E}_n[\mu_6]$ (up to a constant factor when θ is constant).

4.3 Expected Number of Mispredictions in $\frac{3}{2}$ MinMax

Mispredictions in Algorithm 2 can arise in any of the three *if* statements. We first compute the expected number of mispredictions at each of them independently. We start with the *if* statement of line 3, which compares $T[i-1]$ and $T[i]$. For our 1-bit model, there is a misprediction whenever there is a descent at $i-2$ and an ascent at i , or an ascent at i and a descent at $i-2$. A tedious study of all possible cases gives:

Theorem 14 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , the expected number of mispredictions at line 3 of Algorithm 2 satisfies*

$$\begin{aligned} \mathbb{E}_n[\nu_3] = & \frac{n-2}{4} + \frac{\theta(\theta-1)^2}{4} + \frac{\theta^2(\theta-1)^2}{12} \left(\frac{1}{\theta+n-1} - \frac{3}{\theta+n-2} - \frac{1}{\theta+1} \right) \\ & + \frac{\theta^2(\theta-1)^2}{6} \left(\Delta\left(\frac{\theta+1}{2}, \frac{n-2}{2}\right) - \Delta\left(\frac{\theta}{2}, \frac{n-2}{2}\right) \right). \end{aligned}$$

As a consequence, if $\theta = \lambda n$, then $\mathbb{E}_n[\nu_3] \sim \frac{6\lambda^2+8\lambda+3}{12(\lambda+1)^3} n$.

Theorem 15 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , the expected number of mispredictions at line 5 of Algorithm 2 satisfies $\mathbb{E}_n[\nu_7] \leq \frac{2}{\theta} \mathbb{E}_n[\text{record}]$. As a consequence, if $\theta = \lambda n$, then $\mathbb{E}_n[\nu_7] = \mathcal{O}(1)$.*

We now consider the third *if* statement of Algorithm 2. If there is a record (resp. no record) at position $i-3$ or $i-2$, then there is a misprediction when there is no record (resp. a record) at position $i-1$ or i . Studying all the possible configurations at these four positions gives the following result.

Theorem 16 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter θ , the expected number of mispredictions at line 5 of Algorithm 2 satisfies*

$$\begin{aligned} \mathbb{E}_n[\nu_8] = & \frac{(n-2)((2\theta^3 + \theta^2 - 9\theta - 3)n + 2\theta^4 - 5\theta^2 + 9\theta + 3)}{3(\theta+n-1)(\theta+n-2)} \\ & + \frac{\theta(2\theta^3 + \theta + 3)}{3} \Delta\left(\frac{\theta+1}{2}, \frac{n-2}{2}\right) - \frac{\theta(2\theta^3 + \theta - 3)}{3} \Delta\left(\frac{\theta}{2}, \frac{n-2}{2}\right). \end{aligned}$$

As a consequence, if $\theta = \lambda n$, then $\mathbb{E}_n[\nu_8] \sim (2\lambda \log(1 + \frac{1}{\lambda}) - \frac{\lambda(6\lambda^2+15\lambda+10)}{3(\lambda+1)^3})n$.

It follows from Theorems 14, 15 and 16 that:

Corollary 17 *Under the Ewens-like distribution on \mathfrak{S}_n for records with parameter $\theta = \lambda n$, the total number of mispredictions of Algorithm 2 is*

$$\mathbb{E}_n[\nu] \sim \left(2\lambda \log \left(1 + \frac{1}{\lambda} \right) - \frac{24\lambda^3 + 54\lambda^2 + 32\lambda - 3}{12(\lambda + 1)^3} \right) n.$$

Fig. 4 shows that, unlike in the uniform case ($\theta = 1$), Algorithm 2 is more efficient than Algorithm 1 under the Ewens-like distribution for records with $\theta := \lambda n$, as soon as λ is large enough.




Fig. 4: The expected number of mispredictions produced by the naive algorithm (μ) and for $\frac{3}{2}$ -minmax (ν), when $\theta := \lambda n$. We have $\mathbb{E}_n[\mu] \sim \mathbb{E}_n[\nu]$ for $\lambda_0 = \frac{\sqrt{34}-4}{6} \approx 0.305$, and there are fewer mispredictions on average with $\frac{3}{2}$ -minmax as soon as $\lambda > \lambda_0$. However, since $\frac{3}{2}$ -minmax performs $\frac{n}{2}$ fewer comparisons than the naive algorithm, it becomes more efficient before λ_0 . For instance, if a misprediction is worth 4 comparisons, $\frac{3}{2}$ -minmax is the most efficient as soon as $\lambda > 0.110$.

Acknowledgments Many thanks to Valentin Féray for providing insight and references on the classical Ewens distribution and to the referees whose comments helped us clarify the presentation of our work.

References

- [1] R. Arratia, A. D. Barbour, and S. Tavaré. *Logarithmic combinatorial structures: a probabilistic approach*. EMS Monographs in Mathematics. EMS, Zürich, 2003.
- [2] N. Auger, C. Nicaud, and C. Pivoteau. Good predictions are worth a few comparisons. In *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016, Orléans, France*, pages 12:1–12:14, 2016.
- [3] M. Bóna. *Combinatorics of permutations*. Chapman-Hall and CRC Press, 2d edition edition, 2012.
- [4] A. Borodin, P. Diaconis, and J. Fulman. On adding a list of numbers (and other one-dependent determinantal processes). *Bulletin of the American Mathematical Society (N.S.)*, 47(4):639–670, 2010.
- [5] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and Clifford S. *Introduction to Algorithms*. MIT Press, Cambridge, MA, third edition, 2009.
- [6] V. Féray. Asymptotics of some statistics in Ewens random permutations. *Electronic Journal of Probability*, 18(76):1–32, 2013.
- [7] A. Gladkikh and R. Peled. On the cycle structure of Mallows permutations. Preprint available at <http://arxiv.org/abs/1601.06991>.
- [8] J. L. Hennessy and D. A. Patterson. *Computer Architecture, Fifth Edition: A Quantitative Approach*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 5th edition, 2011.
- [9] D. E. Knuth. *The Art of Computer Programming, Volume 3: (2nd Ed.) Sorting and Searching*. Addison Wesley Longman Publish. Co., Redwood City, CA, USA, 1998.
- [10] Heikki Mannila. Measures of presortedness and optimal sorting algorithms. *IEEE Trans. Computers*, 34(4):318–325, 1985.
- [11] O. Petersson and A. Moffat. A framework for adaptive sorting. *Discrete Applied Mathematics*, 59(2):153–179, may 1995.