Introduction to Queueing Networks

Solutions to Problem Sheet 1

---

1. (a) The number of balls in each urn at the next time step depends only on the numbers in each in the current time step, and not on the previous history. The states are $\{0, 1, \dots, n\}$ and the transition probabilities are given by

$$p_{i,i-1} = \frac{i}{n},\ p_{i,i+1} = \frac{n-i}{n}, \quad i = 0, 1, \dots, n. \tag{1}$$

(b) All states are recurrent and belong to a single communicating class, i.e., the Markov chain is irreducible.

(c) As the Markov chain is irreducible and has finitely many states, there is a unique invariant distribution π, which satisfies the global balance equations:

$$\pi_i = \pi_{i+1} p_{i+1,i} + \pi_{i-1} p_{i-1,i}, \quad , i = 1, 2, \dots, n, \tag{2}$$

(where we take $\pi_{-1} = 0$ and $\pi_{n+1} = 0$), as well as the normalisation condition

$$\sum_{i=0}^{n} \pi_i = 1. \tag{3}$$

It is not easy to solve these equations, but if you can guess a solution, it is easy to verify it. The intuition is that each ball is doing a random walk between the two urns, (almost) independent of other balls; the only dependence is that, if one ball is picked at a time step, no other ball is picked in the same step. Moreover, balls are chosen at random, not more or less likely based on which urn they are in, or their own or any other ball's past history. This intuition says that each ball is equally likely to be in either urn, and the positions of different balls are independent. This suggests that the number of balls in either urn should be a Binomial($n, 1/2$) random variable, and hence that

$$\pi_i = \binom{n}{i} \frac{1}{2^n}.$$

It is easy to verify that these π_i satisfy (2) as well as the normalisation condition.

2. (a) N_t is a Markov chain because the number of alleles of each type in generation t suffices to determine the probability distribution for the number of alleles of each type in generation $t+1$; no further knowledge of the past is required.

(b) The state space is $\{0, 1, 2, \dots, N\}$. The states 0 and N are absorbing because, if there are no alleles of one type in some generation, then there can be no alleles of that type in any future generation. All other states belong to a single communicating class since it is possible to go from any number j of alleles of type A to any other number k of such alleles, so long as j isn't 0 or N. Note that k can be 0 or N.

The states 0 and N are obviously recurrent since, if you are ever in one of them, you re-visit them infinitely many times - in fact, you never leave them. All other states are transient. Take state 1 for example. You can't visit it infinitely many times because, on each visit, you have a non-zero chance of hitting state 0 or N in the next step and becoming absorbed. So, after some finite number of visits, this is bound to happen.

(c) The invariant distribution has to be zero outside the set of recurrent states. So, if π is invariant, then $\pi_i = 0$ for all $i \in \{1, 2, \dots, N-1\}$. Moreover, $p_{00} = 1$ and $p_{NN} = 1$, so it is clear that if we take $\pi_0 = \alpha$ and $\pi_N = 1 - \alpha$ for any $\alpha \in [0, 1]$, then π is invariant. These are all the invariant distributions.

3. (a) The states are S, C and R (for sunny, cloudy and rainy), and the transition probabilities are specified by the matrix

$$P = \begin{pmatrix} 0.5 & 0.5 & 0 \\ 0.4 & 0.4 & 0.2 \\ 0 & 0.5 & 0.5 \end{pmatrix}$$

with the states in that order (the first row and column refer to the state S, and so on).

(b) All states form a single communicating class, since there is non-zero probability of going from any state to any other eventually (though not necessarily in one-step: it takes two steps to go from S to R or R to S). All states are also recurrent. Indeed, since there are only finitely many states, not all of them can be transient. (It is not possible that each of the states is only visited finitely many times.) But states in the same communicating class have to all be transient or all be recurrent. As all states of this Markov chain form a single communicating class, they must all be recurrent.

Since there is a single communicating class, the Markov chain is irreducible, and so the invariant distribution is unique. By solving the global balance equations $\pi P = \pi$, together with the normalisation condition $\pi_S + \pi_C + \pi_R = 1$, we find that the invariant distribution is given by $\pi = \left(\frac{4}{11}\ \frac{5}{11}\ \frac{2}{11}\right)$.

(c) If Alice carried an umbrella with her yesterday, then yesterday was either cloudy or rainy. For her to carry an umbrella today, today must be rainy or cloudy. We'll denote the four possibilities for yesterday's and today's joint weather by CC, CR, RC and RR (with the first letter denoting yesterday's weather) and the two possibilities for yesterday's weather by C and R. We want to compute $P(CC \cup CR \cup RC \cup RR | C \cup R)$. Using the invariant distribution calculated in the last part and Bayes' theorem, we have

$$\begin{aligned} P(CC \cup CR \cup RC \cup RR | C \cup R) &= \frac{P(CC \cup CR \cup RC \cup RR)}{P(C \cup R)} \\ &= \frac{\pi_C p_{CC} + \pi_C p_{CR} + \pi_R p_{RC} + \pi_R p_{RR}}{\pi_C + \pi_R} \\ &= \frac{2/11 + 1/11 + 1/11 + 1/11}{5/11 + 2/11} = \frac{5}{7}. \end{aligned}$$

Thus, the probability that Alice carries an umbrella today given that she carried one yesterday is 5/7.

Similarly, if Alice carried an umbrella the last two days, then the weather on these days must have been CC, CR, RC or RR. Hence, the probability that Alice carries an umbrella today given that she did so on the last two days is given by

$$\begin{aligned} &P(CCC \cup CCR \cup CRC \cup CRR \cup RCC \cup RCR \cup RRC \cup RRR | CC \cup CR \cup RC \cup RR) \\ &= \frac{P(CCC \cup CCR \cup CRC \cup CRR \cup RCC \cup RCR \cup RRC \cup RRR)}{P(CC \cup CR \cup RC \cup RR)} \end{aligned}$$

The numerator of the above expression can be evaluated as

$$
\begin{aligned}
&\pi_C\Big(p_{CC}^2 + p_{CC}p_{CR} + p_{CR}p_{RC} + p_{CR}p_{RR}\Big) + \pi_R\Big(p_{RC}p_{CC} + p_{RC}p_{CR} + p_{RR}p_{RC} + p_{RR}^2\Big) \\
&= \frac{5}{11}(0.16 + 0.08 + 0.1 + 0.1) + \frac{2}{11}(0.2 + 0.1 + 0.25 + 0.25) \\
&= \frac{19}{55}.
\end{aligned}
$$

The denominator can be evaluated as

$$
\pi_C p_{CC} + \pi_C p_{CR} + \pi_R p_{RC} + \pi_R p_{RR} = \frac{2}{11} + \frac{1}{11} + \frac{1}{11} + \frac{1}{11} = \frac{5}{11}.
$$

Hence, the probability that Alice carries an umbrella today, given that she did so on the last two days, is given by $\frac{19}{55}/\frac{5}{11} = 19/25$.

(d) From the answer to the previous part,

$$
P(Y_t = 1|Y_{t-1} = 1) = \frac{5}{7}, \text{ whereas } P(Y_t = 1|Y_{t-1} = 1, Y_{t-2} = 1) = \frac{19}{25}.
$$

In other words, the probability distribution of Y_t conditioned on the infinite past (or even two time periods in the past) is not the same as its probability distribution conditioned only on the last time period. Hence, $(Y_t, t \geq 0)$ cannot be a Markov chain.

4. The generating function of X_1 is given by

$$
G_1(z) = E\Big[z^{X_1}\Big] = \sum_{n=0}^{\infty} P(X_1 = n) z^n.
$$

But X_1 is Poisson with parameter λ_1, so $P(X_1 = n) = \lambda_1^n e^{-\lambda_1}/n!$. Substituting this above, we get

$$
G_1(z) = \sum_{n=0}^{\infty} \frac{(\lambda_1 z)^n}{n!} e^{-\lambda_1} = e^{\lambda_1 z} e^{-\lambda_1} = e^{\lambda_1(z-1)}.
$$

Similarly, X_2 has generating function $G_2(z) = e^{\lambda_2(z-1)}$.

Now, we have for the generating function of $X = X_1 + X_2$ that

$$
G(z) = E\Big[z^X\Big] = E\Big[z^{X_1} z^{X_2}\Big] = E\Big[z^{X_1}\Big] E\Big[z^{X_2}\Big] = G_1(z) G_2(z),
$$

since X_1 and X_2 are independent random variables (and hence so are z^{X_1} and z^{X_2}). Substituting for G_1 and G_2, we find that $G(z) = e^{(\lambda_1 + \lambda_2)(z-1)}$, which we recognise as the generating function of a Poisson random variable with parameter $\lambda_1 + \lambda_2$. This completes the proof.

5. (a) By the conditional probability formula, we have for all $t, u \geq 0$ that

$$
P(T > t + u | T > u) = \frac{P(\{T > t + u\} \cap \{T > u\})}{T > u} = \frac{P(T > t + u)}{T > u},
$$

since the event $T > t + u$ is a subset of the event $T > u$, and hence their intersection is the event $T > t + u$. Now, recall that since T is exponentially distributed with parameter μ, $P(T > t) = e^{-\mu t}$ for all $t \geq 0$. Substituting this above,

$$
P(T > t + u | T > u) = \frac{\exp(-\mu(t + u))}{\exp(-\mu u)} = e^{-\mu t} = P(T > t).
$$

(b) i. Since T_1 and T_2 are independent, we have for arbitrary $t > 0$ that

$$\begin{aligned} P(T > t) &= P(\min\{T_1, T_2\} > t) = P(T_1 > t \text{ and } T_2 > t) \\ &= P(T_1 > t)P(T_2 > t). \end{aligned}$$

Now, using the fact that T_1 and T_2 are exponentially distributed with parameters λ_1 and λ_2, we get

$$P(T > t) = e^{-\lambda_1 t} e^{-\lambda_2 t} = e^{-(\lambda_1+\lambda_2)t},$$

from which we recognise that T is exponentially distributed with parameter $\lambda_1 + \lambda_2$.

ii. It is fairly easy to calculate the probability that $T = T_1$. We have

$$P(T = T_1) = P(T_2 \geq T_1) = \int_0^\infty f_{T_1}(x) P(T_2 \geq x) dx.$$

While this calculation yields $P(T = T_1) = \lambda_1/(\lambda_1 + \lambda_2)$, it doesn't tell us that this probability doesn't depend on the value t taken by this random variable.

It may not be immediately obvious why we are making an issue of this point. Consider, for example, that $\lambda_1 = 1$ and $\lambda_2 = 999,999$. So the chance that $T = T_1$ is 1 in a million. Moreover, T_1 typically takes values around 1, whereas T_2 typically takes values around 10^{-6}. The statement that we are asked to prove is that, even if we are told, say, that $T = 1.3$, then conditional on this information, the probability that $T = T_1$ is still 1 in a million. Hopefully, you find that claim counter-intuitive, and see that there is something to be proved here!

In order to get to the result we want, let us compute the conditional probability

$$\begin{aligned} P\big(\{T = T_1\} \cap \{T \geq t\}\big) &= P(t \leq T_1 \leq T_2) = \int_{x=t}^\infty \int_{y=x}^\infty f_{T_1}(x) f_{T_2}(y) dy dx \\ &= \int_{x=t}^\infty \lambda_1 e^{-\lambda_1 x} \Big(\int_{y=x}^\infty \lambda_2 e^{-\lambda_2 y} dy \Big) dx \\ &= \int_{x=t}^\infty \lambda_1 e^{-(\lambda_1+\lambda_2)x} dx \\ &= \frac{\lambda_1}{\lambda_1 + \lambda_2} e^{-(\lambda_1+\lambda_2)t} = P(T = T_1) P(T \geq t), \end{aligned}$$

since T is exponentially distributed with parameter $\lambda_1 + \lambda_2$. This shows that these two events are independent, for any value of t. In other words, the probability that $T = T_1$ is independent of the value of T.