

Class 2: Agenda and Questions

1 Warm-up

There are n pigeons and n pigeon-holes; each pigeon has its own pigeon-hole. However, after a wild night the n pigeons return to a uniformly random pigeon-hole (so it could be that some holes are empty, and some have more than one pigeon). What's the expected number of empty pigeon-holes?

2 Welcome Back and Announcements!

- First homework is up now!

3 Questions?

Any questions from the lecture notes or short lecture videos?

4 Coupon Collecting

This will be a good chance to practice the power of linearity of expectation.

Let W be a collection of n words. For example, $W = \{\text{hello, randomized, algorithms, penguin, spaceship, \dots, mushroom}\}$. There's a button in front of you. Each time you push the button, it will say a word uniformly at random from W . If you push the button multiple times, it answers independently each time. (In particular, it's possible to get the same word twice). The question is:

What is the expected number of times you push the button before you see all n words?

4.1 Group work

You'll answer this question by answering the following questions in your groups.

Let X_i be the time at which you see the i 'th new word. For example, if you push the button six times and see

penguin, penguin, randomized, hello, randomized, spaceship

then $X_1 = 1$, since you saw your first new word ("penguin") on push 1. $X_2 = 3$, since you saw the second new word ("randomized") on push 3. And $X_3 = 4, X_4 = 6$.

Group Work


1. What is $\mathbb{E}X_1$? (This is not a trick question).
2. What is $\mathbb{E}(X_2 - X_1)$? That is, in expectation, how many times do you press the button, after you have seen the first word, before you see a new, second word?
3. What is $\mathbb{E}(X_3 - X_2)$?
4. For any $i = 2, 3, \dots, n$, what is $\mathbb{E}(X_i - X_{i-1})$?
5. Use your answers to the above, plus linearity of expectation, to answer our question: what is the expected number of times you push the button before you see all n words? It's okay if your answer is a summation, but if you have time try to simplify it to get a big-Theta expression.

5 Karger-Stein Algorithm

Next, we'll take a look at a way to speed up Karger's algorithm.

5.1 Motivation for Karger-Stein

Here is one small run of Karger's algorithm:



Group Work

What is the probability of failure at each point in this run? That is, what is the probability that we choose an edge crossing the minimum cut?

Based on your answers, think of ways to improve Karger's algorithm. (In particular, might there be a smarter way to boost the success probability than just to repeat it a bunch of times?)

Your answer perhaps motivates the following algorithm:

MODIFIED-KARGER:

1. Start with a graph G on n vertices.
2. Run Karger's algorithm (once) until there are m vertices remaining. Call the graph you end up with G' . (Note that G' will have some "mega-vertices" comprised of merged vertices).
3. Repeat Karger's algorithm k times independently on G' until we end up with only two mega-vertices. Return the smallest cut we find.

5.2 Group Work

Consider the Modified-Karger algorithm above. In this group work, you will analyze this algorithm.

Group Work

1. Give a bound on the probability, in terms of n and m and k , that MODIFIED-KARGER is successful. (You may not be able to find the probability exactly, but give a decent lower bound, like we did for the original Karger's algorithm; it's okay if your expression is a bit complicated).

You may want to consult the lecture notes on Karger's algorithm. They are available on the course website: cs265.stanford.edu

Hint: You can break up the failure probability into two parts: the probability that we choose an edge crossing the mincut when reducing G to G' , and the probability that we choose an edge crossing the cut in all of the k runs of Karger of G' .

2. Choose $m = \sqrt{n}$ and $k = n \log n$. Use part 1 to show that the success probability of MODIFIED-KARGER is $\Omega(1/n)$.

Hint: The fact that we're looking for a $\Omega(\cdot)$ answer means that it's okay to ignore pesky constants in your analysis.

Hint: You might want to use the useful fact that $1 - x \leq e^{-x}$ for all x .

3. Show that if you repeat MODIFIED-KARGER, with the parameters above, $\Theta(n)$ times, you can obtain a success probability of 0.99.
4. How does this compare to the original Karger's algorithm? More precisely: We saw in the mini-lecture that repeating Karger's algorithm $O(n^2)$ times leads to a success probability of 0.99. This would involve $O(n^3)$ edge contractions (n edge contractions per run of Karger's algorithm).

You've just shown that "repeating MODIFIED-KARGER $\Theta(n)$ times" also has a success probability of 0.99. How many edge contractions does this involve?

5. If you're done with the above and we haven't finished group work yet, you can think about the following two challenge problems:

Above, we saw how to improve Karger's algorithm by breaking one step down into two steps. What if we were to iterate on this idea? (That is, instead of just running Karger on G' , recursively run MODIFIED-KARGER on G'). How would you pick the parameters here? How few edge contractions can you get, if you want success probability 0.999?

Can you derandomize Karger's algorithm?