

EE549: Problem Set #3

Due Wednesday Feb. 13

I. OH, THE TROUBLES I'VE SEEN

Let $N(t)$ be a packet arrival process with rate λ packets/second. Suppose all packet sizes $\{B_i\}$ are i.i.d. with mean $\mathbb{E}\{B\}$ bits. Let $X(t)$ be the corresponding bit arrival process. We want to find the rate of $X(t)$ (in bits/second).

Both of the arguments below (for parts (a) and (b)) are completely incorrect and would receive zero credit if I were to grade such work. For each part, write a paragraph explaining why the arguments are incorrect, and whether or not the result is correct (note that it is possible to have a correct result with incorrect reasoning).

a)

$$\frac{X(t)}{t} = \frac{N(t)\mathbb{E}\{B\}}{t}$$

Thus:

$$\lim_{t \rightarrow \infty} \frac{X(t)}{t} = \lim_{t \rightarrow \infty} \frac{N(t)\mathbb{E}\{B\}}{t} = \mathbb{E}\{B\} \lim_{t \rightarrow \infty} \frac{N(t)}{t} = \mathbb{E}\{B\} \lambda$$

b)

$$\lim_{t \rightarrow \infty} \frac{X(t)}{t} = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^t B_i = \mathbb{E}\{B\} \quad \text{by the law of large numbers}$$

II. STATING THE LAW OF LARGE NUMBERS

State the law of large numbers. Your statement should be complete, in the sense that a reader who is unfamiliar with the law of large numbers should be able to understand it. Think about the kind of statement you would like to see in a book.

III. PRACTICE WITH THE LAW OF LARGE NUMBERS

Let $\{H_i\}_{i=1}^{\infty}$ be an i.i.d. sequence of random variables representing *transmission rates*, and suppose that:

$$H_1 = \begin{cases} 5 \text{ kb/s} & \text{with probability } p_1 \\ 10 \text{ kb/s} & \text{with probability } p_2 \end{cases}$$

where p_1 and p_2 are probabilities such that $p_1 + p_2 = 1$. Let $\mu(t)$ be a server transmission rate process that evolves according to the i.i.d. variables $\{H_i\}$ as follows. Let $\alpha > 0$ be a positive constant. Then:

- $\mu(0) = H_1$.
- $\mu(t)$ decreases linearly to zero during the interval $0 \leq t < \alpha H_1$ (so that $\mu(\alpha H_1^-) = 0$).
- $\mu(\alpha H_1) = H_2$.
- $\mu(t)$ decreases linearly to zero during the interval $\alpha H_1 \leq t < \alpha(H_1 + H_2)$.
- Similarly, for any integer $i \in \{2, 3, \dots\}$, $\mu(\alpha(H_1 + \dots + H_i)) = H_{i+1}$, and $\mu(t)$ decreases linearly to zero during $\alpha(H_1 + \dots + H_i) \leq t < \alpha(H_1 + \dots + H_{i+1})$.

We want to compute the time average rate of the $\mu(t)$ process.

a) Draw a picture representing this situation. The picture should be helpful to solve the problem.

b) Define appropriate notation. Specifically, look on your picture for *renewal times*, where the system is “re-freshed” or “renewed” in the way it behaves. You might use t_i to label these renewal times. Also define any other relevant notation, and label your picture with this notation where appropriate.

c) Write a correct equation (without involving any limits) that will be useful to solve the problem. Write a sentence explaining intuitively what the equation says. *Do not write any equation involving quantities or random variables that you have not clearly defined.*

d) Manipulate the equation and take limits at the very end to show that the law of large numbers can be used to solve this problem. *(Note that the law of large numbers is at the very heart of this problem, and hence your answer should show that. If you do not mention the law of large numbers at all, and if you do not set up a situation where it is clear that the law of large numbers is being used, then you are not demonstrating that you have complete understanding of the problem.)*

e) Check that your answer yields $\bar{\mu} = 2.5$ kb/sec if $p_1 = 1$, and yields $\bar{\mu} = 5$ kb/sec if $p_1 = 0$. Explain why this should be true.

f) Now suppose that $p_1 = p_2 = 0.5$, so that when a new H_i variable is chosen, it is equally likely to be either 5 kb/sec or 10 kb/sec. One might temporarily expect that $\bar{\mu} = (2.5 \text{ kb/sec})(0.5) + (5 \text{ kb/sec})(0.5) = 3.75$ kb/sec, but this is not true. Write a paragraph explaining why this is not true. Your paragraph should be thoughtful. It should compare the correct answer to 3.75, and should provide intuition about the comparison.

IV. MORE PRACTICE WITH THE LAW OF LARGE NUMBERS (WITH LESS TRAINING WHEELS)




Fig. 1. A timeline of $\mu(t)$.

Let $\mu(t)$ evolve according to a *delayed renewal process*. Specifically, assume that $\mu(t) = 1$ for $0 \leq t < 1$. Then assume that $\mu(t)$ evolves according to the random trapezoidal sequence as shown in Fig. 1, where the height H_i of the i th trapezoid is a random variable. Assume that the sequence of random variables $\{H_i\}_{i=1}^{\infty}$ is i.i.d. and let $\mathbb{E}\{H\}$ and $\mathbb{E}\{H^2\}$ represent the first and second moments of H_1 . Assume that each trapezoid i starts out with a constant rate of H_i that lasts for exactly one second, and then the rate decreases linearly to 0 with the same slope of $-\gamma$ for some positive constant γ (thus, for example, the first trapezoid ends at time $t = 2 + H_1/\gamma$).

Compute the time average rate of the $\mu(t)$ process. You should of course draw a picture and define appropriate notation. Your answer should be in terms of $\mathbb{E}\{H\}$, $\mathbb{E}\{H^2\}$, and γ (it turns out that $\mathbb{E}\{H^3\}$ or other higher moments do not matter: the answer is the same for any probability density function $p_H(h)$ that has the same first and second moments).

V. THE DYNAMIC ROUTING CONJECTURE

Consider a system of K parallel queues with server rate processes $\mu_1(t), \dots, \mu_K(t)$ with well defined time averages $\bar{\mu}_1, \dots, \bar{\mu}_K$. Let $X(t)$ be an arrival process with rate r composed of fixed length packets of size B kb (so that $X(t) = N(t)B$). Packets must be routed to one of the K queues immediately upon arrival.

In class we saw that the capacity region of this system is given by the set of all rates r such that $0 \leq r \leq \bar{\mu}_1 + \dots + \bar{\mu}_K$. This involves showing that it is impossible to stabilize the system if $r > \bar{\mu}_1 + \dots + \bar{\mu}_K$, and that it is possible to design a policy that makes all queues rate stable if $r \leq \bar{\mu}_1 + \dots + \bar{\mu}_K$. The policies developed in class involved splitting the traffic so that the individual arrival rates r_1, \dots, r_K into each queue (after splitting) are well defined and satisfy $r_i \leq \bar{\mu}_i$ for all queues i . A group of students also proposed the following policy, and conjectured that it would also stabilize the system when $r \leq \bar{\mu}_1 + \dots + \bar{\mu}_K$:

Proposed Routing policy (for the case $K = 2$): If a packet arrives while queue 1 is empty, route the packet to queue 1. Else, route the packet to queue 2.

In this problem, we show that this policy *does not necessarily* stabilize the system, even in cases when a simpler policy *would* stabilize the system. We construct a simple counter-example: Consider the simple case $K = 2$ and assume $\mu_1(t) = 1$ kb/sec and $\mu_2(t) = 3$ kb/sec for all time (so that the service rates are constant in this example). Assume all packets are size $B = 1$ kb. Let $X(t)$ be a periodic arrival process defined as follows: A period has a duration of T seconds and contains 5 packets. The first period starts at time 0, and packets arrive at times $t \in \{0, \epsilon, 2\epsilon, 3\epsilon, 4\epsilon\}$. The next period starts at time T and packets arrive at times $t \in \{T, T+\epsilon, T+2\epsilon, T+3\epsilon, T+4\epsilon\}$, etc. (see Fig. 2).




Fig. 2. An illustration of the periodic arrival processes used in the counter-example.

a) Design parameters ϵ, T such that the arrival process has rate $r < \mu_1 + \mu_2 = 4$ kb/sec, but that the system under the proposed routing policy is *not* rate stable. This involves showing what will happen in the routing, and proving the system is not rate stable.

b) Consider the same 2-queue example system with constant server rates $\mu_1 = 1$ kb/sec and $\mu_2 = 3$ kb/sec, with the same periodic process as your example in part (a). Consider an alternative deterministic “round-robin” type policy for the 2-queue system that routes every fourth packet to queue 1, and all other packets to queue 2. Show that $U_1(t) \leq U_1^{max}$ and $U_2(t) \leq U_2^{max}$, where U_1^{max} and U_2^{max} are finite constants that you should compute.

c) Show that the JSQ policy in the case of constant server rates μ_1, μ_2 stabilizes the system whenever $r \leq \mu_1 + \mu_2$. Show that in the case $\mu_1 = 1, \mu_2 = 3$, and with the deterministic process $X(t)$ from part (a), we have $U_1(t) + U_2(t) \leq U_{tot}^{max}$, where U_{tot}^{max} is a finite constant that you should compute.

VI. THE LEAKY BUCKET CONSTRAINT

Suppose that $X(t)$ is leaky bucket constrained with rate and burst parameters (r, σ) . The first packet of $X(t)$ arrives at time $t = 1$ and has size $B_1 = 10$ kb. The second packet arrives at time $t = 1.5$ and has size $B_2 = 4$ kb. The third packet has size $B_3 = 7.5$ kb.

a) Assume that $r = 8$ kb/sec and $\sigma = 11$ kb. What is the earliest time that the third packet can arrive?

b) Assume that we have the same packet sizes for all three packets, and the same arrival times for the first two packets. Suppose that $\sigma = 11$ kb as before, but the leaky bucket rate parameter changes to $r = 20.5$ kb/sec. What is the earliest time that the third packet can arrive?

VII. A QUEUE WITH NON-ZERO UNFINISHED WORK

Suppose that we have a single-server work conserving queue $U(t)$ with a constant transmission rate of μ bits/sec. Suppose the queue starts out with U_0 amount of bits (so that $U(0) = U_0$). A leaky bucket input $X(t)$ is applied at time 0 (assume that $X(t)$ does not have a packet arrival at time 0, so that U_0 does not include any data from $X(t)$). Suppose $X(t)$ has rate and burst parameters (r, σ) , where $r < \mu$. Give a time t^* such that the system will definitely empty on or before time t^* .

VIII. PERIODICALLY SPLITTING A LEAKY BUCKET STREAM

Suppose that $X(t)$ is a leaky bucket process with rate and burst parameters (r, σ) , where all packets have constant size B .

a) Show that we must have $B \leq \sigma$.

b) Suppose that we split the process $X(t)$ into a new arrival stream by periodically routing every 5th packet from $X(t)$ to the new stream. Let $\tilde{X}(t)$ be the new arrival process. Prove that $\tilde{X}(t)$ is leaky bucket with parameters $(r/5, \tilde{\sigma})$, and compute $\tilde{\sigma}$.




Fig. 3. Two opportunistic systems operating in discrete time.

IX. OPPORTUNISTIC SCHEDULING

Consider the 3-queue wireless system of Fig. 3a. Each queue has an input rate of r kb/sec, and all packets have a fixed size of 1 kb. The three channels are independent ON/OFF processes that are independent of each other and i.i.d. over slots with ON probability p , so that:

$$p = Pr[\text{Channel 1 ON}] = Pr[\text{Channel 2 ON}] = Pr[\text{Channel 3 ON}]$$

For example, for a particular timeslot, the probability that channel 1 is ON and channel 2 is OFF is $p(1 - p)$. On every timeslot the controller observes all three channel states and decides to allocate the server to one of the queues. If the server is allocated to a queue that is ON, it serves one packet (so the rate is 1 kb/slot). Else, it serves zero packets.

a) Suppose we use the (sub-optimal) server allocation strategy that randomly chooses one of the three queues to serve without looking at the channel state, where the choice is made independently and uniformly every slot. What is the maximum rate r that the system under this scheduling policy can stably support on each of the three arrival streams?

b) Now suppose that we can choose our own scheduling policy, possibly one that looks at the channel states before deciding which queue to serve. What is the maximum rate r that the system can stably support on each arrival stream? You must design an algorithm to support this rate r , and then show it is impossible to design an algorithm that supports a larger rate.

c) Consider now the 2-queue system of Fig. 3b. There are two input streams, one with rate r_1 kb/slot, the second with rate r_2 kb/slot. The first input stream enters queue 1, while packets from the second input stream can be routed to either queue 1 or queue 2 (packets must be routed immediately upon arrival). The channel states are independent of each other and i.i.d. over slots, where $p_1 = Pr[\text{Channel 1 ON}]$ and $p_2 = Pr[\text{Channel 2 ON}]$.

Prove that the capacity region Λ , consisting of all possible rate vectors (r_1, r_2) that can be stably supported considering all possible joint routing and scheduling algorithms, is given by the set of all non-negative rate vectors (r_1, r_2) that satisfy:

$$\begin{aligned} r_1 + r_2 &\leq p_1 + (1 - p_1)p_2 \\ r_1 &\leq p_1 \end{aligned}$$

You must show it is impossible to stabilize the system if one of the above constraints is violated. Then you must design a general stabilizing routing and scheduling policy for any (r_1, r_2) that satisfies the above constraints.