Midterm Exam

- Don't Panic.
- The midterm contains six problems (and one just for fun). You have 120 minutes to earn 100 points.
- The midterm contains 16 pages, including this one and 3 pages of scratch paper.
- The midterm is closed book. You may bring one double-sided sheet of A4 paper to the midterm. (You may not bring any magnification equipment!) You may not use a calculator, your mobile phone, or any other electronic device.
- Write your solutions in the space provided. If you need more space, please use the scratch paper at the end of the midterm. Do not put part of the answer to one problem on a page for another problem.
- Read through the problems before starting. Do not spend too much time on any one problem.
- Show your work. Partial credit will be given. You will be graded not only on the correctness of your answer, but also on the clarity with which you express it. Be neat.
- Draw pictures and give examples.
- Good luck!

Problem #	**Name**	**Possible Points**	**Achieved Points**
1	True, False, Explain	15	
2	Faulty Servers	10	
3	Minimum Cuts for Fun and Profit	12	
4	Room Allocation	30	
3	Morris Likes to Count	16	
6	The World Congress	17	
Total:		100	

Student Number: ______________________________

Problem 1. True, False, and Explain [15 points]

For each statement, indicate whether it is true or false, and briefly explain why. (No credit will be given for a blank or incorrect explanation.)

For any random variable X, assume that $\mathrm{E}[X] \geq \sqrt{\mathrm{Var}[X]} > 0$. Then $\Pr[X \geq 3\mathrm{E}[X]] \leq 1/4$. **TRUE** **FALSE**

Explanation:

For every non-negative random variable X and for every value a, $\Pr[X \geq a] \geq \Pr[X^2 \geq a^2]$. **TRUE** **FALSE**

Explanation:

Stochastic domination is transitive: given three independent random variables A, B, and C, if A stochastically dominates B and B stochastically dominates C, then A stochastically dominates C.

TRUE **FALSE**

Explanation:

The expected search time for Cuckoo Hashing is asymptotically faster than the expected search time for a hash table with chaining (where the table size $m = \Theta(n)$, the number of keys).

TRUE **FALSE**

Explanation:

If you throw n balls randomly into n bins, where $n > 2$, then the expected number of empty bins is $< n/9$.

TRUE **FALSE**

Explanation:

Problem 2. Faulty Servers [10 points]

You are the system administrator for a collection of n servers. Unfortunately, every so often, a server fails. It's your job to detect when servers fail and identify them. To check if a server is running properly, you can *ping* it, sending it a message and getting a response. Unfortunately, communication is faulty, so sometime your ping fails.

- If the server has failed, then there is no response to the ping.
- If the server is functional, then with probability 1/2 you get a response to your ping; with probability 1/2 you get no response.

Thus, when there is no response to a ping, you cannot be sure whether the server failed or whether the ping failed.

This morning, an alarm goes off, indicating that (exactly) one server has failed! Alas, you do not know which server has failed. (You may assume that a randomly chosen server failed.)

Assume that $n = 4$, and you ping server X and get no reponse.

What is the probability that server X has failed? $\boxed{}$

Show that your answer is correct:

Problem 3. Minimum Cuts for Fun and Profit [12 points]

Consider the following graph G containing n nodes and n edges, where $n-1$ nodes are arranged in a ring and one node is attached outside:

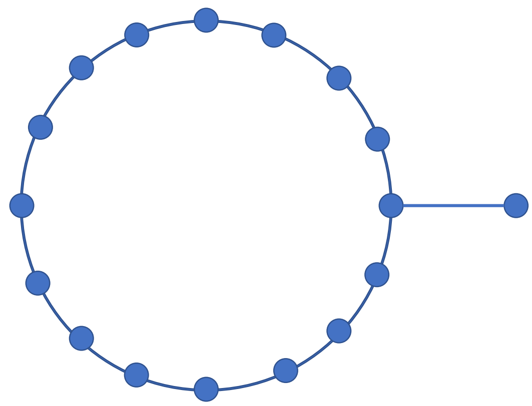

Problem 3.a. What is the minimum cut of G?

Problem 3.b. Prove (carefully, in detail) that the probability that one execution of the contraction algorithm (i.e., executing `Collapse`$(G, n, 2)$, reducing the number of nodes to two) in Karger's Min-Cut algorithm has probability exactly $2/n$ of successfully finding the minimum cut.

Problem 4. Room Allocation [30 points]

Professor Rogammer has n study rooms and n high school students, each of whom needs to be assigned to a room. Her goal is to assign students to study rooms in a uniform fashion, so that no study room has too many students. To do this, she wants to use a random load balancing strategy: when each student arrives, she chooses a room uniformly at random and assigns the student to the room.

It turns out that all the students arrive in pairs. (They were each assigned partners in class. Assume n is even.) Some of the pairs want to work together, while others do not. When a pair of students arrives:

- With probability 1/2, the pair want to study together. Professor Stubbins chooses a random room and assigns them both to the same room.
- With probability 1/2, the pair does not want to study together. Professor Stubbins assigns each of the two students to randomly chosen rooms, independently of each other. (They may, by chance, still end up in the same room, of course.)

Choose one room that we want to analyze (maybe, room number seven). Let X be the number of students assigned to that room, after all n students have arrived.

Problem 4.a. What is the expected value of X?

Show that your answer is correct:

Problem 4.b. What is the variance of X? $\boxed{}$

Show that your answer is correct:

Problem 4.c. Use Chebychev's Inequality to upper bound the probability that the room has more than 9 students. For this approximation, you may assume that $n \geq 5$, which should simplify your variance calculation from the previous part. (If you were not able to solve the previous part, then for partial credit you may use V to represent the variance of X.)

Problem 4.d. Now prove that *every* room has $O(\log n)$ students, with high probability. (Remember to show that this holds for *all* rooms, not just for one room.) For this part, you may again assume that $n \geq 5$.

Hint: You might want to define some new random variables that are independent.

Problem 5. Morris Likes To Count [16 points]

Morris invented the following clever algorithm for (approximate) counting:

```
Algorithm 1: Morris Approximate Counter
c = 0
INCREMENT()
    With probability 1/2^c: set c = c + 1.

ANSWER()
    return 2^c − 1
```

Start with the counter c equal to 0. On each increment, increment the counter c with probability $1/2^c$. Morris showed that after n increment operations:

$$\mathrm{E}\left[2^c\right] = n + 1$$

$$\mathrm{Var}\left[2^c\right] \leq n^2$$

(You can assume this is true for today. As a fun exercise, you can prove it by induction.) This means that, luckily, it is an unbiased estimator, giving the propert expected answer! Unfortunately, the variance is quite high.

Problem 5.a. Design an algorithm using the Morris Counter as a block box that returns a $(1 \pm \epsilon)$ approximation of the correct count with probability at least 3/4. Do this by running α copies of the Morris Counter in parallel and combining the answers in some way. Explain how your algorithm works. (Give your analysis/proof on the next page.)

Problem 5.b. Prove that your algorithm is correct, i.e., after n increment operations it returns a value $n(1 \pm \epsilon)$ with probability at least $3/4$.

Problem 5.c. What value of α did you choose, as a function of ϵ?

Briefly explain your choice of α:

Problem 6. The World Congress[1] [17 points]

In the future, we will all be governed by The World Congress, which has representatives from n countries. The World Congress has many subcommittees, each of which governs an important aspect of daily life (e.g., the Committee on Clean Air, the Committee on the Prevention of War, the Committee on Superhero Management, etc.). Your job is to help design a randomized algorithm for selecting membership of the various committees. There are $k \leq n$ committees in total. Each committee makes decisions based on majority vote.

Unfortunately, some of the representatives represent evil countries that want to overthrow The World Congress and destroy the world.[2] If evil representatives take control of a committee, who knows what harm they may do! It is imperative that we ensure that each committee has a majority of representatives that are good (i.e., *not evil*).

Luckily, we know that at most $1/4$ of the representatives are evil, and at least $3/4$ of the representatives are good. So it shouldn't be too hard to design some good committees, right? Each committee is assigned α members, chosen uniformly and independently at random from the n representatives. (A representative may therefore be on more than one committee, and for simplicity, a representative may be chosen more than once for a single committee.)

Your goal is to choose a value of α so that with probability at least $1-1/n$, all the committees have a good majority.

What value of α do you choose: $\boxed{}$

Prove that all the committees have a good majority with probability at least $1 - 1/n$:

(Extra space on the next page.)

[1]This problem may seem apocryphal, but the basic idea in fact underlies recently popular sharding ideas in cyptocurrencies, where transactions are balanced across smaller committees that can process transactions more quickly.

[2]Their plot is currently unknown, but rumor goes that it involves Thanos, Harvestors, Ice Nine, and burning a lot of oil.

Optional extra space for the previous question:

Scratch Paper

Scratch Paper

Scratch Paper