Week 5: Markov chains
Random access in communication networks
Solutions

A Markov chain model. The model described in the homework defines the following probabilities:

$$\begin{aligned}\mathbb{P}\left[\text{a terminal receives a packet in time slot } n\right] &= \lambda, &&\text{for all } n\\ \mathbb{P}\left[\text{a terminal successfully transmits a packet in time slot } n\right] &= pq, &&\text{for all } n \text{ such that } Q_{jn} > 0\end{aligned}$$

The transition probabilities for the number packets in queue are explained below:

$$\mathbb{P}\left[Q_{j(n+1)} = k+1 \mid Q_{jn} = k\right] = \mathbb{P}\left[\text{terminal receives a packet in time slot } n\right] = \lambda,$$

which holds for all n,

$$\begin{aligned}\mathbb{P}\left[Q_{j(n+1)} = k-1 \mid Q_{jn} = k\right] &= \mathbb{P}\left[\text{terminal successfully transmits a packet in time slot } n\right] = pq,\\ \mathbb{P}\left[Q_{j(n+1)} = k \mid Q_{jn} = k\right] &= \mathbb{P}\left[\text{neither receives nor transmits a packet in time slot } n\right]\\ &= 1 - \mathbb{P}\left[\text{receiving a packet}\right] - \mathbb{P}\left[\text{successfully transmitting a packet}\right]\\ &= 1 - \lambda - pq,\end{aligned}$$

which holds for $k > 0$ since if there is no packet in the queue, then no packet can be transmitted; for $k = 0$ we have

$$\mathbb{P}\left[Q_{j(n+1)} = 0 \mid Q_{jn} = 0\right] = \mathbb{P}\left[\text{terminal doesn't receive a packet in time slot } n\right] = 1 - \lambda.$$

Notice from Figure 1 that this MC is similar to a bounded random walk on the integer line, which you have seen can be recurrent or transient depending on the values of the transition probabilities. We discuss the MC obtained for different values of λ and pq in Table 1.

As far as ergodicity, recall that for a MC to be ergodic it must be irreducible, positive recurrent, and aperiodic. Thus, this MC is only ergodic when $pq > \lambda > 0$.

B Limit distribution. The limit distribution depends on the values of λ and pq. Indeed, recall that only positive recurrent MCs have well-defined limit distribution. Moreover, only ergodic MCs have limit distributions that do not depend on the initial state. In our case, if $\lambda > pq$ the MC consists of one transient class. In this case, the limit distribution does not exist (or if you want, it is not a proper probability distribution since $\pi_k = 0$ for all k). When $\lambda = pq$, the MC is null recurrent and again π_k is null for all k.

In contrast, the MC is positive recurrent for $\lambda < pq$. Not only that, but it is ergodic! In this case, we can calculate what the limit distribution is. Actually, we only need to check that it has

Table 1: MC type depending on the values of λ and pq

Parameters	Type of MC
$\lambda = 0$, $pq = 0$	Every state is a class, every class (state) is (positive) recurrent.
$\lambda = 0$, $pq > 0$	Every state is a class and the only (positive) recurrent class (state) is 0. Every other class (state) is transient.
$\lambda > 0$, $pq = 0$	Every state is a class and every class (state) is transient.
$\lambda > pq > 0$	MC has a SINGLE class and it is transient (there is a positive *drift* toward infinity).
$pq > \lambda > 0$	MC has a SINGLE class and it is (positive) recurrent (there is a negative *drift* toward zero).
$pq = \lambda > 0$	MC has a SINGLE class and it is (null) recurrent (there is no drift: the MC is guaranteed to return to any state, but the expected return time is infinite).

the form proposed in the exercise. To do so, note that any equilibrium distribution π_k must satisfy two equilibrium equations: one for when the queue is empty and another for when it is not. When the queue is empty, i.e., when we are at state 0, we have

$$\begin{aligned}\pi_0 &= \mathbb{P}\left[\text{no new packet is delivered}\right] \times \pi_0 + \mathbb{P}\left[\text{a packet is successfully transmitted}\right] \times \pi_1 \\ &= (1-\lambda)\pi_0 + pq\pi_1.\end{aligned} \tag{1}$$

For an arbitrary state $k > 0$, we get

$$\begin{aligned}\pi_k &= \mathbb{P}\left[\text{new packet is delivered}\right] \times \pi_{k-1} \\ &\quad + \mathbb{P}\left[\text{no new packet is delivered or succesfully transmitted}\right] \times \pi_k \\ &\quad + \mathbb{P}\left[\text{a packet is successfully transmitted}\right] \times \pi_{k+1} \\ &= \lambda\pi_{k-1} + (1-\lambda-pq)\pi_k + (pq)\pi_{k+1}\end{aligned} \tag{2}$$

Let's now postulate that $\pi_k = c\alpha^k$, check whether that is correct, and evaluate the constants. From (1) we get

$$c\alpha^0 = (1-\lambda)c\alpha^0 + (pq)c\alpha^1 \Leftrightarrow c = (1-\lambda)c + (pq)c\alpha \Leftrightarrow \alpha = \frac{\lambda}{pq}.$$

Although we obtained an α that works for $k = 0$, we must check that it also works for an arbitrary $k > 0$. To do so, we substitute α into (2) to get

$$c\left(\frac{\lambda}{pq}\right)^k = \lambda c\left(\frac{\lambda}{pq}\right)^{k-1} + (1-\lambda-pq)c\left(\frac{\lambda}{pq}\right)^k + (pq)c\left(\frac{\lambda}{pq}\right)^{k+1},$$

which simplifies to

$$\frac{\lambda}{pq} = \lambda + (1-\lambda-pq)\frac{\lambda}{pq} + (pq)\left(\frac{\lambda}{pq}\right)^2 \Rightarrow \frac{\lambda}{pq} = \frac{\lambda}{pq}.$$

We have therefore established that $\pi_k = c(\lambda/pq)^k$.

To solve for the constant c, recall that any probability distribution must add up to one. Explicitly,

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

Note that this holds because $pq > \lambda \Rightarrow \lambda/pq < 1$, so that the geometric series is convergent. Thus, in this case,

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

C Probability of empty queue and probability of minimal wait. The probability of the j-th queue being empty for large n is equal to the limit distribution π_0. This is true for the case $pq > \lambda$. Otherwise, the limit distribution is improper (although we could say that the probability of the queue being empty is then zero). Thus, from the previous exercise we obtain

$$\mathbb{P}\left[\text{Queue becomes empty}\right] = \pi_0 = \left(1-\frac{\lambda}{pq}\right), \quad \text{when } pq > \lambda.$$

For a packet to be transmitted in the first slot after it arrives, the queue had to be empty when the packet arrived—otherwise, another packet would have precedence. This occurs with probability π_0 for large n. Furthermore, the terminal must "decide" to transmit the packet, which occurs with probability p. Hence,

$$T_1 = \mathbb{P}\left[\text{queue was empty when packet arrived}\right]\mathbb{P}\left[\text{transmitting packet}\right] = p\pi_0 = p - \frac{\lambda}{q}.$$

Following the same logic, the probability of the packet being successful transmitted in the first slot after arrival is

$$\begin{aligned} S_1 &= \mathbb{P}\left[\text{queue was empty when packet arrived}\right]\mathbb{P}\left[\text{transmitting packet}\right]\mathbb{P}\left[\text{successful transmission}\right] \\ &= (pq)\pi_0 = pq - \lambda. \end{aligned}$$

D Expected queue length. Recall that the expected value is defined as the sum of all possible queue lengths k times their respective probabilities. Explicitly

$$\bar{Q}_j = \lim_{n\to\infty}\mathbb{E}[Q_{jn}] = \lim_{n\to\infty}\left(\sum_{k=0}^{\infty} k\mathbb{P}\left[Q_{jn}=k\right]\right) = \sum_{k=0}^{\infty}\left(\lim_{n\to\infty} k\mathbb{P}\left[Q_{jn}=k\right]\right).$$

Observe that the second equality is not always true, i.e., you are not always allowed to reverse limits and summations. In fact, you typically cannot do that! In this case, however, both the limit of the series and the sum of the limits are well-defined, so our manipulation is valid.

Using the fact that $\pi_k = \lim_{n\to\infty}\mathbb{P}\left[Q_{jn}=k\right]$ (by definition), we get

$$\bar{Q}_j = \sum_{k=0}^{\infty} k\pi_k = \sum_{k=0}^{\infty}\left[k\left(1-\frac{\lambda}{pq}\right)\frac{\lambda}{pq}^k\right] = \left(1-\frac{\lambda}{pq}\right)\sum_{k=0}^{\infty} k\frac{\lambda}{pq}^k. \tag{3}$$

To solve the series above, we will use the classical trick of differentiating the power series. Start by recalling that

$$\sum_{n=0}^{\infty} x^n = \frac{1}{1-x}, \quad \text{for } x < 1.$$

Notice, however, that the series we have in (3) is of the form $\sum_{n=0}^{\infty} nx^n$. To get the "n in front", we can use the linearity of the differential operator to write

$$\sum_{n=0}^{\infty} nx^n = x\sum_{n=0}^{\infty} nx^{n-1} = x\sum_{n=0}^{\infty} \frac{\partial x^n}{\partial x} = x \cdot \frac{\partial}{\partial x}\left(\sum_{n=0}^{\infty} x^n\right) = x \cdot \frac{\partial}{\partial x}\left(\frac{1}{1-x}\right) = \frac{x}{(1-x)^2}, \quad \text{for } x < 1.$$

Using this result back in (3) yields

$$\bar{Q}_j = \left(1 - \frac{\lambda}{pq}\right)\sum_{k=0}^{\infty} k\frac{\lambda}{pq}^k = \left(1 - \frac{\lambda}{pq}\right)\frac{\lambda/pq}{\left(1-\frac{\lambda}{pq}\right)^2} = \frac{\lambda/pq}{\left(1-\frac{\lambda}{pq}\right)} = \frac{\lambda}{pq-\lambda}.$$

E Probability of successful transmission and optimal transmission probability $p^\star$. A transmission is successful when a single terminal tries to transmit a packet, since in this case there is "no collision." Explicitly,

$$\mathbb{P}\left[\text{successful transmission}\right] = \mathbb{P}\left[\text{exactly one terminal attempts to transmit a packet}\right] = p\times(1-p)^{J-1}, \tag{4}$$

where we use the hypothesis (B) to ignore the queue length of the terminals and independence of transmission attempts across terminals. We can now check that $p^\star = 1/J$ indeed maximizes the probability of successful transmission by showing that the derivative of (4) vanishes at $p^\star$[1]:

$$\begin{aligned}\frac{d}{dp}\mathbb{P}\left[\text{successful transmission}\right]\Big|_{p=1/J} &= (1-p)^{J-1} + p(1-J)(1-p)^{J-2}\big|_{p=1/J} \\ &= \left(1-\frac{1}{J}\right)^{J-1} - \frac{J-1}{J}\left(1-\frac{1}{J}\right)^{J-2} = 0.\end{aligned}$$

Observe from Part C that the probability of any queue being empty is an *increasing* function of the probability of successful transmission (i.e., pq). Also, note from Part D that the expected queue length is a *decreasing* functions of this probability. Hence, maximizing the probability of successful transmissions is equivalent to maximizing the probability of empty queues and minimizing the expected queue length.

Recall that a successful transmission occurs when (i) a transmission is attempted (which occurs with probability p) AND (ii) no collision occurs (which occurs with probability q). Hence, the probability of a transmitted packet reaching the AP is given by $q = (1-p)^{J-1}$. Therefore, the asymptotic probability of being able to successfully transmit a packet assuming packets are transmitted with probability $p^\star$ is given by

$$\lim_{J\to\infty} q^\star = \lim_{J\to\infty}\left(1-\frac{1}{J}\right)^{J-1} = e^{-1} \approx 0.36.$$

F Average time occupancies. The main insight in this solution is that the MC is ergodic (assuming $\lambda < pq$). Therefore, the fraction of time there were k packets in the queue is equal to the

[1] Technically, this only shows that $p^\star$ is a stationary point of (4). To show that it is a maximum, we should evaluate its second derivative to make sure it is negative.

limit distribution of state k. Formally,

$$\bar{p}_k = \lim_{n\to\infty} \frac{1}{n} \sum_{m=1}^{n} \mathbb{I}\left(Q_{jm} = k\right) = \pi_k.$$

Indeed, for an ergodic MC, the *ensemble average* (the average across realizations of the MC) and the *ergodic average* (the average across a single realization of the MC) converge as $n \to \infty$.

We can find an analogous way to express the expected value of queue length for leveraging the fact that for any ergodic MC $\{X_n\}$ it holds that

$$\lim_{n\to\infty} \frac{1}{n} \sum_{m=1}^{n} f(X_m) = \sum_{\ell=1}^{\infty} f(\ell)\pi_\ell.$$

Thus, we can write the ergodic average of the queue length as

$$\lim_{n\to\infty} \frac{1}{n} \sum_{m=1}^{n} Q_{jm} = \sum_{k=0}^{\infty} k\pi_k = \bar{Q}_j.$$

G System simulation. The MATLAB function for simulating the communication system is given below. It is written in vectorized form to be faster and more concise. If you're interested in understanding how it works, ask your TAs.

```
function [ R ] = alohaSim( J, p, lambda, N )
%ALOHASIM Simulates the ALOHA protocol for HW 5, Part G
%   Inputs:
%   J - Number of terminals
%   p - Probability of a terminal attempting a transmission
%   lambda - Rate of packet arrivals
%   N - Length of simulation

% Initialize terminal queues to zero
R = zeros(J, N);

% Start simulation
for t = 1:N-1
    % Draw whether each terminal produces a new packet to transmit
    arrivals = binornd(1, lambda, [J, 1]);

    % Update queue length with new packets
    R(:,t+1) = R(:,t) + arrivals;

    % Check which terminals have non-empty queues
    non_empty_queue = R(:,t) > 0;

    % Draw whether each terminal would choose transmit
    transmission = binornd(1, p, [J, 1]);

    % For each terminal, evaluate whether or not they transmit a packet by
    % checking if (i) they have not received a packet AND (ii) their queue
    % is not empty AND (iii) they have chosen to transmit in this time slot
    service = (~arrivals) & non_empty_queue & transmission;

    % Check if transmission is successful, i.e., if a single terminal
    % has attempted to transmit a packet...
    if sum(service) == 1
        % The transmission is successful: remove packet from terminal queue
        R(:,t+1) = R(:,t) - service;
```

```
    end
end

end
```

To obtain the desired plot, we execute the function using the following code, whose result is depicted in Fig. 1.

```
% Delete all variables and close figures
clear all
close all

J = 16;                          % Number of users
p = 1/J;                         % Optimal transmission probability (p^\star)
lambda = 0.9 * p*(1-p)^(J-1);    % Packet arrival rate
N = 10^4;                        % Length of the simulation

% Simulation
R = alohaSim( J, p, lambda, N );

% Queue lengths of terminals 1-4
figure();
stairs(1:N, R(1:4,:)', 'LineWidth', 2);
xlabel('Time slot');
ylabel('Queue length');
legend('Terminal 1','Terminal 2','Terminal 3','Terminal 4', 'Location', 'Best')
grid;

%%% Export figure %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
set(gcf,'Color','w');
export_fig -q101 -pdf HW5_G.pdf
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
```

H Compare numerical and analytical results. Because the MC is ergodic, we can use the time averages from a single trial to estimate its limit distribution. We do so using the following MATLAB code:

```
% Delete all variables and close figures
clear all
close all

J = 16;                          % Number of users
p = 1/J;                         % Optimal transmission probability (p^\star)
lambda = 0.9 * p*(1-p)^(J-1);    % Packet arrival rate
N = 10^4;                        % Length of the simulation

% Simulation
R = alohaSim( J, p, lambda, N );

% We use only terminal 1 from now on
R1 = R(1,:);
maxR1 = max(R1);
prob_sim = histcounts(R1, 0:maxR1+1)/N;

% Analytical distribution
rho = lambda/(p*(1-p)^(J-1));
prob_theo = (1-rho) * rho.^(0:maxR1);

% Plots
figure;
stem(0:maxR1, prob_sim, 'LineWidth', 2);
```

```
hold on
stem(0:maxR1, prob_theo, 'x', 'LineWidth', 2);
xlabel('Queue length');
ylabel('Probability');
legend('Estimated probability (\xi_k)', ...
       'Analytical probability (\pi_k)');
grid;


%%% Export figure %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
set(gcf,'Color','w');
export_fig -q101 -pdf HW5_H.pdf
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
```

The result is shown in Fig. 2. Note that the dominant system assumption used to evaluate the π_k considerably underestimates the performance of the RA policy. Indeed, it is both unnecessary and detrimental to transmit dummy packets when we could be silent and increase the probability of successful transmissions by avoiding collisions.

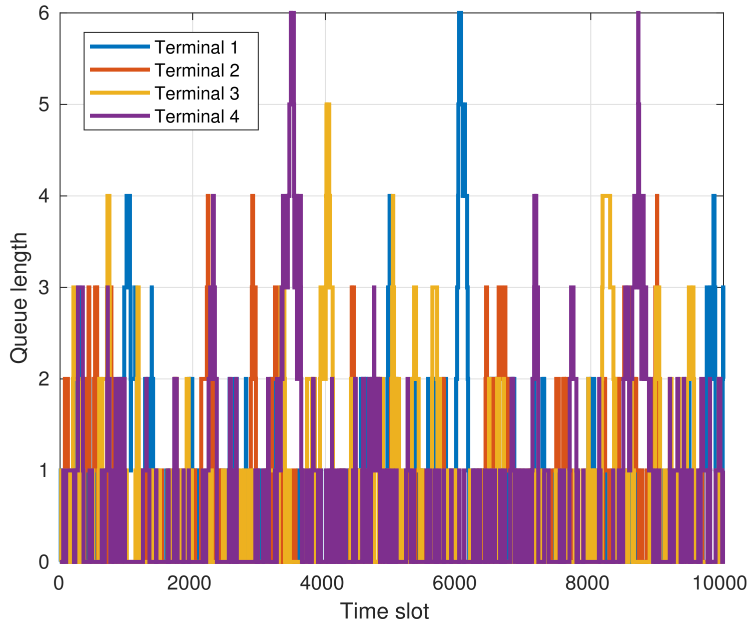

Figure 1: Queue lengths of terminals one through four across $N = 10^4$ time slots (part G).

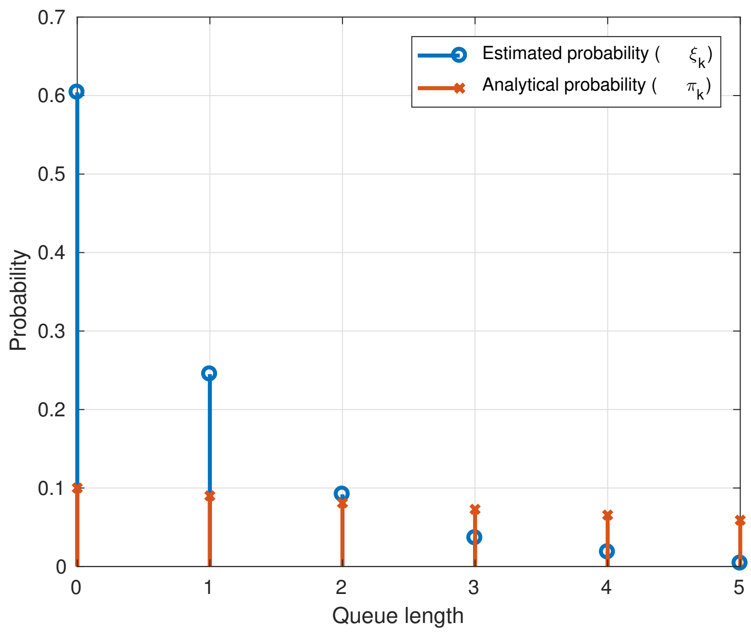

Figure 2: Comparison between the numerical and analytical limit distributions (part H).