SOLUTIONS
Resit Stochastic Modelling
February 12, 2024

Question 1. Consider a discrete-time Markov chain on the state space $\{1, 2, 3, 4, 5\}$.

(a) [5pt] Assume the transition matrix is

$$P = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & \frac{2}{3} & \frac{1}{3} \end{pmatrix}$$

For initial state (i) $X_0 = 1$, determine with which probability the Markov chain ends up in each of the absorbing classes. For initial states (ii) $X_0 = 2$, (iii) $X_0 = 5$, determine whether a limit distribution exists and find the limit distribution in case it exists.

Solution Looking at the transition diagram,

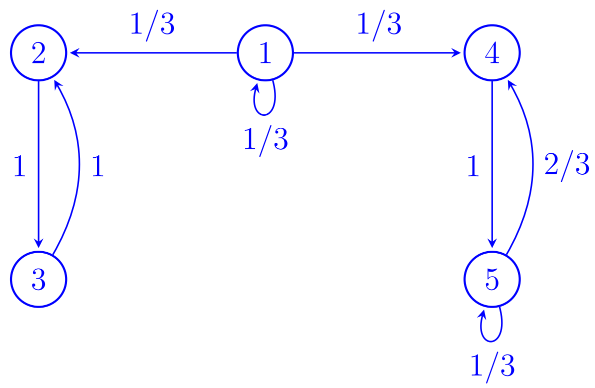

there are three communicating classes:

- {1} is transient,
- {2,3} is absorbing and periodic with period 2,
- {4,5} is absorbing and aperiodic.

(i) If $X_0 = 1$, then the Markov chain ends up in $\{2, 3\}$ with probability

$$P(\text{escape to } 2 \mid \text{escape from } 1) = \frac{1/3}{1/3 + 1/3} = 1/2$$

and, similarly, ends up in $\{4, 5\}$ with probabiltiy 1/2.

(ii) If $X_0 = 2$, then the Markov chain forever stays in the absorbing class {2,3}. It oscillates between states 2 and 3 and and hence π^{lim} does not exist. More formally, the transient distributions repeat the pattern

$$\begin{aligned} &\dots \\ \pi^{(n)} &= (0, 1, 0, 0, 0) \\ \pi^{(n+1)} &= (0, 0, 1, 0, 0) \\ &\dots \end{aligned}$$

and do not converge.

(iii) If $X_0 = 5$, then the Markov chain forever stays in the absorbing class {4,5}. This class is finite and aperiodic, hence the limit distribution within this class exists and is given by the balance and normalization equations

$$\begin{cases} \pi_4 = \pi_5 * 2/3, \\ \pi_5 = \pi_4 + \pi_5 * 1/3, \\ \pi_4 + \pi_5 = 1, \end{cases}$$

which give $\pi_4 = 2/5$, $\pi_5 = 3/5$. Then the limit distribution on the whole state space is

$$\pi^{lim} = (0, 0, 0, 2/5, 3/5).$$

(b) [4pt] Assume the transition matrix is

$$P = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ 0 & 0 & 1 & 0 & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & \frac{2}{3} & \frac{1}{3} \end{pmatrix}$$

Note that the only difference with part (a) is in the transitions out of state 3. What is the expected number of steps it takes to reach state 4 from state 1?

Solution Let $T_4 := \min\{n \geq 0\colon X_n = 4\}$ and $m_i := E(T_4 \mid X_0 = i)$. The question is what is m_1.

By conditioning on the 1st step we get the system

$$\begin{cases} (1) \quad m_1 = 1 + 1/3m_1 + 1/3m_2 + 1/3 * 0, \\ (2) \quad m_2 = 1 + m_3, \\ (3) \quad m_3 = 1 + 1/2m_1 + 1/2m_2. \end{cases}$$

We plug (3) into (2) and get

$$m_2 = 1 + 1 + 1/2m_1 + 1/2m_2 \quad \Leftrightarrow \quad m_2 = 4 + m_1,$$

which we plug into (1) and get

$$m_1 = 1 + 1/3m_1 + 4/3 + 1/3m_1 \quad \Leftrightarrow \quad 1/3m_1 = 7/3 \quad \Leftrightarrow \quad m_1 = 7.$$

Question 2. Every day Bob commutes to work in the morning and then commutes back home in the evening. From time to time he likes to buy a coffee to-go for his commute. As an environmentally conscious person, Bob owns three travel coffee cups and uses them for his to-go coffees when he can.

To be more specific, each commute Bob feels like having a coffee with probability 2/3, independently of his other commutes. If Bob does feel like having a coffee and finds one of the travel cups at the starting location of the commute, he grabs that cup along (so the cup transfers to the end location of the commute). If Bob does feel like having a coffee but finds no travel cup at the starting location, he will buy a coffee in a disposable cup. If Bob does not feel like having a coffee, he does not carry any travel cups along.

(a) [**5pt**] Argue that the following sequence is a discrete-time Markov chain:

$$X_n = \text{number of travel cups at the } \textit{starting location of commute } n.$$

Hint: note that the starting location of the next commute $n+1$ is the end location of commute n. Hence, you need to consider transitions between the starting and end locations of a commute.

Solution The transition diagram is

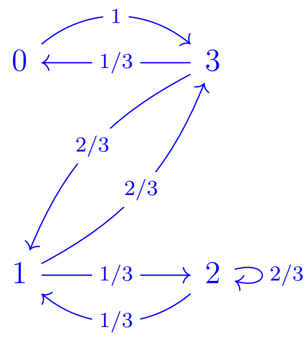

First we explain *explain* how the *transition probabilities* are calculated. Using the hint, the transitions between the starting location of commute n and starting location of commute $n+1$ are the transitions between the starting and end locations of commute n. So we can think of the transitions as follows: if there are i cups at the starting location (say, home) of a commute, then what is the probability that - after this commute - there are j cups at the end location (say, work).

- If there are $i = 2$ cups at home before a commute to work, that means there is 1 cup at work before that commute. Wp 2/3 (1/3) there will be a coffee-wish (no coffee-wish) for that commute and that will transfer 1 extra cup (no extra cup) to work, making it $j = 1 + 1 = 2$ cups at work ($j = 1 + 0 = 1$ cups at work) at the end of the commute.

- Transition probabilities out of states $i = 1$ and $i = 3$ are calculated similarly.

- If there are $i = 0$ cups at home before the commute, there is no cups to transfer to work in any case, regardless of whether there is a coffee-wish or not for that commute. All 3 of the cups are at work before the commute and it will in any case remain $j = 3$ cups at work at the end of the commute.

The *Markov property* follows from the independence of Bob's coffee wishes for the different commutes. The *time-homogeneity* is in place since the transition probabilities in the diagram do not depend on time n. To summarize, X_n, $n \geq 0$, is a time-homogeneous DTMC.

(b) [4pt] What is the long-run fraction of commutes for which Bob finds no travel cup at the starting location?

Solution The question is what is π_0^{occ}.

Since this MC is irreducible and has a finite state space, it has a π^{occ}, which solves the following systems of equations:

$$\begin{cases} \text{balance for state } 0 & \pi_0 = \pi_3 \frac{1}{3}, \quad 1)\ \pi_3 = 3\pi_0 \\ \text{balance for set } \{0,3\} & \pi_3 \frac{2}{3} = \pi_1 \frac{2}{3}, \quad 2)\ \pi_1 = \pi_3 = 3\pi_0 \\ \text{balance for state } 2 & \pi_2 \frac{1}{3} = \pi_1 \frac{1}{3}, \quad 3)\ \pi_1 = \pi_2 = 3\pi_0 \\ \text{norm} & \pi_0 + \pi_1 + \pi_2 + \pi_3 = 1 \quad 4)\ (1+3+3+3)\pi_0 = 1 \Rightarrow \pi_0 = \frac{1}{10}. \end{cases}$$

Hence, the final answer to (b) is $\pi_0^{occ} = \frac{1}{10}$.

Also, as will be relevant for (c), the full occupancy distribution is

$$\pi^{occ} = \left(\frac{1}{10}, \frac{3}{10}, \frac{3}{10}, \frac{3}{10}\right).$$

(c) [**3pt**] Each time Bob uses a travel coffee cup, he gets a discount of €0.25 on his coffee. How much does Bob save on average per commute over the long time-run?

Solution The generic discount in each state is:

$$C_0 = 0, \quad C_1 = C_2 = C_3 = \begin{cases} 0.25 & \text{wp } 2/3, \quad \text{if Bob wants a coffee on that commute} \\ 0 & \text{wp } 1/3. \quad \text{if does not} \end{cases}$$

Hence the long-run average discount (savings) per commute is

$$\pi_0^{occ}\mathbb{E}C_0 + \pi_1^{occ}\mathbb{E}C_1 + \pi_2^{occ}\mathbb{E}C_2 + \pi_3^{occ}\mathbb{E}C_3 = \left(\frac{3}{10} + \frac{3}{10} + \frac{3}{10}\right)\left(\frac{2}{3} \cdot 0.25\right) = €0.15.$$

Question 3. An emergency desk receives true alarms according to a Poisson process with rate 8 per day. It also receives false alarms according to a Poisson process with rate 1 per day. The two Poisson processes are independent.

(a) [**4pt**] What is the probability that the next two alarms are both true alarms?
Hint: what is the distribution of the times between successive true alarms? and between successive false alarms?

Solution Denote by T_1 (F_1) the time until the next true (false) alarm and by T_2 the time between the next and 2nd next true alarms. The question is what is

$$P(T_1 < F_1, T_2 < \text{remaining } F_1).$$

In this probability, T_1 and $T_2 \sim \text{Exp}(8)$, F_1 and remaining $F_1 \sim \text{Exp}(1)$, and there is independence within the pairs of competing exponentials. Hence, the answer is

$$P(T_1 \text{ wins from } F_1, T_2 \text{ wins from remaning } F_1) = \frac{8}{8+1} \cdot \frac{8}{8+1} = \frac{64}{81}.$$

(b) [6pt] A 24-hour day consists of three 8-hour shifts. What is the probability that exactly 3 alarms (in total, true and false) are received on a given day but none of them is received in the middle shift of the day?

Solution Denote by $N(\cdot)$ the total arrival process of true and false alarms together; by merging, it is a Poisson process with rate $\lambda = 9$. The question is what is

$$P\{\overbrace{N(2/3) - N(1/3)}^{\#\text{ alarms in the middle shift}} = 0, \overbrace{\Big(N(1/3) - N(0)\Big) + \Big(N(1) - N(2/3)\Big)}^{\#\text{ alarms in the 1st and 3rd shifts together}} = 3\}?$$

We note that
- the numbers of alarms in the three shifts are independent because the shifts do not overlap,
- the number of alarms in each of the three shifts $\sim \text{Poi}(\lambda * 1/3) = \text{Poi}(3)$;
- due to independence and merging of Poisson random variables, the number of alarms in the 1st and 3rd shift together $\sim \text{Poi}(3+3) = \text{Poi}(6)$.

Hence,the answer is

$$P(\text{Poi}(3) = 0) \cdot P(\text{Poi}(6) = 3) = e^{-3} * e^{-6}\frac{6^3}{3!} = 36e^{-9}.$$

Question 4. Consider a single-server system where *potential* customers arrive according to a Poisson process with rate λ. If, upon arrival, a customer finds $i = 1, 2, \ldots$ other customers in the system, then this customer joins the queue with probability $1/(i+1)$ or leaves immediately with probability $i/(i+1)$. A customer that arrives into an empty system immediately proceeds to the server. The service times are distributed exponentially with rate μ.

(a) **[5pt]** Argue that the number of customers in the system is a continuous-time Markov chain. Intuitively, for which λ and μ is this system stable?

Solution Let $L(t)$ be the number of customers in the system at time t, $t \geq 0$. This is a CTMC since all transitions take Exponential times according to the diagram

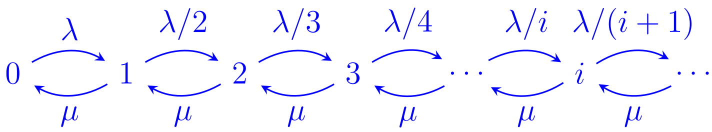

Intuitively, this CTMC is stable for all λ and μ because, in large states, the growth rate λ/i is smaller than the decay rate μ.

(b) **[4pt]** Find the occupancy distribution p^{occ} in the stable scenario.
Hint: the Taylor expansion for the exponential function is $e^x = \sum_{i=0}^{\infty} \frac{x^i}{i!}$.

Solution This is an irreducible MC, and hence the system for p^{occ} is

$$\begin{cases} \text{global balance for sets } \{0, \ldots, i-1\}: \\ p_{i-1} * \lambda/i = p_i * \mu, \quad i = 1, 2, \ldots \\ \text{normalization: } \sum_{i=0}^{\infty} p_i = 1. \end{cases}$$

It follows that,

$$p_i = \frac{\lambda/\mu}{i} p_{i-1} = \frac{\lambda/\mu}{i} \frac{\lambda/\mu}{i-1} p_{i-2} = \frac{\lambda/\mu}{i} \frac{\lambda/\mu}{i-1} \frac{\lambda/\mu}{i-2} p_{i-3}$$
$$= \ldots = \frac{\lambda/\mu}{i} \frac{\lambda/\mu}{i-1} \frac{\lambda/\mu}{i-2} \cdots \cdots \frac{\lambda/\mu}{1} p_0 = \frac{(\lambda/\mu)^i}{i!} p_0, \quad i \geq 1, \quad \text{automatically holds for } i = 0$$

which we plug into the normalization equation and get

$$1 = p_0 \sum_{i=0}^{\infty} \frac{(\lambda/\mu)^i}{i!} = p_0 e^{\lambda/\mu} \quad \Leftrightarrow \quad p_0 = e^{-\lambda/\mu}.$$

Hence,

$$p_i^{occ} = e^{-\lambda/\mu} \frac{(\lambda/\mu)^i}{i!}, \quad i \geq 0.$$

Note that the balance and normalization equations have a solution / p^{occ} exists for any λ and μ. This confirms the intuition on the stability condition from part (a).

(c) [3pt] Express the fraction of lost customers in terms of the probabilities p_i^{occ}. You do not have to further plug in the formulas for p^{occ} that you found in (b).

Solution By PASTA, fraction p_i^{occ} of potential customers upon their arrival find i other customers already in the system. Out of customers who, upon arrival, find i other customers already in the system, fraction $i/(i+1)$ leaves immediately (i.e. they are lost). Hence, in total the fraction of lost customers is

$$\sum_{i=1}^{\infty} p_i^{occ} \frac{i}{i+1}.$$

Question 5. Jobs arrive at a server according to a Poisson process of rate $\lambda = 1/3$ and are served in the order of arrival. For 1/4 of the jobs, their service times have a normal $N(4, 1^2)$ distribution. The remaining 3/4 of the jobs require a fixed service time 2.

(a) [5pt] Find the average waiting time EW.

Solution This is an $M/G/1$ system with arrival rate and service time

$$\lambda = 1/3, \quad B = \begin{cases} N(4, 1^2), & \text{wp } 1/4, \\ 2, & \text{wp } 3/4. \end{cases}$$

Since the order of service is FIFO, the Pollaczek-Khinchine formula applies,

$$\mathbb{E}W = \frac{\rho}{1-\rho} \cdot \frac{\mathbb{E}(B^2)}{2\mathbb{E}(B)}.$$

We have

$$\begin{aligned} \mathbb{E}(B) &= \tfrac{1}{4} * \overbrace{\mathbb{E}(N(4,1^2))}^{4} + \tfrac{3}{4} * 2 = 5/2, \\ \mathbb{E}(B^2) &= \tfrac{1}{4} * \underbrace{\mathbb{E}(N(4,1^2))^2}_{=\mathbb{V}+(\mathbb{E})^2=1^2+4^2=17} + \tfrac{3}{4} * 2^2 = 29/4, \\ \rho &= \lambda \mathbb{E}(B) = 1/3 * 5/2 = 5/6, \end{aligned}$$

and hence

$$\mathbb{E}W = \frac{5/6}{1/6} \cdot \frac{29/4}{2 * 5/2} = 29/4 = 7.25.$$

For part (b), consider a new situation where the server requires *start-up times* of fixed duration 3. That is, when, after an idling period with no jobs to do, the server receives a job, it will only start serving that job 3 time units later.

You can use without proof the fact that the system with such start-up times is still stable, and hence the long-run fraction of time that the server is busy serving jobs is given by ρ. Moreover, denote by Π_{idle} the long-run fraction of time that the server is idling with no jobs to do, and denote by $\Pi_{\text{start-up}}$ the long-run fraction of time that the server is starting-up.

(b) [**6pt**] Find the average waiting time EW in this new situation by doing Mean Value Analysis.

Hint: in the arrival relation, consider what an arriving job has to wait for if it arrives when the server is idling, when the server is starting-up, and when the server is busy serving another job. You can leave Π_{idle} and $\Pi_{\text{start-up}}$ as they are in your solution, you do not have to calculate them.

Solution The MVA equations are:

$$\begin{cases} \text{Little's law} & \mathbb{E}L^q = \lambda \cdot \mathbb{E}W, \\ \text{arrival relation} & \mathbb{E}W = \Pi_{\text{idle}} \cdot 3 + \Pi_{\text{start-up}} \cdot \mathbb{E}R\big|_{X=3} + \rho \cdot \mathbb{E}R\big|_{X=B} + \mathbb{E}L^q \cdot \mathbb{E}B. \end{cases}$$

The arrival relation above comes up as follows:

- Proportion Π_{idle} of jobs (by PASTA) arrive while the server is idling with no jobs to do. Such a job first has to wait for the *full start-up time* 3.
- Proportion $\Pi_{\text{start-up}}$ of jobs (by PASTA) arrive while the server is starting up. Such a job first has to wait for the *remaining start-up time*. Consulting the formula sheet, the average remaining start-up time is $\mathbb{E}R$ where we plug in $X = 3$:

$$\mathbb{E}R\big|_{X=3} = \frac{\mathbb{E}(3^2)}{2\mathbb{E}(3)} = \frac{9}{2 \cdot 3} = \frac{3}{2}.$$

- Proportion ρ of jobs (by PASTA) arrive while the server is busy serving another job. Such a job first has to wait for the *remaining service time* in progress. From (a) we know $\rho = 5/6$. To calculate the average remaining service time, we use the values of $\mathbb{E}(B^2)$, $\mathbb{E}(B)$ that we found in (a):

$$\mathbb{E}R\big|_{X=B} = \frac{\mathbb{E}(B^2)}{2\mathbb{E}(B)} = \frac{29/4}{2 \cdot 5/2} = \frac{29}{20}.$$

- Then each job will have to wait for the full service times of the queue in front of it.

To solve the MVA equations, we plug the Little's law into the arrival relation and get

$$\begin{aligned} \mathbb{E}W &= \Pi_{\text{idle}} \cdot 3 + \Pi_{\text{start-up}} \cdot \frac{3}{2} + \frac{5}{6} \cdot \frac{29}{20} + \frac{5}{6} \cdot \mathbb{E}W, \\ \Rightarrow \quad \frac{1}{6} \cdot \mathbb{E}W &= \Pi_{\text{idle}} \cdot 3 + \Pi_{\text{start-up}} \cdot \frac{3}{2} + \frac{5}{6} \cdot \frac{29}{20}, \\ \Rightarrow \quad \mathbb{E}W &= 18\Pi_{\text{idle}} + 9\Pi_{\text{start-up}} + 29/4. \end{aligned}$$

FORMULA SHEET

Erlang distribution. If S_n has an Erlang(n, μ) distribution, then

$$P(S_n > t) = \sum_{k=0}^{n-1} e^{-\mu t} \frac{(\mu t)^k}{k!} \qquad \text{and} \qquad f_{S_n}(t) = \mu\, e^{-\mu t} \frac{(\mu t)^{n-1}}{(n-1)!}.$$

Residual time till next event. Let X be a generic inter-event time and R the residual time till next event. Then

$$P(R \leq x) = \frac{1}{E(X)} \int_0^x P(X > u)\, \mathrm{d}u \qquad \text{and} \qquad E(R) = \frac{E(X^2)}{2E(X)}.$$

M/G/1 queue. The waiting time W under FIFO and the busy period BP under work-conserving disciplines satisfy

$$E(W) = \frac{\rho}{1-\rho} \frac{E(B^2)}{2E(B)} = \frac{1}{2} \frac{\rho}{1-\rho}(1 + c_B^2)E(B), \quad \text{where } c_B^2 = \frac{V(B)}{(E(B))^2}$$

$$E(BP) = \frac{E(B)}{1-\rho}.$$

M/M/c queue. The probability of waiting Π_W, waiting time W and sojourn time S satisfy

$$\Pi_W = \frac{(c\rho)^c/c!}{(1-\rho)\sum_{i=0}^{c-1}(c\rho)^i/i! + (c\rho)^c/c!},$$

$$E(W) = \Pi_W \frac{1}{c\mu(1-\rho)} \qquad \text{and} \qquad P(W > t) = \Pi_W\, e^{-c\mu(1-\rho)t},$$

$$P(S > t) = \frac{\Pi_W}{1 - c(1-\rho)} e^{-c\mu(1-\rho)t} + \left(1 - \frac{\Pi_W}{1 - c(1-\rho)}\right) e^{-\mu t}.$$

M/G/c/c queue. The blocking probability is

$$B(c, a) = \frac{a^c/c!}{\sum_{i=0}^{c} a^i/i!} \qquad \text{with } a = \lambda\, E(B) = c\rho.$$