

CS 473: Fundamental Algorithms, Fall 2011

Discussion 9

October 18, 2011

9.1 CHOOSING BALLS.

Given n balls and k special balls, consider a random permutation of the n balls. What is the probability that the first ball in this permutation is one of the k special balls? What is the probability that the i th ball in this permutation is a special ball?

9.2 FLIPPING COINS.

You flip a coin n times, where the probability for a head is p . Assume that $p < 1/2n$. Prove that the probability of getting 2 or more heads in these n coin flips is at most $2n^2p^2$.

9.3 HOW MANY COINS TO GET THE FIRST HEAD?

You have a coin that comes up heads with probability p . What is the expected number of times you have to flip the coin till you get a heads?

9.4 FIND k TH SMALLEST NUMBER.

This question asks you to design and analyze a *randomized incremental* algorithm to select the k th smallest element from a given set of n elements (from a universe with a linear order).

In an *incremental* algorithm, the input consists of a sequence of elements x_1, x_2, \dots, x_n . After any prefix x_1, \dots, x_{i-1} has been considered, the algorithm has computed the k th smallest element in x_1, \dots, x_{i-1} (which is undefined if $i \leq k$), or if appropriate, some other invariant from which the k th smallest element could be determined. This invariant is updated as the next element x_i is considered.

Any incremental algorithm can be *randomized* by first randomly permuting the input sequence, with each permutation equally likely.

- (a) Describe an incremental algorithm for computing the k th smallest element.
- (b) How many comparisons does your algorithm perform in the worst case?
- (c) What is the expected number (over all permutations) of comparisons performed by the randomized version of your algorithm? (Hint: When considering x_i , what is the probability that x_i is smaller than the k th smallest so far?) You should aim for a bound of at most $n + O(k \log(n/k))$. Revise (a) if necessary in order to achieve this.

[Notice, that if k is much smaller than n , then your algorithm performs $n + o(n)$ comparisons.]

9.5 SORTING RANDOM NUMBERS

Suppose we pick a real number x_i at random (uniformly) from the unit interval, for $i = 1, \dots, n$.

- (a) Describe an algorithm with an expected linear running time that sorts x_1, \dots, x_n .

To make this question more interesting, assume that we are going to use some standard sorting algorithm instead (say merge sort), which compares the numbers directly. The binary representation of each x_i can be generated as a potentially infinite series of bits that are the outcome of unbiased coin flips. The idea is to generate only as many bits in this sequence as is necessary for resolving comparisons between different numbers as we sort them. Suppose we have only generated some prefixes of the binary representations of the numbers. Now, when comparing two numbers x_i and x_j , if their current partial binary representation can resolve the comparison, then we are done. Otherwise, they have the same partial binary representations (upto the length of the shorter of the two) and we keep generating more bits for each until they first differ.

- (b) Compute a tight upper bound on the expected number of coin flips or random bits needed for a single comparison.
- (c) Generating bits one at a time like this is probably a bad idea in practice. Give a more practical scheme that generates the numbers in advance, using a small number of random bits, given an upper bound n on the input size. Describe a scheme that works correctly with probability $\geq 1 - n^{-c}$, where c is a prespecified constant.

9.6 WHAT HAPPENS ON URANUS STAYS ON URANUS.

There are currently k Uranusians alive on Uranus (they look a lot like cucumbers, but with legs). Each month exactly one of the Uranusians undergoes a critical event. Either it dies with probability p or it splits into two new Uranusians with probability p . With probability $1 - 2p$ nothing happens.

- (A) Let X_i be the size of the population of the Uranus population after the i th month. What is $\mathbf{E}[X_i]$? Here, you need to only provide a good upper bound on $\mathbf{V}[X_i]$, as computing it exactly seems hard.
- (B) Let P_i be the probability that the population of Uranus is non-zero after i months. Prove that $\lim_{i \rightarrow \infty} P_i = 0$. (There is a short and elegant argument showing that, but it is in fact not easy to see.)
- (C) (Harder.) Let X be the number of months till the population of Uranus is extinct. Prove that $\mathbf{E}[X]$ is unbounded.

This is a very simple example of a Galton-Watson process. From wikipedia:

The Galton-Watson process is a branching stochastic process arising from Francis Galton's statistical investigation of the extinction of family names.

There was concern amongst the Victorians that aristocratic surnames were becoming extinct. Galton originally posed the question regarding the probability of such an event