Problem Set 1
Tail Inequalities and Applications

1. Suppose that n balls are independently and uniformly distributed in n bins.
   (a) Show that for large n, the expected number of empty bins approximates to n/e.
   (b) Consider m balls being distributed in n bins. What is the expected number of empty bins?
2. Let X be a random variable with expectation μ and standard deviation σ. Show that, for any $t \in \mathbb{R}^+$
   (a) $\mathbf{Pr}[X - \mu \geq t\sigma] \leq \dfrac{1}{1+t^2}$.
   (b) $\mathbf{Pr}[|X - \mu| \geq t\sigma] \leq \dfrac{2}{1+t^2}$.
   (c) How do the above compare with the bounds obtained by Chebyshev's inequality?
3. Let X be a non-negative integer-valued random variable with positive expectation. Prove the following inequalities.
   (a) $\mathbf{Pr}[X = 0] \leq \dfrac{\mathbf{E}[X^2] - \mathbf{E}[X]^2}{\mathbf{E}[X]^2}$.
   (b) $\dfrac{\mathbf{E}[X]^2}{\mathbf{E}[X^2]} \leq \mathbf{Pr}[X \neq 0] \leq \mathbf{E}[X]$.
4. The *weak law of large numbers* states that, for random variables $\{X_i\}_{i=1}^n$ distributed identically and independently with expectation μ and variance σ^2, we have for any $\varepsilon > 0$,

   $$\lim_{n\to\infty} \mathbf{Pr}\left[\left|\frac{\sum_{i=1}^n X_i}{n} - \mu\right| < \varepsilon\right] = 1.$$

   Prove the weak law of large numbers using Chebbyshev's inequality.
5. Derive the simpler forms of Chernoff bounds (as shown below) from the bounds we obtained in class.

   $$\mathbf{Pr}[X > (1+\delta)\mu] \leq e^{-\delta^2\mu/(2+\delta)} \text{ for } \delta \geq 0, \quad (1)$$

   $$\mathbf{Pr}[X < (1-\delta)\mu] \leq e^{-\delta^2\mu/2} \text{ for } \delta \in (0,1]. \quad (2)$$

   $$\mathbf{Pr}[|X - \mu| \geq \delta\mu] \leq e^{-\delta^2\mu/3} \text{ for } \delta \in (0,1]. \quad (3)$$
6. Consider the following occupancy problem (related to the coupon collector's problem): There are n bins and n players, each player having an infinite supply of balls. The bins are all initially empty. There are a sequence of rounds: in each round, each player throws a ball into an empty bin chosen independently at random from all currently empty bins. Let the random variable Y be the number of rounds before every bin is non-empty. Determine the expected value of Y. What can you say about the tail of Y's distribution?

7. (a) **[Chernoff Bound for a Special Case]** Let $X_1, X_2, \ldots, X_n$ be independently distributed random variables with

$$\mathbf{Pr}[X_i = 1] = \mathbf{Pr}[X_i = -1] = \frac{1}{2}.$$

Let $X = \sum_{i=1}^{n} X_i$. Show that for any $a > 0$,

$$\mathbf{Pr}[X > a] \leq e^{-\frac{a^2}{2n}}.$$

(b) **[Set Balancing]** A set S of m objects, each having zero or more of n potential features is specified by an $n \times m$ matrix $\mathbf{A}$ over $\{0, 1\}$, where $a_{ij} = 1$ iff j-th object has the i-th feature. A partition of the set into S_1, S_2 is given by a vector $\vec{b} \in \{-1, 1\}^m$, with 1 (resp. -1) in position j indicating the presence of object j in S_1 (resp. S_2). If we let $c = \mathbf{A}\vec{b}$, then $|c_i|$ denotes the imbalance in feature i. The imbalance of a partition $\vec{b}$ is given by $\|\mathbf{A}\vec{b}\|_\infty = \max_{i \in [n]} |c_i|$. The set balancing problem is to find a partition $\vec{b}$ that minimizes the imbalance $\|\mathbf{A}\vec{b}\|_\infty$. Consider the following algorithm: choose entries of $\vec{b}$ uniformly and independently at random from $\{-1, 1\}$, completely ignoring $\mathbf{A}$. Show that

$$\mathbf{Pr}[\|\mathbf{A}\vec{b}\|_\infty > \sqrt{4m \ln n}] \leq \frac{2}{n}.$$

Use the bound from the previous part.

8. **[Permutation Routing]** Let n be a positive integer and let $N = 2^n$. A hypercube is a undirected graph over the vertices $\{0, 1, 2, \ldots, N-1\}$. Vertices $x, y \in \{0, 1, 2, \ldots, N-1\}$ are connected by an edge if and only if the Boolean representations of x and y differ in exactly one bit position. In the permutation routing problem, every vertex in the hypercube has exactly one packet to send to some other vertex and receives exactly one packet. The problem hence has an associated permutation π on the set $\{0, 1, 2, \ldots, N-1\}$ such that every vertex x sends a packet to $\pi(x)$ and receives a packet from $\pi^{-1}(x)$.

Every node can send and receive simultaneously. However, at each time step (time is discrete for this problem), at most one packet can be sent along an edge. Hence if some vertex wants to send two packets along the same edge, one packet has to wait till the next time step. The objective of the permutation routing problem is to route all the packets in minimum number of time steps.

A natural algorithm for this problem is *bit fixing*. Consider a packet that needs to be sent from x (its current location) to y. Let $x = (x_{n-1}, x_{n-2}, \ldots, x_0)_2$ and let $y = (y_{n-1}, y_{n-2}, \ldots, y_0)_2$. Let i be the minimum index such that $x_i \neq y_i$ (such an i is guaranteed to exist if $x \neq y$). Let $x' = (x_{n-1}, \ldots, x_{i+1}, y_i, x_{i-1}, \ldots, x_0)_2$. Then the packet is sent along the edge $\{x, x'\}$. Observe that the routing algorithm does not consider other packets for deciding its route. Such routing algorithms are called *oblivious* routing algorithm. Then prove the following.

(a) **[Bit fixing is not good in worst case]** Give an instance of the problem where bit fixing algorithm takes $\Omega(\sqrt{N}/n)$ time steps.

(b) **[2-Phase routing via random intermediate destination]** The idea to bypass worst case instances is to first route packets to a random intermediate instances and then route them to the actual destinations. Formally, for every vertex x, let $\sigma(x)$ denote a random element chosen from $\{0, 1, 2, \ldots, N-1\}$. Now for every x, we first route the packet for x from x to $\sigma(x)$ and then from $\sigma(x)$ to $\pi(x)$ using bit fixing. Prove that, on every instance, the expected number of time steps that the modified algorithm takes is $O(n)$.