Solutions to 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

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

To obtain the stationary distribution $\boldsymbol{\pi} = [\pi_1, \pi_2]^T$, note that the MC is ergodic and hence the unique solution is given by

$$\begin{pmatrix} \pi_1 \\ \pi_2 \end{pmatrix} = \begin{pmatrix} 1/4 & 1/5 \\ 3/4 & 4/5 \end{pmatrix} \begin{pmatrix} \pi_1 \\ \pi_2 \end{pmatrix}, \quad \pi_1 + \pi_2 = 1.$$

Solving the linear system yields $\boldsymbol{\pi} = [4/19, 15/19]^T$. Now, since the MC is ergodic the limiting probabilities converge to the stationary distribution. In particular, one has

$$\lim_{n \to \infty} P_{22}^n = \pi_2 = \frac{15}{19}.$$

Finally, by virtue of the ergodic theorem the time average $\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \mathbb{I}\{X_k = 1\}$ also converges. The long-run fraction of time the MC visits state 1 is thus given by

$$\lim_{n \to \infty} \frac{1}{n} \sum_{k=0}^{n-1} \mathbb{I}\{X_k = 1\} = \pi_1 = \frac{4}{19}.$$

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.

Let $X_{\mathbb{N}} = X_0, X_1, \ldots, X_n, \ldots$ be the random process describing the weather evolution of the given town, with n denoting the day number. Given the nature of the evolution of the process (meaning the weather today only depends on the state of the weather yesterday), $X_{\mathbb{N}}$ is a MC with state space

$$S = \{\text{sunny}, \text{cloudy}, \text{rainy}\} = \{1, 2, 3\}.$$

From the weather state evolution described above, the transition probability matrix is

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

By simple inspection of the transition matrix, it is clear that every state is accessible from both the other remaining states. Thus all states communicate and form a single class. This implies that the MC is irreducible. Because the state space is finite, this class is recurrent and furthermore it has to be positive recurrent. Finally, since $P_{22} > 0$ then state 2 is aperiodic. Since the MC is irreducible and periodicity is a class property, it follows that all states are aperiodic. Overall, the MC is ergodic and therefore a unique stationary distribution $\boldsymbol{\pi} = [\pi_1, \pi_2, \pi_3]^T$ exists satisfying the conditions

$$\begin{pmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \end{pmatrix} = \begin{pmatrix} 0 & 1/4 & 1/4 \\ 1/2 & 1/2 & 1/4 \\ 1/2 & 1/4 & 1/2 \end{pmatrix} \begin{pmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \end{pmatrix}, \quad \pi_1 + \pi_2 + \pi_3 = 1.$$

Solving for $\boldsymbol{\pi}$ yields

$$\boldsymbol{\pi} = [1/5, 2/5, 2/5]^T.$$

This information allows one to answer the questions of interest because the long-run fraction of days that are sunny is $\pi_1 = 1/5$, whereas the long-run fraction of days that are cloudy is $\pi_2 = 2/5$.

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, and we want to derive the transition probabilities of $X_{\mathbb{N}}$. To that end, introduce the i.i.d. sequence of random variables $B_{\mathbb{N}} = B_0, B_1, \ldots, B_n, \ldots$, where B_1 is Bernoulli distributed with parameter $1-\alpha$. This way, the amount of ordered product received before the beginning of day $n+1$ is $qB_{n+1} = 0$ if the order is lost (something that happens with probability $\mathrm{P}[B_{n+1} = 0] = \alpha$), or $qB_{n+1} = q$ if the order is not lost (something that happens with probability $\mathrm{P}[B_{n+1} = 1] = 1 - \alpha$). All in all, for $n \geq 0$, the amount of product in stock at the beginning of day $n+1$ is determined by

$$X_{n+1} = \max\{0, X_n - D_n\} + qB_{n+1}.$$

Notice that $\max\{0, a\} = a$ if $a \geq 0$, and $\max\{0, a\} = 0$ if $a < 0$. Define $\bar{X}_n = \max\{0, X_n - D_n\}$. Hence, the above expression enforces the physical constraint that $\bar{X}_n \geq 0$ always, and $\bar{X}_n = 0$ when the demand D_n exceeds the available product in stock at the beginning of day n, i.e., X_n.

The transition probabilities are given by

$$P_{ij} = \mathrm{P}\left[X_{n+1} = j \,\middle|\, X_n = i\right] = \mathrm{P}\left[\bar{X}_n + qB_{n+1} = j \,\middle|\, X_n = i\right].$$

To further work out the expression, it is convenient to condition on the value of B_{n+1} (meaning the order is received or not), so that using the law of total probability one obtains

$$\begin{aligned}
P_{ij} &= \mathrm{P}\left[\bar{X}_n + qB_{n+1} = j \,\middle|\, X_n = i\right] \\
&= \sum_{b=0}^{1} \mathrm{P}\left[\bar{X}_n + qB_{n+1} = j \,\middle|\, X_n = i, B_{n+1} = b\right] \mathrm{P}[B_{n+1} = b] \\
&= \sum_{b=0}^{1} \mathrm{P}\left[\bar{X}_n + qb = j \,\middle|\, X_n = i\right] \mathrm{P}[B_{n+1} = b] \\
&= \mathrm{P}\left[\bar{X}_n = j \,\middle|\, X_n = i\right] \alpha + \mathrm{P}\left[\bar{X}_n + q = j \,\middle|\, X_n = i\right](1-\alpha).
\end{aligned} \tag{1}$$

Focusing on the calculation of $\mathrm{P}\left[\bar{X}_n = j \,\middle|\, X_n = i\right]$ first, for $j = 0$ one has

$$\mathrm{P}\left[\bar{X}_n = 0 \,\middle|\, X_n = i\right] = \mathrm{P}\left[D_n \geq X_n \,\middle|\, X_n = i\right] = \mathrm{P}[D_n \geq i] = \sum_{k=i}^{\infty} p(k). \tag{2}$$

For $j \geq 0$, then day's n demand does not match the amount of available product i so

$$\mathrm{P}\left[\bar{X}_n = j \,\middle|\, X_n = i\right] = \mathrm{P}\left[X_n - D_n = j \,\middle|\, X_n = i\right] = \mathrm{P}[D_n = i - j] = \begin{cases} p(i-j), & i \geq j \\ 0, & i < j \end{cases}, \quad j > 0. \tag{3}$$

Moving on to the calculation of $\mathrm{P}\left[\bar{X}_n + q = j \,\middle|\, X_n = i\right]$, notice that for $0 \leq j < q$ then

$$\mathrm{P}\left[\bar{X}_n + q = j \,\middle|\, X_n = i\right] = 0, \quad 0 \leq j < q. \tag{4}$$

This is because q units of product were received, so it is impossible to have a stock level X_{n+1} below q. For $j = q$, one has

$$\mathrm{P}\left[\bar{X}_n + q = q \,\middle|\, X_n = i\right] = \mathrm{P}\left[\bar{X}_n = 0 \,\middle|\, X_n = i\right] = \sum_{k=i}^{\infty} p(k), \quad j = q. \tag{5}$$

as derived in (2). Finally, for $j > q$ the probability is

$$\begin{aligned}
\mathrm{P}\left[\bar{X}_n + q = j \,\middle|\, X_n = i\right] &= \mathrm{P}\left[X_n - D_n + q = j \,\middle|\, X_n = i\right] \\
&= \mathrm{P}[D_n = i - j + q] = \begin{cases} p(i-j+q), & i - j + q \geq 0 \\ 0, & i - j + q < 0 \end{cases}, \quad j > q.
\end{aligned} \tag{6}$$

Now, what is left is to carefully put all pieces (2)-(6) together back in (1). For all $i \geq 0$, the transition probabilities are thus given by

$$
\begin{aligned}
P_{i0} &= \alpha \sum_{k=i}^{\infty} p(k), \\
P_{ij} &= \begin{cases} \alpha p(i-j), & i \geq j \\ 0, & i < j \end{cases}, \quad \text{for } 0 < j < q, \\
P_{iq} &= \begin{cases} \alpha p(i-q) + (1-\alpha) \sum_{k=i}^{\infty} p(k), & i \geq q \\ (1-\alpha) \sum_{k=i}^{\infty} p(k), & i < q \end{cases}, \\
P_{ij} &= \begin{cases} \alpha p(i-j) + (1-\alpha) p(i-j+q), & i-j \geq 0 \\ (1-\alpha) p(i-j+q), & i-j+q \geq 0,\ i-j < 0 \\ 0, & i-j+q < 0 \end{cases}, \quad \text{for } j > q.
\end{aligned}
$$

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 $g(2) = g(3) = 2$.

To calculate $\mathrm{P}\left[Y_2 = 1 \,\middle|\, Y_1 = 2, Y_0 = 1\right]$, one can resort to the definition of conditional probability and write

$$
\begin{aligned}
\mathrm{P}\left[Y_2 = 1 \,\middle|\, Y_1 = 2, Y_0 = 1\right] &= \frac{\mathrm{P}[Y_2 = 1, Y_1 = 2, Y_0 = 1]}{\mathrm{P}[Y_1 = 2, Y_0 = 1]} \\
&= \frac{\mathrm{P}[X_2 = 1, X_1 \in \{2,3\}, X_0 = 1]}{\mathrm{P}[X_1 \in \{2,3\}, X_0 = 1]} \\
&= \frac{\mathrm{P}[X_2 = 1, X_1 = 2, X_0 = 1] + \mathrm{P}[X_2 = 1, X_1 = 3, X_0 = 1]}{\mathrm{P}[X_1 = 2, X_0 = 1] + \mathrm{P}[X_1 = 3, X_0 = 1]} \\
&= \frac{\mathrm{P}[X_0 = 1] P_{12} P_{21} + \mathrm{P}[X_0 = 1] P_{13} P_{31}}{\mathrm{P}[X_0 = 1] P_{12} + \mathrm{P}[X_0 = 1] P_{13}} = \frac{5}{12}.
\end{aligned}
$$

Because the function g is not invertible, it turns out that $Y_{\mathbb{N}}$ is not Markov chain. Actually, one can verify that

$$
\mathrm{P}\left[Y_2 = 1 \,\middle|\, Y_1 = 2, Y_0 = 1\right] \neq \mathrm{P}\left[Y_2 = 1 \,\middle|\, Y_1 = 2\right] = \frac{17}{36}.
$$

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) The MC has three communication classes, namely $\mathcal{T}_1 = \{1\}$, $\mathcal{R}_1 = \{2\}$, and $\mathcal{R}_2 = \{3\}$; see also Fig. 1. Notice that both states 2 and 3 are absorbing, hence each of them belongs to its own recurrent class. State 1 on the other hand belongs to a transient class, since given $X_0 = 1$ there is a nonzero probability that the MC transitions to either one of the absorbing states and never again visits state 1.

B) Even though the MC is not ergodic, each of the states is aperiodic so we can determine the limiting probabilities $P_{ij}^{\infty} = \lim_{n \to \infty} P_{ij}^{n}$ for each i and j. First, because states 2 and 3 are absorbing, then it follows that

$$
\begin{aligned}
P_{21}^{\infty} = 0, \quad P_{22}^{\infty} = 1, \quad P_{23}^{\infty} = 0, \\
P_{31}^{\infty} = 0, \quad P_{32}^{\infty} = 0, \quad P_{33}^{\infty} = 1
\end{aligned}
$$

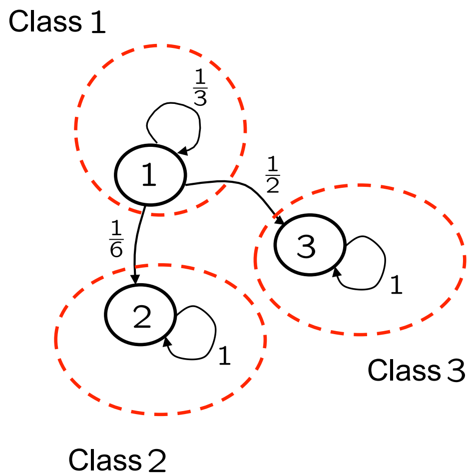

Fig. 1. State transition diagram for the MC with transition probability matrix $\mathbf{P}$. The MC has three communication classes.

Because state 1 is transient, one has $P_{11}^{\infty}=0$. To compute P_{12}^{∞} and P_{13}^{∞}, introduce the matrix of limiting probabilities

$$\mathbf{P}^{\infty}=\begin{pmatrix} 0 & P_{12}^{\infty} & P_{13}^{\infty} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

that must be a fixed point of the recursion $\mathbf{P}^{n}=\mathbf{P}\times\mathbf{P}^{n-1}$, implying $\mathbf{P}^{\infty}=\mathbf{P}\times\mathbf{P}^{\infty}$. This identity gives the equations required to determine P_{12}^{∞} and P_{13}^{∞}, which yield

$$\begin{aligned} \frac{P_{12}^{\infty}}{3}+\frac{1}{6}=P_{12}^{\infty} \;\Rightarrow\; P_{12}^{\infty}=\frac{1}{4} \\ \frac{P_{13}^{\infty}}{3}+\frac{1}{2}=P_{13}^{\infty} \;\Rightarrow\; P_{13}^{\infty}=\frac{3}{4}. \end{aligned}$$

All in all, the matrix of limiting probabilities is

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

C) The matrix of limiting probabilities $\mathbf{P}^{\infty}$ suggests the following three stationary distributions

$$\boldsymbol{\pi}_1=[0,1/4,3/4]^T,\;\; \boldsymbol{\pi}_2=[0,1,0]^T,\;\; \boldsymbol{\pi}_3=[0,0,1]^T.$$

It is straightforward to check that $\boldsymbol{\pi}_i=\mathbf{P}^T\boldsymbol{\pi}_i$, for each $i=1,\ldots,3$.

D) From the first row of $\mathbf{P}^{\infty}$, one can claim that given $X_0=1$ the MC will end up in state 2 (and stay there forever) with probability $1/4$, or else end up in state 3 (and stay there forever) with probability $3/4$. This observation immediately leads to the conclusion that

$$\lim_{n\to\infty}\frac{1}{n}\sum_{k=0}^{n-1}\mathbb{I}\left\{X_k=2\right\}$$

will almost surely converge to a random variable Y that is Bernoulli distributed with parameter $1/4$.

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 $i=1,2,\ldots$. The transition probability matrix is

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

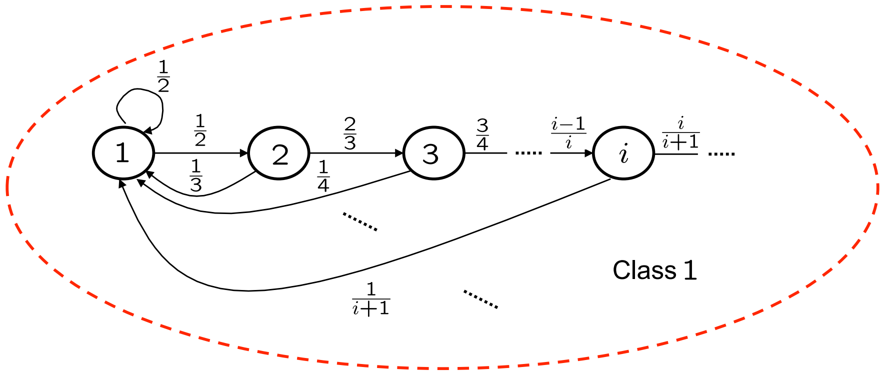

Fig. 2. State transition diagram for the MC with transition probability matrix $\mathbf{P}$. The MC has three communication classes.

The state transition diagram is depicted in Fig. 2, from where it is apparent that state 1 is accessible from all other states (via single-step transitions). Moreover, all other states are clearly accessible from state 1. This shows that all states communicate and form a single (infinitely large) class, so the MC is irreducible.

The goal is to establish the recurrence properties of the MC. Because the MC is irreducible and recurrence is a class property, it suffices to analyze the recurrence properties of a single state, say state 1. To that end, define the return time T_1 to state 1 as

$$T_1 = \min\{n > 0 \mid X_n = 1\}$$

State 1 (and hence the MC) will be recurrent if $\mathrm{P}\left[T_1 < \infty \,\middle|\, X_0 = 1\right] = 1$, which is of course equivalent to $\mathrm{P}\left[T_1 = \infty \,\middle|\, X_0 = 1\right] = 0$. As it can be readily seen from the state transition diagram, the probability that the MC never returns to state 1 given that it started in that state is given by

$$\begin{aligned}\mathrm{P}\left[T_1 = \infty \,\middle|\, X_0 = 1\right] &= P_{12} \times P_{23} \times P_{34} \times \ldots \times P_{i,i+1} \times \ldots \\ &= \lim_{n\to\infty} \prod_{i=1}^{n} P_{i,i+1} = \lim_{n\to\infty} \prod_{i=1}^{n} \frac{i}{i+1} = \lim_{n\to\infty} \frac{1}{n+1} = 0.\end{aligned}$$

In obtaining the second last inequality we have used the fact that terms in successive products of probabilities cancel out, and only the numerator from the first and denominator from the last probability survive. This establishes that the MC is recurrent as desired.

To further show that it is null recurrent, it suffices to focus on first state and verify that $\mathbb{E}\left[T_1 \,\middle|\, X_0 = 1\right] = \infty$. Recalling the definition of T_1, one can obtain the relevant conditional pmf

$$\mathrm{P}\left[T_1 = n | X_0 = 1\right] = \begin{cases} P_{11}, & n = 1 \\ P_{12} \times P_{21}, & n = 2 \\ P_{12} \times P_{23} \times P_{31}, & n = 3 \\ \vdots, & \vdots \\ P_{12} \times P_{23} \times \ldots \times P_{i-1,i} \times P_{i1}, & n = i \\ \vdots, & \vdots \end{cases}$$

whose general term can be simplified as

$$\mathrm{P}\left[T_1 = n \,\middle|\, X_0 = 1\right] = \left(\prod_{i=1}^{n-1} P_{i,i+1}\right) P_{n1} = \left(\prod_{i=1}^{n-1} \frac{i}{i+1}\right) \frac{1}{n+1} = \frac{1}{n(n+1)}.$$

The conditional expectation is from the definition

$$\mathbb{E}\left[T_1 \,\middle|\, X_0 = 1\right] = \sum_{n=1}^{\infty} n \times \mathrm{P}\left[T_1 = n \,\middle|\, X_0 = 1\right] = \sum_{n=1}^{\infty} n \times \frac{1}{n(n+1)} = \sum_{n=1}^{\infty} \frac{1}{n+1} = \infty.$$

The infinite sum diverges, establishing that state 1 (hence, the MC) is null-recurrent as desired.

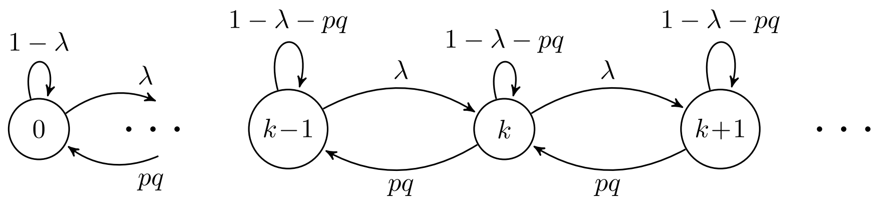

Fig. 3. State transition diagram for the Markov chain representing the queue length at an arbitrary terminal.

7) Random access in communication networks.

A) Model as Markov chain. From the information provided in the problem statement, for each terminal $j = 1, \ldots, J$ then

$$\mathrm{P}\left[\text{Terminal receives packet in time slot } n\right] = \lambda \tag{7}$$

for all n, and

$$\mathrm{P}\left[\text{Terminal successfully transmits a packet in time slot } n\right] = pq \tag{8}$$

for every n such that $Q_{jn} > 0$. We now use (7) and (8) to compute the transition probabilities of the MC. From (7), it follows that

$$\mathrm{P}\left[Q_{j,n+1} = k+1 \,\middle|\, Q_{jn} = k\right] = \lambda \tag{9}$$

since this transition occurs whenever a packet is received in time slot $n+1$. Note that, given the assumption (A) of no concurrence, we do not have to account for the probability of not having a successful transmission because the arrival of a packet already ensures that there is no transmission in that same time step. Similarly,

$$\mathrm{P}\left[Q_{j,n+1} = k-1 \,\middle|\, Q_{jn} = k\right] = pq \tag{10}$$

because the transition from having k packets to $k-1$ packets in the queue occurs whenever a successful transmission takes place in time slot $n+1$. The transition probability in (10) is valid for all $k > 0$. Whenever the queue is empty, i.e. $k = 0$ there is no packet transmission apart from the dummy packets in assumption (B) which do not alter the length of the queue. The probability that the queue backlog remains unchanged must be unfolded into the cases of positive and null backlogs. For $k > 0$, for the length to remain unchanged we must have neither an arrival nor a successful transmission of a packet, which yields

$$\mathrm{P}\left[Q_{j,n+1} = k \,\middle|\, Q_{jn} = k\right] = 1 - \lambda - pq. \tag{11}$$

However, when the queue is empty, by (A) not having an arrival suffices to ensure that the queue length will be null in the next time slot, thus,

$$\mathrm{P}\left[Q_{j,n+1} = 0 \,\middle|\, Q_{jn} = 0\right] = 1 - \lambda. \tag{12}$$

All other transition probabilities are zero since the queue length may not vary by more than one unit in any given time slot. The complete MC model is presented in Figure 3.

Moving on to recurrence and ergodicity analysis, a quick inspection reveals that recurrence depends on the values of the parameters λ and pq. For example, if we assume that $\lambda = pq = 0$, then every state in the MC is an absorbing state. Consequently, no state is accessible from any other state, implying that every state is a different class. Moreover, being absorbing states, every class is positive recurrent. A different situation arises when $\lambda = 0$ and $pq > 0$. In this case, every state is still a class since there is no communication (although there is one sided accessibility) between states but the only positive recurrent class is state 0 with every other state being a transient class. In the opposite situation where $\lambda > 0$ and $pq = 0$, every state is still a class but all of them are transient since the queue length inevitable grows towards $+\infty$. Whenever both $\lambda > 0$ and $pq > 0$, the MC is irreducible, since every state communicates with each other. However, the recurrence of this class depends on the relation between λ and pq. Whenever $\lambda > pq > 0$, similar to right-biased random walk studied in class the MC has a positive drift towards infinity. Thus, after being in an arbitrary state i, there is a positive probability of never returning to this state. Consequently, the MC is composed of one transient class. Without going into details of a formal proof, some

TABLE I
SUMMARY OF THE RECURRENCE ANALYSIS. RECURRENCE DEPENDS ON THE PARAMETERS λ AND pq.

Parameters	Effect on recurrence
$\lambda = 0$, $pq = 0$	Every state is a class. Every class is positive recurrent.
$\lambda = 0$, $pq > 0$	Every state is a class. Class 0 is positive recurrent. All other classes are transient.
$\lambda > 0$, $pq = 0$	Every state is a class. Every class is transient.
$\lambda > pq > 0$	The MC is irreducible and transient.
$pq > \lambda > 0$	The MC is irreducible and positive recurrent.
$pq = \lambda > 0$	The MC is irreducible and null recurrent.

intuition can be gained if we analyze the expected behavior of the MC. If at time n the MC is at state $i > 0$, then for time $n + 1$ we have that

$$\mathbb{E}\left[X_n + 1 \,\middle|\, X_n = i\right] = pq(i-1) + (1 - \lambda - pq)i + \lambda(i+1) > i \tag{13}$$

where the strict inequality comes from the fact that $\lambda > pq$. The bound (13) implies that the expected state of the MC always increases, which is an intuitive way of explaining there is a positive probability of never returning to a given state. On the contrary, if $pq > \lambda > 0$, the MC suffers a negative drift towards zero. By analyzing the expected value of the MC as in (13), it follows that state 0 is always re-visited making it positive recurrent. Moreover, since the MC is irreducible, the whole chain must be positive recurrent. For the special case in which $pq = \lambda > 0$, an intermediate behavior is obtained. Although the probability of returning to any state is 1, the expected time it takes to return is infinite. Consequently, the chain is null recurrent. The recurrence analysis is summarized in Table I.

Given the detailed recurrence analysis in Table I, studying ergodicity of this MC is immediate. By definition, for a state to be ergodic it must be positive recurrent and aperiodic. For a MC to be ergodic it must be irreducible with ergodic states. Hence, the only parameter combination that results in an ergodic MC is $pq > \lambda > 0$.

B) Limit distribution. Since the MC is ergodic for the parameter configuration $pq > \lambda > 0$, we can compute the limit distribution π_k given by

$$\pi_k := \lim_{n\to\infty} \mathrm{P}\left(Q_{jn} = k\right), \quad k \geq 0. \tag{14}$$

Intuitively, π_k is the probability of finding k packets in the queue of any given terminal, assuming that the communication system has been operating for a long time.

From the transition probabilities in Figure 3, the limit distribution we are looking for must satisfy

$$\pi_0 = (1 - \lambda)\pi_0 + pq\pi_1 \tag{15}$$

for the state when the queue is empty, and

$$\pi_k = \lambda\pi_{k-1} + (1 - \lambda - pq)\pi_k + pq\pi_{k+1}, \quad k > 0. \tag{16}$$

To solve the system of difference equations (15)-(16), we propose a solution of the form $\pi_k = c\alpha^k$ for some constants $c > 0$ and $\alpha > 0$ to be determined. Because the solution must be unique, if that candidate works out we can be sure that those are the limiting probabilities. We now use equations (15) and (16) plus the fact that the π_k must sum up to 1 in order to validate the proposed functional form of π_k, and to find the value of the constants α and c. Specifically, by substituting $\pi_k = c\alpha^k$ into (15) we obtain

$$c\alpha^0 = (1 - \lambda)c\alpha^0 + pqc\alpha^1 \tag{17}$$

from where it follows that $\alpha = \lambda/pq$. Similarly, substituting $\pi_k = c\alpha^k$ into (16) yields

$$c\alpha^k = \lambda c\alpha^{k-1} + (1 - \lambda - pq)c\alpha^k + pqc\alpha^{k+1}, \tag{18}$$

which after some minor manipulations (divide all summands by $c\alpha^{k-1}$) becomes the following quadratic equation in α

$$pq\alpha^2 - (\lambda + pq)\alpha + \lambda = 0. \tag{19}$$

The solutions of (19) are $\alpha = 1$ and $\alpha = \lambda/pq$, confirming the validity of the latter since it solves both (17) and (18).

In order to obtain the value of the constant c, we use the fact that the limit probabilities π_k of every state $k = 0, 1, 2, \ldots$ in the MC must add up to 1. Formally,

$$1 = \sum_{k=0}^{\infty} \pi_k = \sum_{k=0}^{\infty} c \left(\frac{\lambda}{pq}\right)^k = c \sum_{k=0}^{\infty} \left(\frac{\lambda}{pq}\right)^k = c \, \frac{1}{1 - \frac{\lambda}{pq}}, \tag{20}$$

where we used the fact that $\lambda < pq$ for the convergence of the infinite geometric series. From (20) we obtain $c = 1 - \lambda/pq$ resulting in the limit probabilities

$$\pi_k = c\alpha^k = \left(1 - \frac{\lambda}{pq}\right)\left(\frac{\lambda}{pq}\right)^k. \tag{21}$$

Observe that in (21), as λ tends to pq from below, probabilities π_k tend to 0 for all k, which is consistent with the fact that when $\lambda = pq$ the MC is null recurrent. Similarly, when $\lambda > pq > 0$ and the MC is transient, probabilities $\pi_k = 0$ for all k (this follows from the definition of transient state).

C) Probability of empty queue and probability of minimal wait. Now that we know the limit distribution of the MC, we can compute some performance metrics for the long run operation of the system. For example, a first performance metric is the probability of the queue of an arbitrary terminal j being empty after a large number of time slots n. This probability converges to π_0 for large n,

$$\mathrm{P}\left[\text{Empty queue}\right] = \pi_0 = 1 - \frac{\lambda}{pq}. \tag{22}$$

Another related important performance metric is the probability T_1 a packet being transmitted (either successfully or not) in the first slot after arrival. For this to occur, two independent events must take place: the queue must be empty upon arrival of the packet and the terminal must transmit in the next time slot. Consequently, T_1 may be computed as

$$\begin{aligned} T_1 &= \mathrm{P}\left[\text{Empty queue}\right] \mathrm{P}\left[\text{Transmission in next time slot}\right] \\ &= \pi_0 p = \left(1 - \frac{\lambda}{pq}\right) p = p - \frac{\lambda}{q}. \end{aligned} \tag{23}$$

Also important is the probability S_1 of the packet being successfully communicated in the first slot after arrival. Using the same argument behind (23) we obtain that

$$\begin{aligned} S_1 &= \mathrm{P}\left[\text{Empty queue}\right] \mathrm{P}\left[\text{Successful transmission in next time slot}\right] \\ &= \pi_0 pq = \left(1 - \frac{\lambda}{pq}\right) pq = pq - \lambda. \end{aligned} \tag{24}$$

D) Expected queue length. Yet another performance metric is the expected queue length $\mathbb{E}\left[Q_{jn}\right]$ when the system is operating in steady state. Since we are interested in the steady-state behavior of the system, we will rely on the limit probabilities π_k to compute the expected queue length. Formally, we want to find

$$\lim_{n\to\infty} \mathbb{E}\left[Q_{jn}\right], \quad j = 1, \ldots, J. \tag{25}$$

From the definition of expected value, we obtain

$$\lim_{n\to\infty} \mathbb{E}[Q_{jn}] = \lim_{n\to\infty} \sum_{k=0}^{\infty} \mathrm{P}\left[Q_{jn} = k\right] k = \sum_{k=0}^{\infty} \lim_{n\to\infty} \mathrm{P}\left[Q_{jn} = k\right] k, \tag{26}$$

where we exchanged the summation and the limit in the last equality. We can do this exchange since we know that $\lim_{n\to\infty} \mathrm{P}\left[Q_{jn} = k\right]$ exist for all k. In fact, $\lim_{n\to\infty} \mathrm{P}\left[Q_{jn} = k\right] = \pi_k$ by definition, so (26) can be written as

$$\lim_{n\to\infty} \mathbb{E}[Q_{jn}] = \sum_{k=0}^{\infty} \pi_k k = \left(1 - \frac{\lambda}{pq}\right) \sum_{k=0}^{\infty} \left(\frac{\lambda}{pq}\right)^k k. \tag{27}$$

To solve the infinite series in (27), we start with the well-known geometric series

$$\sum_{k=0}^{\infty} x^k = \frac{1}{1-x}, \tag{28}$$

which is true for $|x| < 1$ and that we used in (20). After differentiating both sides of (28) and multiplying both of them by x, we obtain

$$\sum_{k=0}^{\infty} kx^k = \frac{x}{(1-x)^2}. \tag{29}$$

By using (29) for $x = \lambda/pq$, we can evaluate the expected value expression in (27), to obtain

$$\lim_{n\to\infty} \mathbb{E}[Q_{jn}] = \left(1 - \frac{\lambda}{pq}\right) \frac{\lambda/pq}{\left(1 - \frac{\lambda}{pq}\right)^2} = \frac{\lambda/pq}{\left(1 - \frac{\lambda}{pq}\right)} = \frac{\lambda}{(pq - \lambda)}. \tag{30}$$

Notice that the expression obtained in (30) for the expected queue length is quite intuitive. The expected length grows with increasing arrival rate λ and decreases when the difference between the successful transmission rate pq and the arrival rate λ increases.

E) Probability of successful transmission and optimal transmission probability p. So far, we have assumed that the probability q that a transmission by an arbitrary terminal j does not experience a collision with any other terminal is given. However, under the dominant system assumption (B), for a transmitted packet not to collide it must be that none of the remaining $J-1$ transmitted any packet. Since every terminal acts independently, we may write

$$q = (1-p)^{J-1}. \tag{31}$$

This allows us to compute the probability p that maximizes the probability of successful transmission pq. In order to do this, we differentiate pq with respect to p and look for the roots of the corresponding equation, i.e.

$$\frac{d}{dp}[pq] = \frac{d}{dp}\left[p(1-p)^{J-1}\right] = (1-p)^{J-1} - (J-1)p(1-p)^{J-2} = 0. \tag{32}$$

Since we obviously assume $p < 1$ because otherwise the probability of successful transmission is trivially null, we may divide (32) by $(1-p)^{J-2}$ to obtain a linear expression that yields the optimal probability

$$p* = \frac{1}{J}. \tag{33}$$

Since the probability of the queue being empty is an increasing function of pq [cf. (22)] and the expected queue length is a decreasing function of pq [cf. (30)], the optimal probability p^* in (33) entails simultaneously shorter queues and higher probabilities of these queues being empty. Moreover, for this p^*, the probability of no collision becomes [cf. (31)]

$$q^* = \left(1 - \frac{1}{J}\right)^{J-1}. \tag{34}$$

If we consider a system with many terminals, we may estimate the probability of no collision as the limit of q^* when J tends to infinity, i.e.,

$$\lim_{J\to\infty} q^* = \lim_{J\to\infty}\left(1 - \frac{1}{J}\right)^{J-1} = \lim_{J\to\infty}\left(1 - \frac{1}{J}\right)^{J}\left(1 - \frac{1}{J}\right) = \frac{1}{e} \approx 0.368. \tag{35}$$

The above result implies that RA communications utilizes approximately 37% of the available access point resources without any coordination overhead among terminals.

F) Average time occupancies. It is possible to argue that the limit probabilities in (21) as well as the performance indicators in (22)-(24), and (27) 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 stochastic 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

$$\bar{p}_k := \lim_{n\to\infty} \frac{1}{n}\sum_{m=1}^{n} \mathbb{I}\{Q_{jm} = k\}, \qquad \text{for all } k, \tag{36}$$

which tells you the fraction of time there were k packets awaiting transmission in the queue of terminal j. Different from the π_k in (21) the $\bar{p}_k$ in (36) are a performance metric associated with each particular experiment. However, given that we are in the ergodic parameter configuration ($pq > \lambda > 0$), by the ergodic theorem we can assure that $\bar{p}_k = \pi_k$ for all f, i.e., the limit as $n \to \infty$ of the fraction of time there were k packets in the queue is just the limit probability k.

Similarly, we can find an analogous way of expressing the expected value of the queue length for any terminal j, namely

$$\lim_{n\to\infty} \mathbb{E}\left[Q_{jn}\right] = \lim_{n\to\infty} \frac{1}{n} \sum_{m=1}^{n} Q_{jm} = \sum_{k=1}^{\infty} k\, \pi_k. \tag{37}$$

G) System simulation. The Matlab function to implement a simulation of the system follows. Notice that enforcing the no-concurrence assumption is quite difficult. Instead, we assume arrivals have precedence over services.

```
function [x] = aloha_uplink_simulation(J,p,lambda,N)
    % Implementing no concurrence hypothesis is very difficult.
    % Instead, we assume arrivals have precedence over service.
    x=zeros(J,N);
    for t=1:N-1
            arrivals=binornd(1,lambda,J,1);
            is_there_packets=x(:,t)>0;
            decide_to_transmit=binornd(1,p,J,1);
            service = is_there_packets & decide_to_transmit &  ~arrivals;
            % service = is_there_packets & decide_to_transmit;
            % Uncomment the line above to relax the non-concurrence assump.
            if sum(service)<=1
                x(:,t+1)=x(:,t)+arrivals-service;
            else
                x(:,t+1)=x(:,t)+arrivals;
            end
    end
end
```

Next, we run the function aloha_uplink_simulation for $J = 16$ terminals, optimal transmission probability $p = 1/J$ as computed in Part *E*, $\lambda = 0.9pq$ and $N = 10^4$. A graph with the path followed by terminals 1 through 4 is shown in Fig. 4, and the results obtained using the following Matlab script.

```
clear all; close all; clc;

J=16;
p=1/J;
N=10^5;
lambda=0.9*p*(1-p)^(J-1);

x = aloha_uplink_simulation(J,p,lambda,N);

figure
stairs(1:1000,x(1:4,1:1000)','LineWidth',2);
title('Evolution of queues 1..4 over the first 1000 time slots','FontSize',12)
xlabel('time','FontSize',12)
ylabel('packets in queues 1..4','FontSize',12)
legend('queue 1','queue 2','queue 3','queue 4','Location','Best')
```

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

$$\xi_k := \lim_{n\to\infty} \mathrm{P}\left(R_{jn} = k\right), \quad \text{for all } k, \tag{38}$$

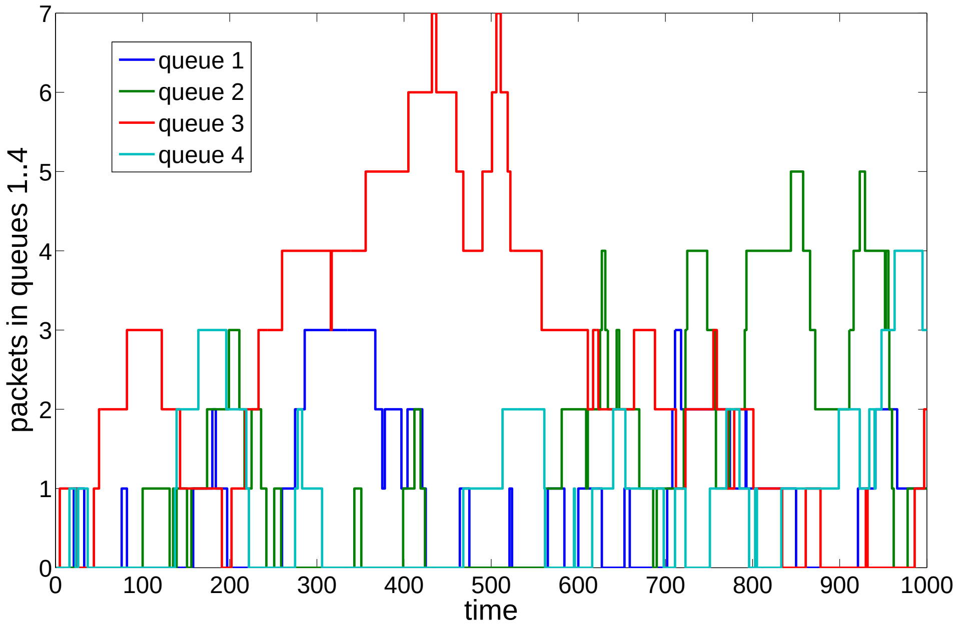

Fig. 4. Evolution of the first four queues over the first 1000 time slots. Queue 3 achieves a maximum queue length of seven packets. Every other queue length remained under this value.

these probabilities cannot be computed in closed form. However, we may use our simulation results to estimate the probability distribution function in (38). Moreover, since the MC is ergodic, we may compute the probabilities in (38) as the time averages over a single simulation run, that is

$$\xi_k \approx \frac{1}{N} \sum_{n=1}^{N} \mathbb{I}\{R_{jn} = k\}, \qquad \text{for all } k, \tag{39}$$

and for large N. The Matlab script to generate the requested plot follows:

```
clear all; close all; clc;

J=16;
p=1/J;
N=10^5;
lambda=0.9*p*(1-p)^(J-1);

x = aloha_uplink_simulation(J,p,lambda,N);

Q1=max(x(1,:));
frequencies=zeros(1,Q1+1);
for i=0:Q1
    frequencies(1,i+1)=sum(x(1,:)==i);
end

rho=lambda/(p*(1-p)^(J-1));

A=[frequencies/N;(1-rho)*(rho.^(0:Q1))];
A=A';
figure
```

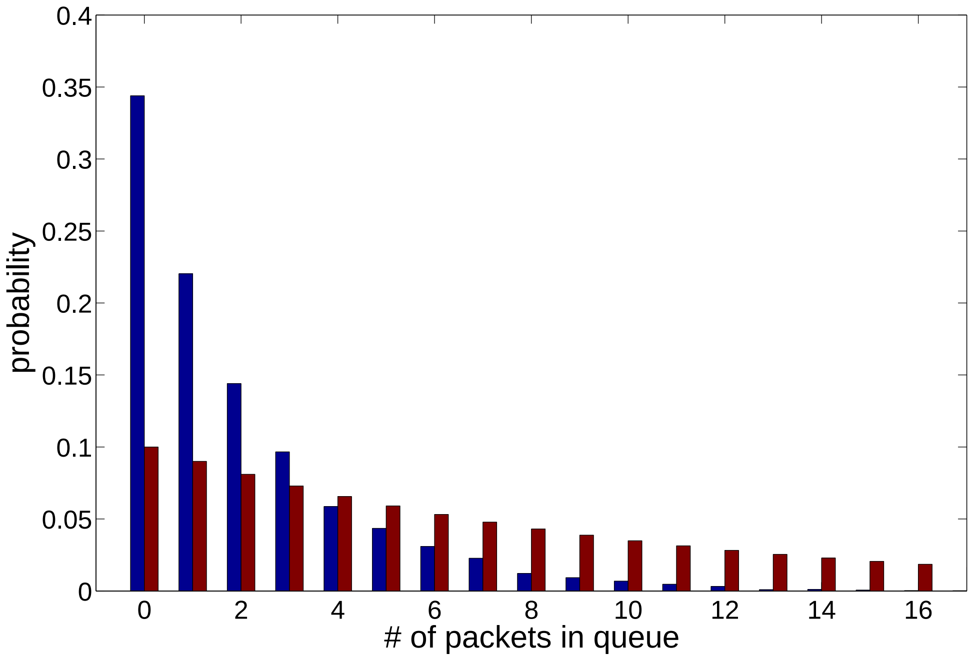

Fig. 5. Comparison of ξ_k (blue) and π_k (red). Probabilities π_k, under the dominant system assumption, underestimate the RA system performance.

```
bar(0:Q1,A, 1);
axis([-1 Q1 0 0.4])
xlabel('# of packets in queue','FontSize',12)
ylabel('calculated probability','FontSize',12)
```

In Figure 5, we plot our estimate of ξ_k and compare it with the π_k in (21) found theoretically under the dominant system assumption. From the figure, it is immediate that the dominant system assumption considerably underestimates the performance of the RA policy. For example, the probability of a queue being empty under the dominant system hypothesis is 0.1 whereas the probability when relaxing this assumption is close to 0.35. The performance difference is attributable to the detrimental effect of dummy packets. Without the dominant system assumption, empty queues remain silent, increasing the probability of successful transmission of other queues by avoiding unnecessary collisions and positively impacting the system performance.