Homework 5 - Markov chains

1) Stationary distribution. Consider a Markov chain (MC) $X_{\mathbb{N}} = X_0, X_1, \ldots, X_n, \ldots$ with state space $S = \{1, 2\}$ and transition probability matrix

$$\mathbf{P} = \begin{pmatrix} 1/4 & 3/4 \\ 1/5 & 4/5 \end{pmatrix}.$$

Compute the stationary distribution of $X_{\mathbb{N}}$. What is $\lim_{n\to\infty} P_{22}^n$? What is $\lim_{n\to\infty} \frac{1}{n}\sum_{k=0}^{n-1} \mathbb{I}\{X_k = 1\}$? Justify your answers.

2) A cloudy town. A certain town never has two sunny days in a row. Each day is classified as being either sunny, cloudy (but dry), or rainy. If it is sunny one day, then it is equally likely to be either cloudy or rainy the next day. If it is rainy or cloudy one day, then there is one chance in two that it will be the same the next day, and if it changes then it is equally likely to be either of the other two possibilities. In the long run, what proportion of days are sunny? What proportion are cloudy?

3) A store's supply and demand chain. During each day, a non-negative integer number of customers arrives to a store to purchase a particular product. Each customer purchases a unit of the product when the product is in stock. Customers who do not find the product in stock depart without making a purchase. The store orders $q > 0$ new units of the product from its supplier at the end of the day (after that day's demand has materialized). However, the supplier is not completely reliable, and each day, with probability α independent of everything else, the order is permanently lost in which case the order does not arrive to the store. If the order is not lost, it arrives to the store before the beginning of the next day. Suppose the sequence of daily demands $D_{\mathbb{N}} = D_0, D_1, \ldots, D_n, \ldots$ is i.i.d. with $\mathrm{P}[D_n = d] = p(d)$ for $d \geq 0$.

Let $X_{\mathbb{N}} = X_0, X_1, \ldots, X_n, \ldots$ be the MC that represents the amount of product in stock at the beginning of each day. Derive the transition probabilities of $X_{\mathbb{N}}$.

4) Non-invertible function of a Markov chain. Suppose that $X_{\mathbb{N}} = X_0, X_1, \ldots, X_n, \ldots$ is a MC with state space $S = \{1, 2, 3\}$, transition probability matrix

$$\mathbf{P} = \begin{pmatrix} 0 & 2/3 & 1/3 \\ 1/4 & 1/4 & 1/2 \\ 3/4 & 1/4 & 0 \end{pmatrix}$$

and initial distribution $\mathrm{P}[X_0 = 1] = 1/5$, $\mathrm{P}[X_0 = 2] = 2/5$ and $\mathrm{P}[X_0 = 3] = 2/5$. Suppose that the random process $Y_{\mathbb{N}} = Y_0, Y_1, \ldots, Y_n, \ldots$ satisfies $Y_n = g(X_n)$, $n \geq 0$, where $g(1) = 1$ and $g(2) = g(3) = 2$. You are asked to calculate $\mathrm{P}\left[Y_2 = 1 \,\middle|\, Y_1 = 2, Y_0 = 1\right]$.

5) A non-irreducible Markov chain. Consider a MC $X_{\mathbb{N}} = X_0, X_1, \ldots, X_n, \ldots$ with state space $S = \{1, 2, 3\}$ and transition probability matrix

$$\mathbf{P} = \begin{pmatrix} 1/3 & 1/6 & 1/2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

A) Determine the communication classes and classify them as transient or recurrent.

B) Determine $\lim_{n\to\infty} P_{ij}^n$ for each i and j.

C) Give three examples of stationary distributions for $\mathbf{P}$.

D) Suppose that $X_0 = 1$. Does $\lim_{n\to\infty} \frac{1}{n}\sum_{k=0}^{n-1} \mathbb{I}\{X_k = 2\}$ exist? If so, what is its distribution?

6) A null-recurrent Markov chain. Consider a MC with state space $S = \{1, 2, \ldots\}$ and transition probabilities $P_{i,i+1} = i/(i+1)$ and $P_{i1} = 1/(i+1)$ for $i = 1, 2, \ldots$ Show that the MC is null recurrent.

7) Random access in communication networks. A pervasive situation in communication systems is to have a common infrastructure access point (AP) serving a group of physically distributed devices. This is, e.g., the situation of your cellphone sending information to the cellular base station, your laptop transmitting to the 802.11 wireless

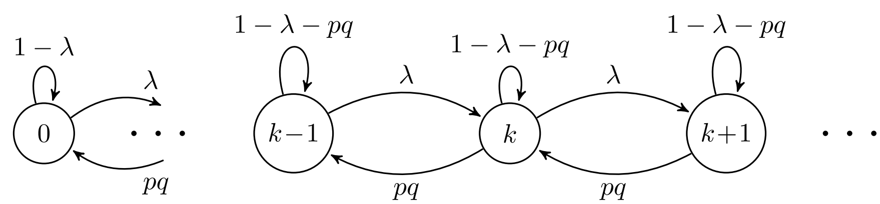

Fig. 1. State transition diagram for queue lengths.

router, or a group of satellites communicating with a common ground station. In these examples there is a common AP, the base station, the router or the ground station, serving a group of distributed devices, cellphones, laptops or satellites. A problem in the "uplink" communication from physically distributed terminals to the AP is how to separate the information transmitted by different devices. One possibility is to assign different times or frequencies to different terminals. This is called time or frequency division multiplexing. Another possibility is to let terminals transmit at random and hope for luck to avoid simultaneous transmissions. This is called random access (RA). RA is appealing because it requires little coordination between terminals. However, it is not clear that RA is a viable communication strategy. In this exercise, we will define RA protocols and study pertinent performance indicators.

Start considering J distributed terminals $1, 2, \ldots J$ and time slots $1, 2, \ldots, n, \ldots$ during which they attempt communication with the AP. Each terminal maintains a buffer in which packets are stored to await transmission to the AP. Denote as Q_{jn} the number of packets in the j-th terminal queue at time n. In each time slot n terminal j receives a packet with probability λ and if the transmission buffer is not empty it attempts to transmit a packet with probability p. If the communication is successful, something that for now we say happens with probability q, the packet leaves the terminal's queue since it was delivered to the AP. If not the packet stays in queue to await retransmission at a later time.

To simplify analysis we introduce two assumptions:

(A) *No concurrence*. There are no simultaneous transmission and packet arrivals in the same time slot.

(B) *Dominant system*. If the transmission queue of terminal j is empty at time n, i.e., if $Q_{jn} = 0$ terminal j decides to transmit a dummy packet with probability p. This packet carries no information and is transmitted for the sole purpose of interfering with other terminals. The purpose of the dummy packet is to render the probability of receiving interference from terminal j independent of the number Q_{jn} of packets in its queue.

A) Model as Markov chain. The number Q_{jn} of packets in queue is a MC. Assuming no concurrence as per (A), transition probabilities are given by

$$\begin{aligned} &\mathrm{P}\left[Q_{j,n+1} = k+1 \,\middle|\, Q_{jn} = k\right] = \lambda && \mathrm{P}\left[Q_{j,n+1} = k-1 \,\middle|\, Q_{jn} = k\right] = pq \\ &\mathrm{P}\left[Q_{j,n+1} = k \,\middle|\, Q_{jn} = k\right] \quad = 1 - \lambda - pq && \mathrm{P}\left[Q_{j,n+1} = 0 \,\middle|\, Q_{jn} = 0\right] \quad = 1 - \lambda \end{aligned} \tag{1}$$

where the first equation is valid for all k, and the second and third equation are valid for $k \neq 0$. A state transition diagram is shown in Fig.1. Explain how to obtain these transition probabilities. Is this MC recurrent for all or some combination of parameter values? Explain. Argue that the MC is ergodic for $0 < \lambda < pq$.

B) Limit distribution. Assuming the communication system has been operating for a long time n, we want to find the probability distribution $\mathrm{P}\left[Q_{jn} = k\right]$. Formally, we want to find

$$\pi_k := \lim_{n \to \infty} \mathrm{P}\left[Q_{jn} = k\right], \qquad \text{for all } k. \tag{2}$$

To find the limit distributions in (2) notice that the system is fundamentally different if $0 < \lambda < pq$, $0 < \lambda = pq$ and $\lambda > pq > 0$. For $0 < \lambda < pq$ (ergodic MC) show that a solution of the form $\pi_k := c(\lambda/pq)^k$ is correct. Determine c. What are the limiting probabilities for $0 < \lambda = pq$ and $\lambda > pq > 0$?

C) Probability of empty queue and probability of minimal wait. An important performance metric is the probability of a queue becoming empty. What is the value of this probability for n sufficiently large? Another related important metric is the probability T_1 of the packet being transmitted in the first slot after arrival. Also important is the probability S_1 of the packet being successfully communicated in the first slot after arrival. Compute T_1 and S_1 for n large.

D) Expected queue length. Yet another performance metric is the expected queue length $\mathbb{E}[Q_{jn}]$. Compute this metric for n large. Formally, we want you to report

$$\lim_{n\to\infty} \mathbb{E}[Q_{jn}]. \tag{3}$$

E) Probability of successful transmission and optimal transmission probability p. Assuming validity of the dominant system hypothesis as per (B), compute the probability of successful transmission pq (note that this will happen when none of the other terminals transmits in the same time slot). Prove that making $p = 1/J$ maximizes the probability of successful transmission pq and the probability of having empty queues $\mathrm{P}[Q_{nj} = 0]$ while minimizing the expected queue length $\mathbb{E}[Q_{nj}]$. For this value of p write down the corresponding q and show that as the number of terminals $J \to \infty$ the number of successful transmissions converges to $1/e \approx 0.36$. This is remarkable. RA utilizes a decent 36% of the available resources without any coordination overhead between terminals.

F) Average time occupancies. It is possible to argue that the probabilities in Parts *B* and *C* as well as the expected value computed in Part *D* are of little practical value. What these probabilities express is an average across all possible paths of the communication system. Say we run the system once and obtain a certain path $Q_{jn}^{(1)}$, a second run yields a path $Q_{jn}^{(2)}$ and so on. The computed probabilities measure how likely different events are across these different realizations of the random process. In a practical implementation however, we need a performance guarantee for every run of the system. One such performance metric is, e.g., the following time average

$$\bar{p}_k := \lim_{n\to\infty} \frac{1}{n} \sum_{m=1}^{n} \mathbb{I}\{Q_{jm} = k\}, \qquad \text{for all } k \tag{4}$$

which tells you the fraction of time there were k packets in the queue Q_{jm}. Different from the π_k in (2) the $\bar{p}_k$ in (4) is a performance metric associated with each particular experiment. Find $\bar{p}_k$ and explain why it is useful. What similarly useful performance metric can you define that is equivalent to the expected value in (3)?

G) System simulation. Write a function to simulate this system assuming validity of the no concurrence hypothesis, but without assuming validity of the dominant system hypothesis. The function should accept as parameters the number of users J, transmission probability p, arrival rates λ and a time horizon parameter N. The function returns the history of the number of packets in queue R_{jn} (denoted differently from Q_{jn} where (B) is assumed) between times 0 and N for all J queues. Queue lengths are initialized at 0 for all queues. Run your function for $J = 16$ terminals, optimal transmission probability p computed in Part *E*, $\lambda = 0.9pq$ and $N = 10^4$. Report a graph with the path followed by terminals 1 through 4.

H) Compare numerical and analytical results. Define the limit distribution of the simulated system without the dominant system hypothesis as

$$\xi_k := \lim_{n\to\infty} \mathrm{P}[R_{jn} = k], \qquad \text{for all } k. \tag{5}$$

These probabilities cannot be computed in closed form. Use your simulation code to estimate the probability distribution function in (5). You are required to do this from a single run of the function in Part *G*. Plot your estimate of the limit distribution ξ_k and compare with π_k. Comment on the usefulness of the dominant system approach.