Homework 5

(Network Optimization by Randomization)

1. (50 pts.) Consider the balls-and-bins experiment with 3 balls being thrown (independently and uniformly) into 3 bins. Describe an associated probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and, on this space:

   (a) Construct at least two Bernoulli random variables.

   (b) Construct at least two Binomial random variables.

   (c) How do your random variables relate to each other?

   (d) Compute the probability that the first bin is empty.

   (e) Pick one bin uniformly, and compute the probability that it's empty. If you get an equal/different probability from d), explain why.

   (f) Pick a bin with a 3:2:1 bias for the first/second/third bin and compute the probability that it's empty. If you get an equal/different probability from d,e), explain why.

   (g) Estimate the probability that any of the bins is empty.

2. (100 pts. Bonus) Consider the load-balancing experiment in which m identical jobs are to be assigned to n processors. Consider the following allocation algorithm

   - Initially, all the processors are empty.
   - For each job select k processors (independently, uniformly, and without replacement), and assign the job to the least loaded processor in terms of the number of jobs already assigned.

   Simulate this algorithm for the following values $m = 1000$, $n = 100$, $k = 1, 2, 3, 5, 10, 50, 100$.

   (a) For each triplet (m, n, k), record the *maximum* number of jobs a_k encountered at the processors. Generate a plot with the values of k's on the x-axis and the values of a_k's on the y-axis, and describe what you observe.

   (b) Same as in 1), except that for each triplet (m, n, k) you run 100 simulations, and on the y-axis you display the average values of the a_k's. Do you get same observations as in the plot from 1)? Explain why!

Notes: Homework due on June 28^{th}; points over 100 are considered to be Bonus