Enumeration and shape of 1324-avoiding permutations

A permutation $\pi \in \mathcal{S}_n$ is said to avoid a pattern $\sigma \in \mathcal{S}_k$ if there is no increasing subsequence $1 \le i_1 < \cdots < i_k \le n$ such that the subsequence $(\pi_{i_1}, \ldots, \pi_{i_k})$ has the same order type as σ (meaning $\pi_{i_r} < \pi_{i_s}$ iff $\sigma_r < \sigma_s$). In particular, $\pi \in \mathcal{S}_n$ is 1324-avoiding if there is no $1 \le a < b < c < d \le n$ such that $\pi_a < \pi_c < \pi_b < \pi_d$. For $\pi \in \mathcal{S}_n$, its permutation matrix $M(\pi)$ is the $n \times n$ matrix with $M(i,j) = 1$ if $\pi_i = j$ and Zero otherwise. Then π avoids 1324 iff $M(\pi)$ has no 4×4 minor (choose any four rows and any four columns) equal to $M(1324)$. Let $\mathcal{A}_n$ denote the set of 1324-avoiding permutations of length n and let $N_n := N_n(1324) := |\mathcal{A}_n|$ count the number of 1324-avoiding permutations of length n. The Marcus-Tardos theorem affirms the existence of a finite limit

$$c := c_{1324} := \lim_{n\to\infty} N_n(1324)^{1/n} .$$

Submultiplicativity immediately implies that $N_n(1324)^{1/n} \le c$ for all n. Best current bounds are $9.81 < c < 13.74$, the lower bound due to Bevan (2015) and the upper to Bona (2015). These improved upon bounds of 9.35 (Albert et al, 2006) and $7 + 4\sqrt{3} \approx 13.93$ (Bona, 2012) respectively.

Problem 1. *Improve these bounds.*

Problem 2. *Shown that N_{n+1}/N_n converges to c.*

For $\pi \in \mathcal{A}_n$, let $L(\pi)$ denote the length of the longest initial segment of π that avoids 132. There is a simple bijection from pairs (π, m) with $\pi \in \mathcal{A}_n$ and $0 \le m \le L(\pi)$ to $\mathcal{A}_{n+1}$ obtained by inserting the element $n+1$ into π in position $m+1$. Let μ_n denote the uniform measure on $\mathcal{A}_n$ and let $\mathbb{E}_n$ denote expectation with respect to μ_n. Then Problem 2 is equivalent to showing that $\mathbb{E}_n(L+1) \to c$.

Let $M(n) := \mathbb{E}_n M$ denote the average permutation matrix over $\mathcal{A}_n$. The left-hand figure shows an MCMC simulation of $M(350)$. The only known shape result (Poh, 2015) is the exponential region of decay depicted in black in the figure on the right.

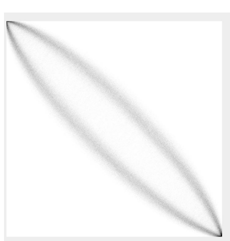

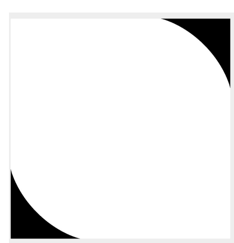

Problem 3. *Does all the mass in $M(n)$ converge to the diagonal as $n \to \infty$?*

Problem 4. *Let R_n be the least value such that $\sum_{k=n/2-R}^{n/2+R} M(n)_{n/2,k} \ge 1/2$. In other words, R_n is the μ_n-median absolute value of $\pi_{n/2} - n/2$. How does R_n scale with n? This is a refinement of the previous problem, which asks whether $R_n = o(n)$.*

For $\pi \in \mathcal{A}_{2n}$, let $\pi|_n$ denote the subsequence of π consisting of the values less than or equal to n. Thus, necessarily, $\pi|_n \in \mathcal{A}_n$.

Problem 5. *What is the typical value of* $\mathbb{E}_{2n} L(\pi|_n)$*?*

On the one hand, we know $\mathbb{E}_{2n} L$ converges to c at least in as a logarithmic cesaro sense even if Problem 2 fails. On the other hand, the data makes it appear that $\mathbb{E}_{2n} L(\pi|_n)$ may be of order n (if Problem 3 fails) and in any case seems extremely unlikely to be $O(1)$.

The next figure shows a sample permutation in $\mathcal{A}_{100}$ (red) against the intensity plot of $M(100)$. It is known (Claesson et al., 2012) that permutations in $\mathcal{A}_n$ split into two subsequences, one avoiding 132 and one avoiding 213. The lower arc of the red permutation avoids 132 and the upper arc avoids 213. The set of 132-avoiding permutations is counted by Catalan numbers and its shape is well understood (Miner and Pak, 2014).

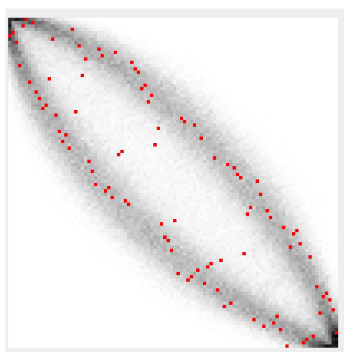

Problem 6. *Can one characterize the distribution of the* 132*-avoiding permutation obtained by taking the lower arc of a typical element of* $\mathcal{A}_n$*?*

The previous figure was generated by an MCMC algorithm which attempts to switch the values of $\pi(i)$ and $\pi(j)$, rejecting if this forms a 1324 pattern to occur. The Markov chain is doubly stochastic, hence uniform on $\mathcal{A}_n$. The chain was run until various functionals appeared to be roughly in stationarity.

Problem 7. *What is the mixing time of the Metropolis-Hastings chain on* $\mathcal{A}_n$ *moving by random transpositions?*

Deletion channels

A deletion channel is a map whose input is a string of message bits and whose output is a string with some of the bits deleted, but no indication as to which. Formally, suppose $\{X_n\}$ are a string of message bits, assumed to be IID Bernoulli(1/2), and let $\{U_n\}$ be independent uniform $[0, 1]$ random variables which we use to keep or retain each bit according to a tunable deletion parameter p. The n^{th} retained bit is $Z_n := X_{\tau_n}$ where $\tau_n = \inf\{k : S_k \geq n\}$ and S_k are the partial sums of retentions $S_k := \sum_{i=1}^{k} \mathbf{1}_{U_i \leq 1-p}\}$.

For $N \geq 1$, let $\mathbf{X}^N := (X_1, \ldots, X_N)$ be the first N bits of message and let $\mathbf{Z}^N := (Z_1, \ldots, Z_{S_N})$ be the string of bits received when the first N bits are sent. The main problem is to determine the rate at which information is transmitted through the channel. The amount of information that $\mathbf{Z}^N$ reveals about $\mathbf{X}^N$ is $h_N := h(\mathbf{X}^N) + h(\mathbf{Z}^N) - h(\mathbf{Z}^N, \mathbf{X}^N)$.

Problem 1. *Determine the transmission rate* $h := h(p) := \lim_{N\to\infty} \dfrac{h_N}{N}$.

This problem and related ones are discussed at length by Mitzenmacher (Probability Surveys, 2009). This problem statement could just be a pointer to that reference. However, I will try to summarize what I think are the most interesting sub-problems. One is to improve the upper and lower bounds. The best known upper bounds are roughly $h(p) \leq (4/5)(1 - p)$ (see page 20 of Mitzenmacher, who refers to Diggavi et al. and Feratoni et al.). The best known lower bound (Drinea and Mitzenmacher) is $h(p) \geq (1 - p)/9$; this is notable because it shows a transmission rate proportional to the retention rate even when this is near zero, a fact which is not at all obvious.

Problem 2. *Improve either bound.*

Of greater interest, perhaps, is the question of how to extract what information is there. In principle, this is the same as the question of computing the distribution of $\mathbf{X}^N$ given $\mathbf{Z}^N$. The case $p = 1/2$ allows for a particularly simple version of this problem. Because $\mathbf{X}^N$ is uniform on all words of length 2^N and the deletion is uniform over all substrings, this boils down to computing the distribution of $\#S(\mathbf{X}^N, \mathbf{Z}^N)$, the number of ways in which $\mathbf{Z}^N$ appears as a substring of $\mathbf{X}^N$. Ander Holroyd and I have some partial results related to the following problem − ask if interested.

Problem 3. *Characterize the distribution of* $\#S(\mathbf{X}^N, \mathbf{Z}^N)$ *given* $\mathbf{Z}^N = W$*, for typical words* W*.*

Existing lower bounds for general p involve constructive extraction of information and sometimes construction of codebooks. One coding schemem that appears in Mitzenmacher's survey is Poisson repetition: the message string is obtained from an initial code word by replaing each bit by a string of consecutive bits equal to it, with block lengths chosen to be IID Poissons of some mean λ.

Problem 4 (Mitzenmacher, 2009)**.** *Find an efficient encoding and decoding algorithm for a Poisson-repeat channel with parameter, say, 1.*