CS 286a | **HW 4: End game** | **Guru: Adam Wierman**

Assigned: 3/10/09 | Due: 3/18/09 5pm

I encourage you to discuss these problems with others, but you need to write up the actual solutions alone. At the top of your homework sheet, please list all the people with whom you discussed. Crediting help from other classmates will not take away any credit from you. Also, please limit your use of the web as it is difficult to develop good problems, and so enough search will probably lead you to find the solutions online.

1 Course evaluations [10 points]

I'll be very happy if you're willing to fill out 2 course evaluations for the term: the official online Caltech evaluation and my own informal evaluation form (attached). The official one is clearly important and useful, but my additional questions will help me to adapt some specific parts of the course for the students next year.

2 Conspiracy or catastrophe? [15 points]

We've mentioned the conspiracy principle and the catastrophe principle quite a lot recently, but we haven't studied them very explicitly under specific distributions. To illustrate the principles a little more concretely, let's focus on the case of the Weibull distribution.

Our goal will be to understand the most likely way for a sum of n i.i.d. Weibull's to be bigger than na.

(a) Let's first determine the probability that all of the Weibull's are approximately a. We will approximate this quantity by calculating $f(a)^n$. (Why is this a reasonable approximation?)

(b) Next, we'll determine the probability that 1 is very large (approximately na) and the rest are very small (approximately 0). We will approximate this quantity by calculating $nf(an)$. (Why is this a reasonable approximation?)

(c) Contrast $f(a)^n$ with $nf(an)$ for the Weibull. Interpret the results in terms of the conspiracy principle and the catastrophe principle.

(d) Comment on how accurate the two approximations we used actually are, and what the exact calculation would look like.

3 Large deviations practice [20 points]

In this problem your goal is to derive the tail behavior of waiting time, W, under non-preemptive Last-Come-First-Served (LCFS) in the M/GI/1 heavy-tailed setting. Specifically, your goal is to prove that when the job size distribution is regularly varying with index $\alpha \in (1, 2)$,

$$Pr(W^{LCFS} > x) \sim \rho Pr(\mathcal{E} > (1-\rho)x) \text{ as } x \to \infty.$$

To do this, I suggest applying the Tauberian theorem that we studied in class.

You may start from the following formula for the Laplace transform of the waiting time under LCFS:

$$\widetilde{W}(s) = 1 - \rho + \lambda \frac{1 - \widetilde{B}(s)}{s + \lambda - \lambda \widetilde{B}(s)}$$

and make use of the result we mentioned in class stating that the tail behavior of a busy period in the case of regularly varying job sizes with index α is:

$$Pr(B > x) \sim \frac{1}{1-\rho} Pr(X > (1-\rho)x) \sim \frac{x^{-\alpha} L(x)}{(1-\rho)^{\alpha+1}}$$

4 Not all networks are product-form [15 points]

Give an example of a network which does not have a product-form limiting distribution. Solve your network for the limiting distribution and show that the limiting distribution is in fact not product-form, e.g., prove that the number of jobs at the two servers are not independent.

5 Classed queueing networks [40 points]

In this problem, you will prove that classed queueing networks have a product form (as we mentioned in class).

Specifically, define the *state* of the network to be $z = (z_1, z_2, \ldots, z_k)$, where z_i is the state of server i.

$$\begin{aligned} z &= (z_1, z_2, \ldots, z_k) \\ &= \Big((c_1^{(1)}, c_1^{(2)}, \ldots, c_1^{(n_1)}), (c_2^{(1)}, c_2^{(2)}, \ldots, c_2^{(n_2)}), \ldots, (c_k^{(1)}, c_k^{(2)}, \ldots, c_k^{(n_k)}) \Big) \end{aligned}$$

where n_i denotes the number of packets at server i, and $c_i^{(1)}$ denotes the class of the 1st packet at server i (the one serving) and $c_i^{(n_i)}$ denotes the class of the last packet queued at server i.
Claim: The classed network has product-form, namely,

$$\pi_{(z_1, z_2, \ldots, z_k)} = Pr(\text{state at server 1 is } z_1) \cdot Pr(\text{state at server 2 is } z_2) \cdots Pr(\text{state at server } k \text{ is } z_k)$$

where, server i behaves like an *M/M/1 classed queue*, i.e.

$$Pr(\text{state at server } i \text{ is } z_i = (c_i^{(1)}, c_i^{(2)}, \ldots, c_i^{(n_i)})) = \frac{\lambda_i(c_i^{(1)}) \lambda_i(c_i^{(2)}) \ldots \lambda_i(c_i^{(n_i)})}{{\mu_i}^{n_i}} \cdot (1 - \rho_i)$$

Here are some steps to make it easier:

(a) First consider just a single FCFS server. There are ℓ classes of packets arriving at the server. Let $\lambda(c)$ denote the arrival rate of class c, where $c = 1 \ldots \ell$. The state of the server is just a vector of the form $(c_1, \ldots, c_n)$ where there are n packets at the server, and c_i denotes the class of the ith packet in line.

(i) Prove (by substituting into the balance equations) that the limiting solution has the form

$$\pi_{c_1, \ldots, c_n} = \frac{\lambda(c_1) \cdots \lambda(c_n)}{\mu^n} \cdot C,$$

where C is a normalizing constant.

(ii) Determine the limiting probability that there are n jobs at the server: $\Pr(N_S = n)$. Simplify your answer by using λ to denote $\sum_{c=1}^{\ell} \lambda(c)$. Your answer will still be expressed in terms of the normalizing constant C.

(iii) Use the solution to the previous step to determine the normalizing constant C.

(b) Now consider a network of servers. Use the notation from class. Your goal is to determine the limiting probability of being in state

$$z = \left((c_1^{(1)}, c_1^{(2)}, \ldots, c_1^{(n_1)}), (c_2^{(1)}, c_2^{(2)}, \ldots, c_2^{(n_2)}), \ldots, (c_k^{(1)}, c_k^{(2)}, \ldots, c_k^{(n_k)}) \right).$$

You will use the local balance approach. The equations will look very similar to the un-classed case we did in class. Keep in mind that jobs are allowed to change class.

(i) Show that $A = A'$, namely the rate of leaving state z due to an outside arrival is equal to the rate of entering state z due to a departure to outside.

(ii) Show that $B_i = B_i'$, namely the rate of leaving state z due to a departure from server i is equal to the rate of entering state z due to an arrival at server i.

(c) Verify that the product-form holds. Namely, use summations to determine the limiting probability that the state at server i is z_i and show that

$$\pi_{(z_1, z_2, \ldots, z_k)} = \prod_{i=1}^{k} Pr(\text{state at server } i \text{ is } z_i)$$

COURSE EVALUATION SHEET
(CS/SS 286AB) PERFORMANCE EVALUATION

COMMENTS ON THE COURSE IN GENERAL

1. What topics did you most enjoy across the two terms?

2. What topics did you least enjoy across the two terms?

3. What topics should we have spent more time on?

4. What topics should we have spent less time on?

5. What topics should we have covered that we didn't't?

6. What topics should not be missed in a one term version of the course?

7. Were the lectures paced appropriately?

8. Was attending the lectures useful?

9. Were the assignments interesting and/or helpful?

10. Were the assignments at the right level of difficulty?

11. Were the assignments the right length?

12. How would you rank the course overall

    (bad) 1 2 3 4 5 6 7 8 9 10 (good)

COMMENTS ON ADAM'S LECTURING

Please comment on the strengths and the weaknesses. What should he keep doing and what should he improve upon? Did he improve or get worse in the second term?

-
-
-

COMMENTS ON THE COURE IN GENERAL

Please comment on the strengths and the weaknesses. What should change in future years? What should remain the same?

-
-
-

SUGGESTIONS FOR NEXT YEAR

If the course is offered in a 1 term version where undergraduates and first year graduate students are encouraged to take the course, what topics are most important to cover? What topics should be skipped? Should there be student presentations? Should there be a project? Should the course be less theoretical and consider more applications?

-
-
-