Large Deviations for Increasing Subsequences of Permutations and a Concurrency Application

Yuliy Baryshnikov
Depts. of Mathematics and ECE
University of Illinois at Urbana-Champaign
Urbana, IL 61801
ymb@illinois.edu

Abram Magner
Coordinated Science Lab
University of Illinois at Urbana-Champaign
Urbana, IL 61801
anmagner@illinois.edu

ABSTRACT

The study of concurrent processes with conflict points is connected with the geometry of increasing subsequences of permutations – a permutation encodes the transactions of two processes that conflict (i.e., must be executed serially), and a given increasing subsequence encodes one particular serialization of the executions of two processes. This motivates the study of random increasing subsequences of random permutations. Here, we give a large deviation principle which implies that such a subsequence never deviates too far from the identity permutation: a random serialization of two concurrent processes will not delay either process too much at any given time. We then give an efficient exact algorithm for uniform random sampling of an increasing subsequence from a given permutation. We leave to future work generalizations to larger numbers of processes, wherein conflict sets may take on more interesting geometries.

Categories and Subject Descriptors

G.2.1 [**Combinatorics**]: Combinatorial algorithms, Counting problems, Permutations and combinations. D.4.1 [**Process management**]: Concurrency.

General Terms

Algorithms, Performance, Theory.

Keywords

Large deviations, permutations, increasing subsequences.

1. INTRODUCTION

It is well known in the theory of concurrency control that the conflict points of a pair of processes can be represented geometrically as points in $[0,1]^2$ [4]. Given such a representation, a *serialization* of the two processes – a sequential scheduling of the constituent transactions of each process – corresponds exactly to a homotopy class of increasing paths from $(0,0)$ to $(1,1)$ which avoids all conflict points.

This geometric picture, in turn, corresponds exactly to the following more combinatorial setting: the n conflict points correspond to a partial order σ on a set of size n, and a serialization corresponds to an increasing subsequence of σ. That is, we have the following:

Proposition 1. *The number of distinct serializations of a pair of processes is equal to the number of increasing subsequences in the coordinatewise-domination partial order on the set of conflict points.*

It is natural to study the *typical* behavior of serializations of a *typical* pair of processes, where the conflict points are now distributed uniformly at random in the $[0,1]^2$. Now σ is a permutation (i.e., a total order) with probability 1, and the goal is to study random increasing subsequences of a random permutation.

Here we establish upper and lower bounds on large deviation probabilities for random increasing subsequences of a random permutation. For a pair of concurrent processes, this has the desirable consequence that a randomly selected serialization (usually) does not favor one process or another too much over any given time interval, so that neither process is subject to starvation.

Finally, we give an efficient exact algorithm for uniformly sampling an increasing subsequence of a given permutation.

2. PRIOR RESULTS

Increasing subsequences of random permutations (and, more generally, in random partial orders [1]) have been extensively studied. E.g., the moments of the number of increasing subsequences of a random permutation were precisely studied in [3] (notably implying poor concentration). Additionally, the length of the longest increasing subsequence of a random permutation has been characterized: its expected value is $2\sqrt{n}(1+o(1))$ and is tightly concentrated around this value [2].

3. MODEL AND MAIN RESULTS

Consider a permutation $\sigma \in S_n$, the symmetric group on n letters. An *increasing subsequence* of σ is a finite sequence $1 \leq t_1 < t_2 < \cdots < t_k \leq n$ (possibly of length $k=0$ or 1) of positive integers such that for all $j \in [k] = \{1,...,k\}$, we have $\sigma(t_j) < \sigma(t_{j+1})$. We denote by $\tau(\sigma)$ the number of increasing subsequences of σ.

Here, we are interested in σ sampled uniformly at random from S_n, along with a uniformly random increasing subsequence $T = (t_1,...,t_k)$ of σ.

Our first main result precisely characterizes the probability that any element of T is far from the diagonal:

Theorem 1 (LDP for increasing subsequences). *Let σ be a random permutation from S_n, and let T be a random increasing subsequence of σ. Then we have, for any*

$\epsilon > 0$,

$$\Pr\left[\max_{i \in T} |\sigma(i) - i| > \epsilon n\right] = \exp(-\Theta(\sqrt{n})), \qquad (1)$$

where the constants hidden in the $\Theta(\cdot)$ *may be explicitly bounded in terms of* ϵ.

This theorem implies that a random increasing subsequence is very likely to be close to the identity permutation restricted to the elements of the subsequence.

It is also of interest to be able to sample a uniformly random increasing subsequence of a given permutation. Our next result says that this can be done efficiently.

THEOREM 2 (EXACT SAMPLING ALGORITHM). *There is an algorithm for sampling a uniformly random increasing subsequence of a permutation from* S_n *in time* $\Theta(n^2)$, *assuming that arithmetic operations take constant time.*

4. PROOF OF THEOREM 1

The proof is guided by the following insight: if any given element of a random increasing subsequence strays too far from the diagonal (say, i is such that $\sigma(i)$ is much larger than i), then the typical number of indices $j > i$ for which $\sigma(j) > \sigma(i)$ is small; thus, the number of increasing subsequences containing i (and, hence, the probability that a random increasing subsequence contains such an i) cannot be very large. Presently, we make this intuition rigorous.

For brevity, we restrict ourselves to upper bounding the probability in (1). To do this, we union bound over all $i \in [n]$:

$$\Pr\left[\max_{i \in T} |\sigma(i) - i| > \epsilon n\right] \leq \sum_{i=1}^{n} \Pr[i \in T \cap |\sigma(i) - i| > \epsilon n].$$

For simplicity, we will only upper bound the probability

$$\Pr[i \in T \cap \sigma(i) - i > \epsilon n], \qquad (2)$$

since this gives an upper bound for the probability of the desired event by symmetry.

To calculate this probability, we sum over all permutations π satisfying $\pi(i) - i > \epsilon n$ and over all increasing subsequences t of π:

$$\Pr[\sigma(i) - i > \epsilon n \cap i \in T] = \sum_{\pi \,:\, \pi(i) - i > \epsilon n} \sum_{t} \Pr[\sigma = \pi \cap T = t]$$

$$= \frac{1}{n!} \sum_{\pi \,:\, \pi(i) - i > \epsilon n} \frac{\tau(\pi, i)}{\tau(\pi)}, \qquad (3)$$

where we have defined $\tau(\pi, i)$ to be the number of increasing subsequences of π that include the index i.

Let $A_\epsilon = A_\epsilon(\sigma, i)$ denote the event that $\sigma(i) - i > \epsilon n$. Then we can express (3) as a conditional expectation:

$$\Pr[\sigma(i) - i > \epsilon n \cap i \in T]$$
$$= \Pr[A_\epsilon(\sigma, i)] \cdot \mathbb{E}\left[\frac{\tau(\sigma, i)}{\tau(\sigma)} \;\middle|\; A_\epsilon(\sigma, i)\right]. \qquad (4)$$

To compute the remaining expectation, we note that $\tau(\sigma, i)$ can be expressed as a product as follows:

$$\tau(\sigma, i) = \beta(\sigma, i, \sigma(i)) \cdot \alpha(\sigma, i, \sigma(i)), \qquad (5)$$

where we define $\beta(\pi, i, j) = 1+$ # of nonempty increasing subsequences of π ending before i and whose last element x satisfies $\pi(x) < j$, and $\alpha(\pi, i, j) = 1+$ # of nonempty increasing subsequences of π beginning after i and whose initial element x satisfies $\pi(x) > j$.

Now, both β and α are constrained by the number of *indices* x for which, say, $x < i$ and $\pi(x) < \sigma(i)$. Thus, we introduce the following quantities:

$$L_{i,k} := |\{x \in [n] \;:\; x > i, \sigma(x) > k\}|,$$
$$B_{i,k} := |\{x \in [n] \;:\; x < i, \sigma(x) < k\}|.$$

Here, $B_{i,k}$ is hypergeometrically distributed with population size n, number of trials $i-1$, and number of successes $k-1$. Similarly, for any given b, $[L_{i,k} | B_{i,k} = b]$ is hypergeometrically distributed with population size $n-i$, number of trials $n-i$, and number of successes $n-k-i+b$.

With this in mind, we define an event W, under which $L_{i,\sigma(i)}$ and $B_{i,\sigma(i)}$ are well-behaved: for an arbitrary fixed $t > 0$,

$$W = [|B_{i,i+\epsilon n} - \mathbb{E}[B_{i,i+\epsilon n}]| < t(i-1)$$
$$\cap |L_{i,i+\epsilon n} - \mathbb{E}[L_{i,i+\epsilon n}| < t(n-i)]] \,.$$

Then from standard hypergeometric concentration results,

$$\Pr[W] \geq 1 - e^{-2t^2(i-1)} - e^{-2t^2(n-i)}. \qquad (6)$$

Then we can rewrite the expectation in (4) by conditioning on W:

$$\mathbb{E}[\tau(\sigma, i)/\tau(\sigma) | A_\epsilon(\sigma, i)]$$
$$= \mathbb{E}[\tau(\sigma, i)/\tau(\sigma) I[W] | A_\epsilon(\sigma, i)] + \mathbb{E}[\tau(\sigma, i)/\tau(\sigma) I[\neg W] | A_\epsilon(\sigma, i)]$$
$$\leq \mathbb{E}[\tau(\sigma, i)/\tau(\sigma) I[W] | A_\epsilon(\sigma, i)] + \Pr[\neg W | A_\epsilon(\sigma, i)],$$

where the inequality is from the fact that $\tau(\sigma, i)/\tau(\sigma) \leq 1$. The second term decays exponentially, by our choice of W, so we can focus on the first term.

To proceed, we will need a lemma to the effect that we can split the expected fraction into a ratio involving two expectations that are easy to bound:

LEMMA 3. *We have*

$$\mathbb{E}\left[\frac{\tau(\sigma, i)}{\tau(\sigma)} I[W] \middle| A_\epsilon(\sigma, i)\right]$$
$$\leq \frac{\mathbb{E}[\tau(\sigma, i) I[W] | A_\epsilon(\sigma, i)]}{2^{\mathbb{E}[L(\sigma)]}} + e^{-\Theta(n/\log^3 n)}. \qquad (7)$$

Here, $L(\sigma)$ *denotes the length of the longest increasing subsequence of* σ, *so the denominator becomes*

$$2^{\mathbb{E}[L(\sigma)]} = 2^{2\sqrt{n}(1+o(1))}. \qquad (8)$$

We can then estimate the numerator of the ratio (7) as follows:

$$\mathbb{E}[\tau(\sigma, i) I[W] \mid A_\epsilon(\sigma, i)]$$
$$= \mathbb{E}[\beta(\sigma, i, \sigma(i)) \cdot \alpha(\sigma, i, \sigma(i)) I[W] \mid A_\epsilon(\sigma, i)].$$

We can evaluate the remaining expression by observing that, conditioned on particular values for $B_{i,i+\epsilon n}$ and $L_{i,i+\epsilon n}$ (say, b and ℓ, which are constrained by T), $\beta(\sigma, i, \sigma(i))$ and $\alpha(\sigma, i, \sigma(i))$ are independent and have the same distribution as the number of increasing subsequences in random permutations from S_b and S_ℓ, respectively (the first and second moment of this distribution are given in [3]).

This implies that

$$\mathbb{E}[\tau(\sigma, i) I[W] | A_\epsilon(\sigma, i)]$$
$$\leq 2^{2\sqrt{i(i+\epsilon n)/n(1+O(t/n))} + 2\sqrt{n-(i+\epsilon n)-i+i(i+\epsilon n)/n}}.$$

Now, suppose that $i \sim cn$, for $c \in (0,1)$. Then the expression in the exponent becomes

$$2\sqrt{n}(\sqrt{c(c+\epsilon)} + \sqrt{1-2c-\epsilon+c(c+\epsilon)}),$$

and we want to choose ϵ so that this is $< 2\sqrt{n}$. We have the following:

$$\sqrt{c(c+\epsilon)} + \sqrt{1-2c-\epsilon+c(c+\epsilon)} < \sqrt{c(1+\epsilon/c)+(1-c)(1-\frac{\epsilon}{1-c})} = 1,$$

where the inequality is by Jensen and the concavity of the square root function, and it is strict because $\epsilon \in (0, 1-c)$.

Now, note that the function

$$F(c) = \sqrt{c(c+\epsilon)} + \sqrt{1-2c-\epsilon+c(c+\epsilon)}$$

satisfies $F(0) = 1-\epsilon$ and $F(1-\epsilon) = \sqrt{1-\epsilon}$ and is concave as a function of c.

This implies that *any* nonzero value of ϵ yields exponential decay of the probability $\Pr[\sigma(i) - i > \epsilon n \cap i \in T]$: i.e., it is at most $e^{-\Theta(\sqrt{n})}$, where the constant in the $\Theta(\cdot)$ is bounded away from 0 (for any fixed ϵ) for all i.

Now, to upper bound $\Pr[i - \sigma(i) > \epsilon n \cap i \in T]$, we note the following symmetry property: an increasing subsequence of σ containing i is equivalent to an increasing subsequence of the *counterclockwise rotation* of the permutation ρ corresponding to the graph of σ by π radians containing $n-i$. Moreover, $i - \sigma(i) > \epsilon n$ if and only if $\rho(n-i) - (n-i) > \epsilon n$, and ρ is uniformly distributed on S_n. This implies that

$$\sum_{i=1}^{n} \Pr[i \in T \cap \sigma(i) - i > \epsilon n] = \sum_{i=1}^{n} \Pr[i \in T \cap i - \sigma(i) > \epsilon n], \tag{9}$$

so that

$$\Pr\left[\max_{i\in T} |\sigma(i) - i| > \epsilon n\right] \leq 2\sum_{i=1}^{n} \Pr[i \in T \cap \sigma(i) - i > \epsilon n] \leq 2\sum_{i=1}^{n} 2^{-2\sqrt{n}(1-F(i/n))} = 2^{-2(1-\sqrt{1-\epsilon})\sqrt{n}(1+o(1))},$$

since $1 - F(c)$ is minimized at $c = 1-\epsilon$. This concludes the proof of the upper bound. ∎

Remark: Strictly speaking, we did not address the bounding of terms for which $i = o(n)$ or $i = n(1-o(n))$. These may be shown to be negligible compared to our stated upper bound by a slight variation of our argument.

5. PROOF OF THEOREM 2

We analyze an algorithm, which takes advantage of the following decomposition, for any increasing subsequence t of a given permutation σ:

$$\frac{1}{\tau(\sigma)} = \Pr[T^\sigma = t] = \prod_{i\in t} \Pr[i \in T^\sigma | \mathcal{F}_{i-1}(t)] \cdot \prod_{i\notin t}(1 - \Pr[i \in T^\sigma|\mathcal{F}_{i-1}(t)]),$$

where $\mathcal{F}_{i-1}(t)$ is the event that $T^\sigma \cap [i-1] = t \cap [i-1]$, and we denote by T^σ a uniformly random increasing subsequence of σ.

Recall that $\alpha(\pi, i, j)$ is the number of increasing subsequences of π beginning after i, whose initial element x satisfies $\pi(x) > j$, plus 1 for the empty sequence (for the special case of $i = j = 0$, we define $\alpha(\pi, 0, 0) = \tau(\pi)$). Then each conditional probability can be computed as

$$\Pr[i \in T^\sigma|\mathcal{F}_{i-1}(t)] = \frac{\alpha(\sigma, i, \sigma(i))}{\alpha(\sigma, \lambda_t(i), \sigma(\lambda_t(i)))}, \tag{10}$$

where $\lambda_t(i)$ denotes the largest $j < i$ for which $j \in t$ (when $i = 1$, $\lambda_t(i) = 0$, and we define $\sigma(0) = 0$ for convenience). Thus, the core of the algorithm is to compute these counts (i.e., $\alpha(\sigma, i, j)$ for $i, j \in [n]$) for a given σ by dynamic programming. Essentially, we use the fact that $\alpha(\sigma, i, j) = 1 + (\alpha(\sigma, i+1, j) - 1) \cdot (I[\sigma(i+1) > j] + 1)$.

function ComputeAlphaTable($\sigma \in S_n$)
  $A \leftarrow [0]^{(n+1)\times(n+1)}$
  for $i = 1$ to n **do**
    $A[i, n] \leftarrow 1$
    $A[n, i] \leftarrow 1$
  for $i = n-1$ to 0 **do**
    for $j = n-1$ to 0 **do**
      $A[i, j] \leftarrow 1+(A[i+1, j]-1)\cdot(I[\sigma(i+1) > j]+1)$
  return A

Using this, we can sample T using the following algorithm:

function SampleIncreasingSubsequence($\sigma \in S_n$)
  $A \leftarrow$ ComputeAlphaTable(σ)
  $t \leftarrow \emptyset$
  $\lambda \leftarrow 0$
  for $i = 1$ to n **do**
    With probability $A[i, \sigma(i)]/A[\lambda, \sigma(\lambda)]$, add i to t and set $\lambda = i$.
  return t

That the algorithm correctly samples from the set of increasing subsequences of σ can be verified easily by induction.

Now, computing A takes time $\Theta(n^2)$, and this is followed by a linear number of operations (since t can be represented as a linked list, we need only to append to the end at each step). Thus, the running time is $\Theta(n^2)$, as desired. It is simple to extend this algorithm to higher dimensional analogues (e.g., as in [1]). For dimension d, the running time becomes $\Theta(dn^{d+1})$.

6. REFERENCES

[1] B. Bollobás and P. Winkler. Box spaces and random partial orders. *Transactions of the American Mathematical Society*, 324(1):59–72, 1991.

[2] J. Deuschel and O. Zeitouni. On increasing subsequences of i.i.d. samples. *Combinatorics, Probability, and Computing*, 8:247–263, 1999.

[3] V. Lifschitz and B. Pittel. The number of increasing subsequences of the random permutation. *Journal of Combinatorial Theory, Series A*, 31:1–20, 1981.

[4] C. Papadimitriou. Concurrency control by locking. *SIAM J. Comput.*, 12(2):215–226, 1983.