SOLUTIONS
Resit Stochastic Modelling
February 14, 2022

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,

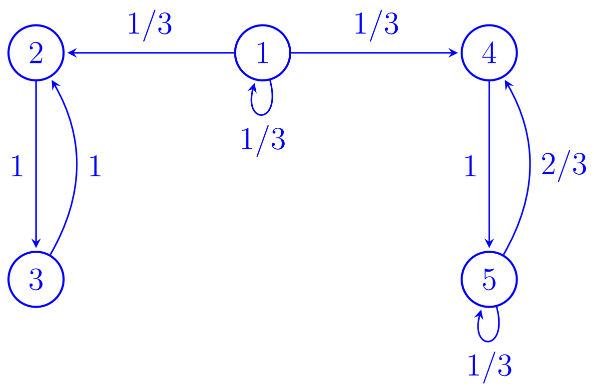

there are three communicating classes:

- {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 limit distribution on the whole state space is

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

(b) [3pt] 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 in (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. A furniture store carries in its range a certain type of bed for which the monthly demands are independent and distributed uniformly between 0 and 4 (i.e., the demand during each month is i with probability 1/5 for $i = 0, 1, \dots, 4$). Due to limited storage, the store can have up to 4 beds in stock at a time. The demand that arrives while beds are out of stock is lost. The inventory control is as follows. If there are less than 2 beds in stock at the end of a month, then the shop restocks up to full capacity and starts the next month with 4 beds in stock (the delivery time from the supplier to the shop is negligible). Otherwise the shop does not restock and starts the next month with what is left from the previous month.

(a) [5pt] Argue that the inventory levels at the end of each month form a discrete-time Markov chain. Provide the transition matrix and a system of equations that determines the occupancy and limit distribution. Without solving this system, argue why the occupancy and limit distributions exist.

Solution Denote by X_n the number of beds in stock at the end of month n, $n \geq 0$. This is a DTMC with state space

$$S = \{0, 1, 2, 3, 4\}$$

and transition matrix

$$P = \begin{array}{c} 0 \\ 1 \\ 2 \\ 3 \\ 4 \end{array} \begin{pmatrix} 1/5 & 1/5 & 1/5 & 1/5 & 1/5 \\ 1/5 & 1/5 & 1/5 & 1/5 & 1/5 \\ 3/5 & 1/5 & 1/5 & 0 & 0 \\ 2/5 & 1/5 & 1/5 & 1/5 & 0 \\ 1/5 & 1/5 & 1/5 & 1/5 & 1/5 \end{pmatrix} \begin{array}{l} \text{repeats row for 4} \\ \text{repeats row for 4} \\ \\ \\ \\ \end{array}$$

To explain some of the transition probabilities,

$$P(X_{n+1} = 0 | X_n = 0) = P(\text{demand in month } n+1 \text{ is } 4) = 1/5,$$
$$\text{stock up to 4}$$
$$P(X_{n+1} = 0 | X_n = 2) = P(\text{demand in month } n+1 \text{ is 2, 3, or 4}) = 3/5.$$
$$\text{remains 2 start of next month}$$

The *Markov property* follows from the independence of the monthly demands. We also have the *time-homogeneity* as the transition probabilities do not depend on n.

Regarding the occupancy and limit distribution, this MC is irreducible, aperiodic and has a finite state space. Hence both π^{lim} and π^{occ} exist and are determined by this system:

$$\begin{cases} \text{balance} & \pi_0 = \frac{1}{5}\pi_0 + \frac{1}{5}\pi_1 + \frac{3}{5}\pi_2 + \frac{2}{5}\pi_3 + \frac{1}{5}\pi_4, \\ & \pi_1 = \frac{1}{5}\pi_0 + \frac{1}{5}\pi_1 + \frac{3}{5}\pi_2 + \frac{2}{5}\pi_3 + \frac{1}{5}\pi_4, \\ & \pi_2 = \frac{1}{5}\pi_0 + \frac{1}{5}\pi_1 + \frac{1}{5}\pi_2 + \frac{1}{5}\pi_3 + \frac{1}{5}\pi_4, \\ & \pi_3 = \frac{1}{5}\pi_0 + \frac{1}{5}\pi_1 \qquad\quad + \frac{1}{5}\pi_3 + \frac{1}{5}\pi_4, \\ & \pi_4 = \frac{1}{5}\pi_0 + \frac{1}{5}\pi_1 \qquad\qquad\qquad + \frac{1}{5}\pi_4, \\ \text{norm} & \sum_{i=0}^{4} \pi_i = 1. \end{cases}$$

(b) [3pt] If the distribution from part (b) were known, how would you calculate the long-run average number of lost sales per month?

Solution

$$\pi_2^{occ}\Big(P(\text{demand} = 3) * 1 + P(\text{demand} = 4) * 2\Big) + \pi_3^{occ} P(\text{demand} = 4) * 1 =$$
$$= \frac{3}{5}\pi_2^{occ} + \frac{1}{5}\pi_3^{occ}$$

(c) [5pt] Consider a new situation where customers that arrive while beds are out of stock are not lost. Instead backorders are placed for these customers. *At the end of a month with backorders*, the store orders (and immediately receives) 4 beds from the supplier. Out of these 4 beds, the backorders

are immediately covered. The remaining beds are the stock that the store starts the next month with. *For months without backorders*, the replenishment rule is the same as before. Model this new situation as a discrete-time Markov chain assuming additionally that, when the shop opened for the very first time, there was a full stock of 4 beds.
Hint: the final assumption guarantees that the number of backorders never exceeds a certain threshold.

Solution The final assumption guarantees that there can be at most 2 backorders per month. Let

$$X_n = \begin{cases} i, & \text{if there are } i = 0,1,2,3,4 \text{ beds in stock} \\ -i, & \text{if there are } i = 1,2 \text{ backorders} \end{cases} \quad \text{at the end of month } n.$$

This is a DTMC with state space

$$S = \{-2, -1, 0, 1, 2, 3, 4\}$$

and transition matrix

$$P = \begin{matrix} -2 \\ -1 \\ 0 \\ 1 \\ 2 \\ 3 \\ 4 \end{matrix} \begin{pmatrix} 1/5 & 1/5 & 1/5 & 1/5 & 1/5 & 0 & 0 \\ 0 & 1/5 & 1/5 & 1/5 & 1/5 & 1/5 & 0 \\ 0 & 0 & 1/5 & 1/5 & 1/5 & 1/5 & 1/5 \\ 0 & 0 & 1/5 & 1/5 & 1/5 & 1/5 & 1/5 \\ 1/5 & 1/5 & 1/5 & 1/5 & 1/5 & 0 & 0 \\ 0 & 1/5 & 1/5 & 1/5 & 1/5 & 1/5 & 0 \\ 0 & 0 & 1/5 & 1/5 & 1/5 & 1/5 & 1/5 \end{pmatrix} \begin{matrix} \text{repeats row for 2} \\ \text{repeats row for 3} \\ \text{repeats row for 4} \\ \text{repeats row for 4} \\ \\ \\ \\ \end{matrix}$$

Like in (a), the Markov property follows from the independence of the monthly demands and we have the time-homogeneity.

Question 3. In a three-component computer system, the components A, B, and C work in parallel and experience failures. A failed component is replaced immediately. All lifetimes are independent and distributed exponentially. The rates are $\lambda_A = 2$, $\lambda_B = 1$, $\lambda_C = 1$ per day for components of type A, B, C, respectively.

(a) [5pt] What is the joint probability that the following happens on a given day: there are two failures, both of them happen in the second half of the day, and none of them is of component A?

Solution From the problem description we recognise that the failures of the three components follow three independent Poisson processes. We denote
- by $N_A(t)$, $t \geq 0$ the Poisson failure process for component A, its rate is $\lambda_A = 2$;
- by $N_{B+C}(t)$, $t \geq 0$, the merger of the Poisson failure processes for components B and C, it is a Poisson process as well, its rate is $\lambda_B + \lambda_C = 2$, and it is independent from $N_A(t)$, $t \geq 0$.

The question is what is

$$P\{\overbrace{N_A(1) - N_A(0)}^{\#\ A \text{ failures in } (0,1]} = 0, \overbrace{N_{B+C}(1/2) - N_{B+C}(0)}^{\#\ B\&C \text{ failures in } (0,1/2]} = 0, \overbrace{N_{B+C}(1) - N_{B+C}(1/2)}^{\#\ B\&C \text{ failures in } (1/2,1]} = 2\}?$$

We note that
- the three events are independent because the processes $N_A(\cdot)$ and $N_{B+C}(\cdot)$ are independent and because the intervals (0,1/2] and (1/2,1] are non-overlapping;
- $N_A(1) - N_A(0) \sim Poi(\lambda_A * (1-0)) = Poi(2)$,
- $N_{B+C}(1/2) - N_{B+C}(0) \sim Poi((\lambda_B + \lambda_C) * (1/2 - 0)) = Poi(1)$,
- $N_{B+C}(1) - N_{B+C}(1/2) \sim Poi((\lambda_B + \lambda_C) * (1 - 1/2)) = Poi(1)$.

Hence,the answer is

$$e^{-2} * e^{-1} * e^{-1}/2! = e^{-4}/2.$$

(b) [3pt] What is the probability that the first failure in this system is of component A?

Solution Here and in part (c), we will use notations $A_1, A_2, \ldots,$ for the consecutive lifetimes (times between failures) of component A, and similarly for B and C.

The question is what is

$$P\{\underbrace{A_1}_{\sim \text{Exp}(\lambda_A)} \text{ wins from } \underbrace{\min(B_1, C_1)}_{\sim \text{Exp}(\lambda_B + \lambda_C)}\} = \frac{2}{2+2} = \frac{1}{2}.$$

(c) [3pt] What is the joint probability that the first, second and third failure in this system are of component A, B and C, respectively?

Solution The question is what is

$$P\{A_1 \text{ wins from } \min(B_1, C_1), \\ \quad \text{remaining } B_1 \text{ wins from } \min(A_2, \text{remaining } C_1), \\ \quad \text{remaining } C_1 \text{ wins from } \min(\text{remaining } A_2, B_2)\}.$$

Using the memoryless property of the Exponential distribution, the probability in question equals

$$\begin{aligned} & P(\text{Exp}(\lambda_A) \text{ wins from Exp}(\lambda_B + \lambda_C)) \\ & * P(\text{Exp}(\lambda_B) \text{ wins from Exp}(\lambda_A + \lambda_C)) \\ & * P(\text{Exp}(\lambda_C) \text{ wins from Exp}(\lambda_A + \lambda_B)) \\ = & \frac{2}{2+2} * \frac{1}{1+3} * \frac{1}{1+3} = \frac{1}{32}. \end{aligned}$$

Question 4. Consider a stable $M/M/2$ system with arrival rate λ, service rate μ, and the load per server $\rho := \lambda/(2\mu) < 1$.

(a) [6pt] Argue that the number of customers in the system is a continuous-time Markov chain. Determine the occupancy and limit distribution. In particular, derive that

$$p_0^{occ} = p_0^{lim} = \frac{1-\rho}{1+\rho}.$$

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

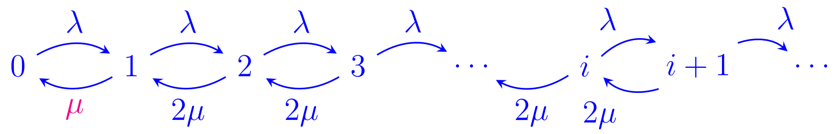

Below we find a solution to balance and normalization equations. This solution is both p^{lim} and p^{occ} since $L(t), t \geq 0$ is an irreducible CTMC.

The system for p^{lim} and p^{occ} is,

$$\begin{cases} \text{global balance for sets } \{0, \dots, i-1\}: \\ p_0 * \lambda = p_1 * \mu, \quad i = 1, \\ p_{i-1} * \lambda = p_i * 2\mu, \quad i = 2, 3, \dots \\ \text{normalization: } \sum_{i=0}^{\infty} p_i = 1. \end{cases}$$

With $\rho := \lambda/(2\mu)$, it follows that

$$p_1 = 2\rho p_1$$

and, for $i = 2, 3, \ldots$,

$$p_i = \rho p_{i-1} = \rho^2 p_{i-2} = \ldots = \rho^{i-1} p_1 = 2\rho^i p_0.$$

To summarize, for all $i \geq 1$, we have $p_i = 2\rho^i p_0$, which we plug into the normalization equation and get

$$1 = p_0 + \sum_{i=1}^{\infty} p_i = p_0 + \sum_{i=1}^{\infty} 2\rho^i p_0 = p_0\Big(1 + 2\sum_{i=1}^{\infty} \rho^i\Big) = p_0\Big(1 + 2(\sum_{i=0}^{\infty} \rho^i - 1)\Big)$$
$$= p_0\Big(2\frac{1}{1-\rho} - 1\Big) = p_0 \frac{2-(1-\rho)}{1-\rho} = p_0 \frac{1+\rho}{1-\rho}.$$

Hence,

$$p_0^{occ} = p_0^{lim} = \frac{1-\rho}{1+\rho},$$
$$p_i^{occ} = p_i^{lim} = 2\frac{1-\rho}{1+\rho}\rho^i, \quad i \geq 1.$$

(b) [2pt] Express the fraction Π_W of customers that experience waiting in terms of the probabilities p_i^{occ}. You do not have to plug in the solution from (a) and further work out the formula.

Solution Customers that do not have to wait (the fraction of such customers is $1 - \Pi_W$) are those who, upon arrival, see at least one of the two servers free, i.e. those who see 0 or 1 other customers in the system. By PASTA,

$$1 - \Pi_W = p_0^{occ} + p_1^{occ}, \quad \text{i.e. } \Pi_W = 1 - p_0^{occ} - p_1^{occ}.$$

(c) [5pt] Use Mean Value Analysis to find the customer-average waiting time EW and the time-average number of customers in the queue EL^q.
Remark: a reference to the formula sheet for EW will not suffice. The waiting probability Π_W can be left in the answer as is.

Solution The MVA equations are:

$$\begin{cases} \text{Little's law} & EL^q = \lambda EW, \\ \text{arrival relation} & EW = \Pi_W * E\text{Exp}(2\mu) + EL^q * E\text{Exp}(2\mu) \\ & \quad\;\; = \Pi_W * \dfrac{1}{2\mu} + EL^q * \dfrac{1}{2\mu}. \end{cases}$$

The arrival relation above comes up as follows:
- Proportion Π_W of customers arrive while both servers are busy and one of the two servers has to free up *before the queue starts moving forward.* This takes an Exp(2μ) amount of time (the minimum of the two remaining Exp(μ) service times, we also use the memorylessness here).
- *Then the entire queue has to move forward.* For each customer in the queue this again takes an $Exp(2\mu)$ amount of time (as one of the two servers has to become free for the next customer in the queue to move forward to the servers).

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

$$EW = \frac{\Pi_W}{2\mu} + EL^q * \rho \quad \Leftrightarrow \quad EW = \frac{\Pi_W}{2\mu(1-\rho)}$$

and

$$EL^q = \frac{\Pi_W \rho}{1-\rho}.$$

Question 5. A service desk handles two types of customers that arrive according to two independent Poisson processes, both of unit rate, $\lambda_1 = \lambda_2 = 1$. The service times are distributed exponentially, with rates $\mu_1 = 3$ and $\mu_2 = 6$ for the two customer types, respectively. The customers are helped one at a time in the order of arrival. The waiting room is unlimited.

(a) [4pt] The number of customers present can be viewed as an $M/G/1$ model. Specify the arrival rate and the service time distribution in this $M/G/1$ model, and find the customer-average waiting time.
Reminder: for a random variable $X \sim$ Exponential(α), the variance is $1/\alpha^2$.

Solution Each next customer is of type 1 if the next type 1 arrival wins from the next type 2 arrival, i.e., with probability $\lambda_1/(\lambda_1 + \lambda_2) = 1/2$. Similarly, each next customer is of type 2 with probability 1/2. Hence, the service time distribution in the $M/G/1$ model representing the total number of customers present is (hyperexponential)

$$B = \begin{cases} \text{Exp}(3) & \text{wp } 1/2, \\ \text{Exp}(6) & \text{wp } 1/2. \end{cases}$$

The arrival rate in this $M/G/1$ model is the rate of the merger of the two arrival processes,

$$\lambda = \lambda_1 + \lambda_2 = 2.$$

We use the PK formula to find the customer-average waiting time,

$$EW = \frac{\rho}{1-\rho}\frac{E(B)^2}{2EB},$$

where

$$\begin{aligned} EB &= \tfrac{1}{2}E\text{Exp}(3) + \tfrac{1}{2}E\text{Exp}(6) = \tfrac{1}{2} * \tfrac{1}{3} + \tfrac{1}{2} * \tfrac{1}{6} = \tfrac{1}{4}, \\ E(B^2) &= \tfrac{1}{2}E(\text{Exp}(3)^2) + \tfrac{1}{2}E(\text{Exp}(6)^2) = \tfrac{1}{2} * \tfrac{2}{3^2} + \tfrac{1}{2} * \tfrac{2}{6^2} = \tfrac{5}{36}, \\ &\text{(we use } E(\text{Exp}(\alpha))^2 = V + (E)^2 = \tfrac{2}{\alpha^2}) \\ \rho &= \lambda EB = 2 * \tfrac{1}{4} = \tfrac{1}{2}. \end{aligned}$$

Altogether, we have

$$EW = \frac{1/2}{1/2} * \frac{5/36}{1/2} = \frac{5}{36} * \frac{2}{1} = \frac{5}{18}.$$

(b) [5pt] Model the situation at the service desk as a continuous-time Markov chain where the state has two components one of which is the number of customers present.

Solution Let

$$X(t) = (\#\text{ customers present, type of customer in service}) \quad \text{at time } t.$$

This is a CTMC with transition diagram

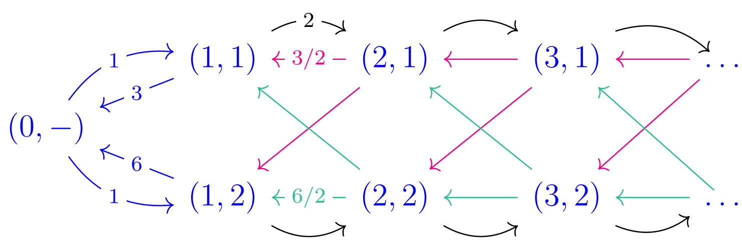

(with all $\xrightarrow{2}$, $\xrightarrow{3/2}$, $\xrightarrow{6/2}$).

To explain: in states $(i, 1)$, once the Exp(3) service is completed, the next customer in queue/to go into service is of type 1 or 2 with probability 1/2. This is equivalent to two separate transitions Exp$(3 * 1/2)$ to $(i-1, 1)$ and Exp$(3 * 1/2)$ to $(i-1, 2)$ since Exp$(3) \sim \min(\text{Exp}(3 * 1/2), \text{Exp}(3 * 1/2))$.

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

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