Universitat Politècnica de Catalunya
Facultat d'Informàtica de Barcelona

Degree: Grau en Enginyeria Informàtica **Academic year:** 2023–2024 (Final Exam)
Course: Randomized Algorithms (RA-MIRI) **Date:** January 18th, 2024
Time: 2h 30m

1. **(2.5 points)** We need to send a signal S which might be $S = -1$ or $S = +1$ over a wireless network. Because of other sources emitting signals S_i, $1 \le i \le n$, at the same time, the received signal R can be expressed as

$$R = S + \sum_{i=1}^{n} p_i S_i,$$

where the $p_i \ge 0$ measures the strength of signal S_i; the p_i's are not probabilities, since $\sum_i p_i$ might be $\neq 1$. If $R > 0$, we assume that the original signal $S = +1$; conversely, if $R < 0$ then we assume that $S = -1$ (if $R = 0$, we choose at random). We want to bound the probability of that we identify S wrongly. That will happen whenever $|R - S| > 1$.

Let $X = \sum_{i=1}^{n} p_i S_i$ denote the "'noise", and assume the S_i's are i.i.d. with

$$\mathbb{P}[S_i = +1] = \mathbb{P}[S_i = -1] = \frac{1}{2}, \qquad 1 \le i \le n.$$

(a) Compute $\mathbb{E}[X]$ and $\mathbb{V}[X]$.

(b) Compute the moment generating function $\mathbb{E}\left[e^{tX}\right]$ and show that it is bounded by $e^{(\sum_i p_i^2)t^2/2}$. Useful formula: $(e^x + e^{-x})/2 \le e^{x^2/2}$ (it can be shown using the Taylor series expasions of both sides of the inequality).

(c) Using Markov's inequality we can derive a Chernoff-like bound as

$$\mathbb{P}[X \ge a] = \mathbb{P}\left[e^{tX} \ge e^{at}\right] \le \mathbb{E}\left[e^{tX}\right] e^{-at}.$$

Use the bound on $\mathbb{E}\left[e^{tX}\right]$ and set $t = 1/\sum_i p_i^2$ to obtain an exponentially decaying upper bound for $\mathbb{P}[X \ge a]$.

(d) Using analogous arguments, the upper bound above also applies to $\mathbb{P}[-X \ge a]$, and then we can combine this result to obtain a bound for $\mathbb{P}[|X| \ge a]$. Using that bound, give a lower bound for the probability of a correct identification of S.

---

2. **(2.5 points)** We have a computer monitoring a sensor, by requesting data from the sensor from time to time. It does so at randomly picked moments, to avoid any easily predictable pattern which could be exploited by a malicious adversary. However, we are guaranteed that the computer will monitor the sensor $\lambda = 3$ times on each interval of 10 minutes on average (we will call a *time frame* or just a *frame*, each such 10-minutes interval).

   (a) Give a formula for the probability that the computer monitors the sensor exactly j times in a frame. To compute it, consider that the frame is subdivided in a big number n tiny time intervals, each one a potential moment in which the computer issues a monitoring request to the sensor. Thus, each tiny time interval contains, with some probability, a monitoring request, independently of the others, and of those n intervals, on average, λ of them have monitoring requests (and the other don't).

   (b) If there is no monitoring request during a frame we say that it is *non-monitored.* Consider m consecutive non-overlapping frames. What is the expected number of non-monitored frames? Let Y denote the number of non-monitored frames out of m. Prove

$$\mathbb{P}\Big[|Y - \mathbb{E}[Y]\,| \geq b\sqrt{\mathbb{E}[Y]}\Big] \leq \frac{1}{b^2}.$$

---

3. **(2.5 points)** A certain city has N bus lines numbered $1, 2, \ldots, N$. Walking around the city you have seen buses with numbers $1 \leq i_1 \leq i_2 \leq \cdots \leq i_k \leq N$. You might have observed less than k different bus lines, because you could have observed more than one bus of the same line. You do not know N, that is, how many bus lines there are in the city, but you can give an estimate $\hat{N}$ of N as a function of k and the observed numbers $i_1, \ldots, i_k$, such that $\mathbb{E}\big[\hat{N}\big] \sim N$. Here, the expectation is on the sample of k lines that you have observed; each one of the N^k possible choices is assumed equally likely.

   (a) Compute the probability that $X \equiv i_k$, the largest of k randomly drawn numbers from $\{1, \ldots, N\}$ is $\leq j$, for $1 \leq j \leq N$. The k draws are independent and "with replacement" as any particular bus line can be observed several times.

   (b) Compute the expected value of X. To that end, prove first that $\mathbb{E}[X] = \sum_{1 \leq j < N} \mathbb{P}[X > j]$. Useful fact: $\sum_{i=1}^{n} i^r = \frac{n^{r+1}}{r+1} + \mathcal{O}(n^r)$.

   (c) Propose an asymptotically unbiased estimator $\hat{N}$ for N: $\hat{N} := f(k, X)$ and $\mathbb{E}\big[\hat{N}\big] = N + o(N)$ as $N \to \infty$.

---

4. **(2.5 points)** Modern hardware tries to optimize the execution of instructions in a pipelined fashion by predicting on each conditional instruction which of the two branches will be taken. Many solutions have been proposed, but branch predictions must be carried out at a very low level, so very sophisticated solutions must be avoided. One such mechanism is using a finite automaton that keeps information about the behavior of the conditional instruction on the last k times it has been executed. One such particular automaton for $k = 2$ is the so-called 2-bit flip-on-consecutive counter. To analyze the performance of this branch prediction mechanism we are lead to consider the Markov chain below

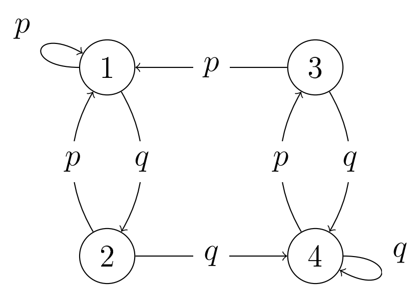

where $0 \leq p \leq 1$ and $q = 1 - p$.

(a) Write the transition matrix $P^{(2)}$ for two steps of the Markov chain. That is, $p_{uv}^{(2)}$ is the probability that we are at state v after two steps of the Markov chain if we started at state u, for all u and v.

(b) Find the stationary distribution $\pi = \pi(p)$ for the Markov chain. Identities such as $p^2q + pq^2 = pq$ or $p^2 + q = 1 - pq$ might be helpful here and in the next question.

(c) Compute a closed form for the probability of a misprediction, which is, by definition

$$P_{\text{misprediction}} = \pi(p) \cdot (q, q, p, p)^{\mathrm{T}}$$

Prove that $P_{\text{misprediction}} = 0$ if $p = 0$ or $p = 1$. Prove also that it is maximum if $p = q = 1/2$; for that case, $P_{\text{misprediction}} = 1/2$.
