Exam 1

CSCI 4030 / CSCI 6220: Randomized Algorithms

Fall 2026

Rensselaer Polytechnic Institute

Name: _____ **RIN:** _____

Instructions. Answer all questions. You have 100 minutes. You may use one 8.5×11 inch double-sided crib sheet. No electronic devices are permitted. Show your reasoning clearly; unsupported answers may receive little or no credit. You may use any of the inequalities below without proof.

Useful inequalities. For any nonnegative random variable X and $t > 0$,

$$\Pr(X \geq t) \leq \frac{\mathbb{E}[X]}{t} \qquad \text{(Markov)}.$$

For any random variable X with finite variance and $t > 0$,

$$\Pr\big(|X - \mathbb{E}[X]| \geq t\big) \leq \frac{\mathrm{Var}(X)}{t^2} \qquad \text{(Chebyshev)}.$$

For events $A_1, \ldots, A_m$,

$$\Pr\left(\bigcup_{i=1}^{m} A_i\right) \leq \sum_{i=1}^{m} \Pr(A_i) \qquad \text{(union bound)}.$$

If $X = \sum_i X_i$ is a sum of independent Bernoulli random variables and $\mu = \mathbb{E}[X]$, then for $0 < \delta \leq 1$,

$$\Pr\big(X \geq (1+\delta)\mu\big) \leq \exp\left(-\frac{\delta^2 \mu}{3}\right), \qquad \Pr\big(X \leq (1-\delta)\mu\big) \leq \exp\left(-\frac{\delta^2 \mu}{2}\right).$$

[12 points] Problem 1 (How much randomness do we really need?).

A randomized-algorithms library was originally implemented using fully independent random choices. Because random bits and stored random seeds are resources too, an engineer wants to replace full independence by weaker randomness whenever the analysis still goes through.

For parts (a)–(c), identify the *weakest* assumption among

no independence, pairwise independence, 3-wise independence, full independence

that is sufficient for the stated argument. Give a brief justification.

(a) (2 points) The running time is $T = \sum_{i=1}^{n} X_i$, and the analysis only computes $\mathbb{E}[T]$ from the individual expectations $\mathbb{E}[X_i]$.

(b) (3 points) The analysis uses $\mathrm{Var}(\sum_{i=1}^{n} X_i) = \sum_{i=1}^{n} \mathrm{Var}(X_i)$.

(c) (3 points) Three specified distinct keys x, y, z are hashed into m buckets, and the analysis needs

$$\Pr\big(h(x) = h(y) = h(z)\big) = \frac{1}{m^2}.$$

(d) (4 points) A Las Vegas algorithm always returns the correct answer and has expected running time at most T. We convert it to a Monte Carlo algorithm by stopping it after $10T$ steps and returning FAIL if it has not finished. Give an upper bound on the probability that the new algorithm returns FAIL, and justify your answer.

[12 points] Problem 2 (How often do we update the current best?).

A system has n candidate configurations with distinct performance scores. The configurations are tested in a uniformly random order.

The first configuration tested is deployed. After that, whenever a newly tested configuration has a higher score than *all configurations tested before it*, the system replaces the currently deployed configuration with the new one.

Let R be the total number of deployments, including the initial deployment.

(a) (7 points) For $i = 1, \ldots, n$, let

$$X_i = \begin{cases} 1, & \text{if the } i\text{th configuration tested causes a deployment,} \\ 0, & \text{otherwise.} \end{cases}$$

Express R in terms of $X_1, \ldots, X_n$, and compute $\Pr(X_i = 1)$.

Hint: The scores of the first i configurations tested are distinct and are in uniformly random order.

(b) (5 points) Compute $\mathbb{E}[R]$ exactly as a sum, and give its asymptotic order in n.

[28 points] Problem 3 (Keeping a server farm balanced).

A service has n identical servers and receives n jobs. As a very low-overhead routing rule, each arriving job chooses one server independently and uniformly at random. Let L_j be the final load of server j.

You may use, without proof, the standard estimate

$$\binom{n}{k} \le \left(\frac{en}{k}\right)^k, \qquad 1 \le k \le n.$$

(a) (6 points) Fix a server j. Show that for every integer $1 \le k \le n$,

$$\Pr(L_j \ge k) \le \binom{n}{k}\left(\frac{1}{n}\right)^k \le \left(\frac{e}{k}\right)^k.$$

For the first inequality, use the following observation: if $L_j \ge k$, then there is some set of k jobs that all chose server j.

(b) (5 points) Suppose k is chosen so that

$$\left(\frac{e}{k}\right)^k \le \frac{1}{n^2}.$$

Deduce from part (a) that

$$\Pr\left(\max_{1\le j\le n} L_j \ge k\right) \le \frac{1}{n}.$$

Remark. For all sufficiently large n, $k = \lceil 4\log n/\log\log n\rceil$ satisfies the displayed condition. One way to verify this is to use $\log k - 1 \ge \frac{1}{2}\log\log n$. Thus one random choice gives maximum load $O(\log n/\log\log n)$ with high probability.

(c) (3 points) Now change the routing rule: each arriving job samples $d \geq 2$ servers independently and uniformly at random and joins a least-loaded sampled server. After t jobs have been placed, suppose at most a β_i fraction of the servers have load at least i. What is the probability that all d servers sampled by job $t+1$ have load at least i?

(d) (6 points) For $i \geq 0$, let $B_i(t)$ be the number of servers with load at least i after t jobs have been placed, and write $B_i = B_i(n)$ for the final value. First show deterministically that

$$B_4 \leq \frac{n}{4}.$$

Set $\beta_4 = 1/4$ and, for $i \geq 4$, define

$$\beta_{i+1} = 2\beta_i^d.$$

Suppose that throughout the process $B_i(t) \leq \beta_i n$. Using part (c), explain why each arriving job has conditional probability at most β_i^d of being placed at load level at least $i+1$. Why does this make $\beta_{i+1} n = 2\beta_i^d n$ a natural threshold for B_{i+1}?

(e) (8 points) Define

$$E_i = \{B_i \le \beta_i n\}.$$

You may assume the concentration step from the lecture analysis: for every relevant level $i \ge 4$,

$$\Pr\big(E_{i+1}^c \mid E_4 \cap E_5 \cap \cdots \cap E_i\big) \le \frac{1}{n^3}.$$

Confirm that part (d) above implies that after

$$r = O\left(\frac{\log\log n}{\log d}\right)$$

additional levels, the recurrence above reaches a level r for which $\beta_r n < 1/2$.

Show that

$$\Pr(E_4 \cap E_5 \cap \cdots \cap E_r) \ge 1 - \frac{r-4}{n^3},$$

and use $\beta_r n < 1/2$ to deduce a high-probability upper bound on the maximum server load. State the resulting asymptotic dependence on n and d when we place a job on the least loaded of d uniformly randomly selected servers.

[22 points] Problem 4 (Reliable estimates from cheap random seeds).

A streaming algorithm maintains a small randomized sketch of a massive data stream in order to estimate F_0, the number of distinct items in the stream. One copy of the sketch returns an estimate within a factor $1 \pm \varepsilon$ of F_0 with probability at least $3/4$.

To increase confidence, we run k copies of the sketch and return the median of their estimates. If more than half of the copies are successful, then the median is also a $(1 \pm \varepsilon)$-approximation. Generating k fully independent hash seeds is costly, so an engineer proposes using seeds whose resulting success indicators are only *pairwise independent*.

Let

$$X_i = \mathbf{1}\{\text{copy } i \text{ is successful}\}, \qquad X = \sum_{i=1}^{k} X_i.$$

You may assume k is even, and count $X = k/2$ as a failure.

(a) (4 points) Compute $\mathbb{E}[X]$.

(b) (6 points) Compute $\mathrm{Var}(X)$. Be explicit about where pairwise independence is used.

(c) (6 points) Use Chebyshev's inequality to show that the median estimate fails with probability at most $3/k$.

(d) (4 points) If the copies were fully independent, use the supplied Chernoff bound to show that the median estimate fails with probability at most $\exp(-k/24)$.

(e) (2 points) Suppose the target failure probability is at most $1/n$. Determine, up to constant factors, the number of copies k sufficient under pairwise independence and under full independence. Which assumption reaches the target with asymptotically fewer copies?

[26 points] Problem 5 (A hash table that looks random locally but can fail globally).

A dictionary stores p distinct keys, where $p \geq 3$ is prime, in p buckets using separate chaining. Storing a truly random hash function is impractical, so we want a compact randomized hash family.

Part I: What does a universal-family guarantee buy us?

Suppose first that h is drawn from a universal family, so that for every pair of distinct keys x, y,

$$\Pr\big(h(x) = h(y)\big) \leq \frac{1}{p}.$$

For a fixed stored key x, let C_x be the number of stored keys in the same chain as x, including x itself.

(a) (4 points) Express C_x as 1 plus a sum of collision indicators.

(b) (4 points) Show that $\mathbb{E}[C_x] < 2$.

(c) (3 points) Use Markov's inequality to give an explicit upper bound on $\Pr(C_x \geq 6)$.

(d) (4 points) An engineer concludes: "Part (c) shows that every chain is shorter than 6 with high probability." Determine the upper bound that parts (a)–(c) actually give on

$$\Pr\Big(\max_x C_x \geq 6\Big).$$

Explain why this bound does *not* establish the engineer's claim.

Problem 5 continued.

Part II: Pairwise-random outputs can hide global structure.

Now identify both the keys and buckets with the finite field $\mathbb{F}_p$ and choose a, b independently and uniformly from $\mathbb{F}_p$. Define

$$h_{a,b}(x) = ax + b \pmod{p}.$$

(e) (5 points) Let $x \neq y$ and fix arbitrary buckets $u, v \in \mathbb{F}_p$. Show that there is exactly one pair $(a, b) \in \mathbb{F}_p^2$ satisfying

$$ax + b = u, \qquad ay + b = v.$$

Deduce that $h(x)$ and $h(y)$ are independent and uniform over $\mathbb{F}_p$.

(f) (3 points) What happens to the hash table when $a = 0$? Determine the probability of this event and use it to give a lower bound on

$$\Pr\big(\text{maximum chain length} = p\big).$$

(g) (3 points) For a fixed bucket j, write

$$L_j = \sum_{x \in \mathbb{F}_p} \mathbf{1}\{h(x) = j\}.$$

The summands are pairwise independent. Identify the exact step in the standard MGF proof of a Chernoff bound at which replacing full independence by pairwise independence is not sufficient, and explain why.