**ECE/CS 541: Computer System Analysis**

# Homework #3 (**Due 11/01/18**)

*Student Name: Student's name* *Netid: netid*

The intent of this homework is to get you some hands on experience with Markov chains and serve as introduction to queueing theory. The problems will vary in difficulty from those that you should be able to solve in a couple of minutes, to those that require more thought.

You are to submit this homework in **hard-copy at the start of class** on Thursday November 1, 2018. The number of points on each problem should reflect how much time you should be spending on the problem, this should serve as a good exercise for the midterm.

---

## Problem 1 (15 points): Modeling with Exponentials

A professor recently hired $n$ graduate students and is interested in assigning them to a set of $n$ projects, with one student assigned for one project. When assigning student $i$ to project $j$, the professor will incur a cost $C_{ij}$. For the rest of the problem, we assume that the $C_{ij}$'s are i.i.d. random variables and $C_{ij} \sim Exp(1)$, $\forall i, j, 1 \leq i \leq n, 1 \leq j \leq n$.

**Data:** $S, T, n, C$
**Result:** Assignment of $S$ to $T$ with total cost $C^*$

$C^* \leftarrow 0$
**for** $i \in S$ **do**
    $j^* \leftarrow \underset{j \in T}{\operatorname{argmin}}\, C(i,j)$
    Assign $j^*$ to $i$ with cost $C(i, j^*)$
    $C^* \leftarrow C^* + C(i, j^*)$
    $T \leftarrow T \setminus \{j^*\}$
**end**

(a) Algorithm $A$

**Data:** $S, T, n, C$
**Result:** Assignment of $S$ to $T$ with total cost $C^*$

$C^* \leftarrow 0$
**while** $S \neq \Phi$ **do**
    $(i^*, j^*) \leftarrow \underset{(i,j) \in S \times T}{\operatorname{argmin}}\, C(i,j)$
    Assign $j^*$ to $i^*$ with cost $C(i^*, j^*)$
    $C^* \leftarrow C^* + C(i^*, j^*)$
    $S \leftarrow S \setminus \{i^*\}$
    $T \leftarrow T \setminus \{j^*\}$
**end**

(b) Algorithm $B$

Figure 1: The two greedy algorithms provided by the professor.

The professor suggests two greedy algorithms shown in Figure 1. Algorithm $A$ is a greedy algorithm that attempts to find the local best at every step of the way. Starting with student 1, the algorithm selects the project $j_1$ that will minimize the cost $C(1, j_1) = \min_j C(1, j)$. Next, it will consider student 2 and select the project $j_2$ that will minimize the cost over the remaining projects, i.e., $C(2, j_2) = \min_{j \neq j_1} C(2, j)$, and so on. The algorithm continues until all students have been assigned to projects.

The second algorithm, algorithm $B$, is also a greedy algorithm that takes a more "global" view at each step. It first considers all $n^2$ possible assignments and selects the one with minimum cost, i.e., it assigns student $i^*$ to task $j^*$ such that $C(i^*, j^*) = \min_{(i,j)\in\{1,\ldots n\}\times\{1,\ldots,n\}} C(i,j)$. In the following step, it considers the remaining $(n-1)^2$ assignments and selects the assignment that has the minimum cost, and so on until all assignments have been made.

Seeing that you have taken ECE541, the professor asks you to evaluate these two algorithms and report on the expected cost of the assignments produced by each algorithm. Note that we are not interested in runtime analysis, we are rather interested in the expected cost of the final assignment of each algorithm. Which of the two algorithms is better? Show the set up you used for obtaining your answer.

## Problem 2 (20 points): Modeling with Poisson Processes

I am interested in procuring all $m$ records of my favorite music band. However, given my current graduate student salary, it is either that or rent, for which the choice is (almost?) clear. Luckily, the band is known to give out free records at every one of their concerts. My strategy therefore is to linger around every one of the band's concerts in the area and collect the free records given out. Assume that at every concert, the band gives out record $k$ with probability $p_k$ independently of previous concerts, where $k \in \{1, 2, \ldots, m\}$ and $\sum_{i=1}^{m} p_k = 1$.

Let $N$ be the random variable that denotes the number of concerts that I have to attend until I collect all of the records. Find $E[N]$, i.e., find the expected number of concerts that I have attend in order to collect all of the band's records.

In order to help you obtain the solution, answer the following questions:

1. Let $N_j$ be the random variable representing the number of concerts that I have to attend until I procure record $j$, for $j \in \{1, 2, \ldots, m\}$. What is the distribution of $N_j$?
2. Express $N$ as a function of the $N_j$'s.
3. Why is your expression in question (2) not suitable for finding $E[N]$ (not suitable does not mean impossible, it just means there is an easier way).

Let's now try to overcome this limitation. Since we only care about the number of concerts I need at attend collect all records and not the time it takes me to collect them, we'll assume that the time between any two successive concerts is exponentially distributed with rate 1, independently of previous concerts. Now let $\{N(t), t \geq 0\}$ be the random process representing the number of records collected up to time $t$. Answer the following questions.

1. Why is $\{N(t), t \geq 0\}$ a Poisson process?
2. Let $X_j$ be the random variable representing the time until I first collect record $j$, for $j \in \{1, 2, \ldots, m\}$. Find the distribution of $X_j$.
3. Let $X$ be the random variable representing the time until I collect all records. Express $X$ as a function of the $X_j$'s and obtain $E[X]$. It is okay to express this quantity as an integral. Try to simplify it as much as possible.
4. Express $X$ as a function $N$ and find $E[N]$.

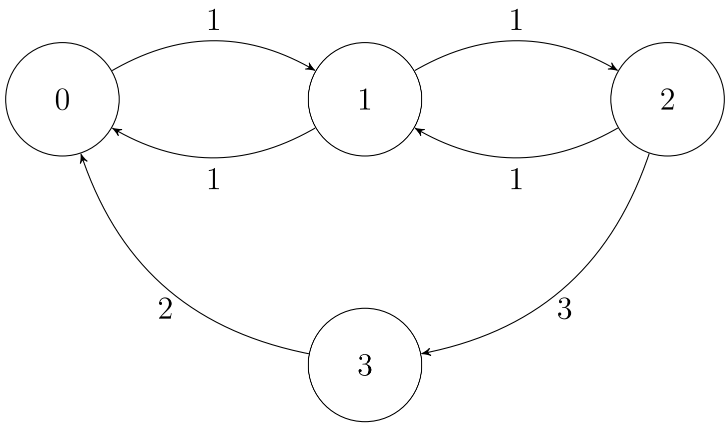


Figure 2: CTMC for problem 4.

## Problem 3 (10 points): CTMC Warmup

Consider a system where customer arrivals form a Poisson process with rate $\lambda$. The system has two servers, each of which gives an exponentially distributed amount of service, with service rate $\mu$. This system is unique in the following ways:

- There is space to hold only two customers in the system.
- The only time an arriving customer is admitted into the system is when both servers are idle.
- The only time when customers enter service is when **all** of the following conditions are satisfied:
  - There is one customer waiting, and
  - a second customer arrives, and
  - both servers are idle.

  In this case, both customers will enter service.

a) Describe this system as a continuous time Markov chain **using at most four states**. Make sure that you state what each state represents and mark the transition rates on the edges.

b) Write the flow balance equations for this system that express the long-term average fraction of time the chain remains in each state.

c) For each state, give the long-term average fraction of time the chain remains in that state.

## Problem 4 (5 points): Uniformization

Consider the CTMC illustrated in Figure 2.

a) What is the transition matrix of the DTMC based on uniformizing this chain with $\lambda = 4$?

b) Use uniformization to compute the state occupancy vector at time 32, given $\pi_0 = (1, 0, 0, 0)$. What is the error bound? You can use any software tools you like.

c) What is the asymptotic state occupancy probability vector for the uniformized DTMC?

d) What do you think (in general) is the relationship between the asymptotic state occupancy probability vector for a CTMC and the asymptotic state occupancy probability vector for the corresponding unifromized DTMC? Why?

## Problem 5 (20 points): The classical car wash problem

A “wash-by-hand” car wash employs four workers. The time required by a worker to wash a car is exponentially distributed with mean 10 minutes. Car arrivals to the car wash form a Poisson process with rate 0.2 per minute. If a car arrives and there is an available worker, the car begins immediately to be washed by that worker. If a car arrives and all available workers are busy, then the car will wait for an available work *unless* there is another car waiting. In such a case, the newly arriving car will leave. Thus at most one car is waiting at a given time.

When a worker finishes washing a car, if there is no car waiting for service, then he slips behind the building for a smoke with probability 0.1. The time it takes a worker to finish his smoke is exponentially distributed with mean 4 minutes. Any worker behind the building smoking is considered not to be available in the event that a car approaches the car wash for service.

Develop a CTMC for this system:

a) What are the states and what do they represent?

b) How would you go about deriving the probability that when a car approaches the car wash, it will eventually be serviced?

c) How would you go about computing the fraction of time that all workers are behind the building smoking?

d) How would you go about deriving the revenue lost by the owner due to cars arriving when an otherwise available worker is behind the building smoking?

## Problem 6 (5 points): Queues Warmup

Given an M/M/1 queue with $\rho = \frac{\lambda}{\mu} < 1$, let $L$ be the long-run average number of jobs in the system.

a) Show that $P[L \geq k] = \rho^k$ for all $k \in \{0, 1, 2, \ldots\}$

b) Using the result in part (a), formulate the CDF of $L$.

c) Using the result in part (a), given an arrival rate $\lambda$, a positive integer $k$, and a probability $p$ such that $0 < p < 1$, find the service rate $\mu$ for which

$$P[L \geq k] = p$$

Specifically, find $\mu$ when $\lambda = 100$, $k = 6$, and $p = 0.2$.

## Problem 7 (10 points): Unhappy students

Due to budget cuts at the state level, the U of I has decided to cut down its number of EWS machines to just 4. Students arrive at the lab according to a Poisson process with average rate of 6 students per hour. Students who find all machines occupied leave impatiently while cursing. A student who finds an empty machine will use it for an average time of 30 minutes, where the use time is exponentially distributed.

1. What is the steady-state fraction of students that will be cursing?
2. What is the average number of machines that are in use?
3. Find the minimum number of machines that the university must buy in order to reduce the fraction of cursing students to a value below $\frac{1}{10}$?

## Problem 8 (15 points): Transactions at the bank

A local Urbana bank utilizes a central server to process online transactions. Incoming transactions form a Poisson process with a rate of 4 transactions per minute. Whenever the server is busy, the transactions are stored in an input buffer of *infinite capacity*. The time to process a single transaction is exponentially distributed with a mean processing time of 10 seconds.

a) What is the average number of transactions in the system? What is the average response time (time spent in the system) for a transaction? What is the probability of finding at least 5 transactions in the system?

b) Consider now the transactions that are stored in the buffer (i.e., the ones that are not immediately serviced). What is the average time that such a transaction spends in the buffer? What is its average response time? How does this average response time (for transactions that must be buffered) compare to the average response time over all transactions?

c) For a given transaction and at a given time, what is the probability that the response time exceeds the average response time?

d) At the time the server was purchased, the bank specified that performance would be satisfactory if the response time for at least 90% of the incoming transactions was no greater than a specified time $t_s$. Assuming that the system had barely met this specification, what is the chosen value of $t_s$?

e) You are now hired as a new IT consultant for the bank. You decide to revise the value of $t_s$ and chose a value of 20 seconds. Assuming that the bank's workload remains the same, must the system be upgraded to meet your new specification? If so, what is the average processing rate required of the new server? What % increase in processing rate does this represent? What is the resulting decrease in server utilization, if any?

## Problem 9 (Bonus Question, 20 points): Edgar Allan Markov

In this problem, you are tasked with creating an algorithm that, given a number $n > 0$, will use Markov chains to generate $n$ lines of song lyrics. For training purposes, you are provided with a file that contains lyrics from your instructor's favorite musicians (Sorry, it's not Beyoncé). In addition, we have provided you with a `Python` script that will strip and tokenize the lyrics.

Explain in details how you set up your solution and any *reasonable* assumptions you make. Submit your code along with your top generated song containing 15 lines.