Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

HS 2021
Prof. R. Wattenhofer / Z. Meng

Discrete Event Systems

Solution to Exercise Sheet 6

1 Quiz Questions

a) Yes, by using states that consist of the current and the previous value in our sequence of random variables (or, more specifically, by using all possible combinations (current value, previous value) as states). For a concrete example (which is even a bit more complex since only for one state the probabilites for the next state depend also on the previous state) compare the solution to exercise 2c).

b) Yes, since all states have the same period if a Markov chain is irreducible, which yields the required aperiodicity.

c) Using the argument for the previous answer again, the existence of two states with different periods implies that the Markov chain is (always) not irreducible. It immediately follows that the Markov chain is not ergodic. Obviously, the Markov chain is not aperiodic.

d) We only give some intuition here, no rigorous proofs.

Removing $v \to u$ will decrease PageRank(u): "Less PageRank flows to u and more to w instead."

Removing $u \to v$ will decrease PageRank(u): "Less PageRank flows to v which links to u."

Removing $u \to w$ will increase PageRank(u): "More PageRank flows to v instead of w and the former links to u while the latter does not."

Removing $v \to w$ will increase PageRank(u): "More PageRank flows to u instead of w."

e) No, e.g., if you have a graph consisting of two nodes and a connecting edge and you start in one node, then, in each time step, you will switch to the other node and never converge to the stable distribution (which has probability of 0.5 for each of the two nodes).

2 Soccer Betting

a) The following Markov chain models the different transition probabilities (W:Win, T:Tie, L:Loss):

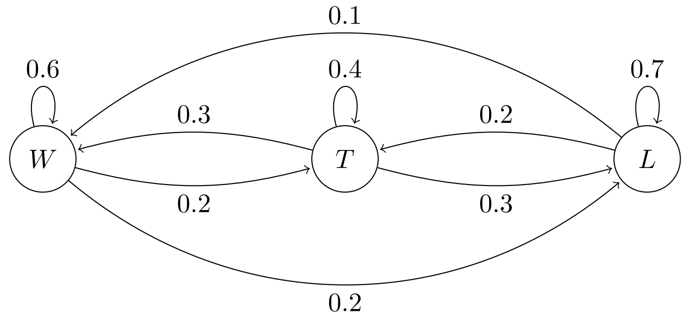

b) The transition matrix P is

$$P = \begin{pmatrix} 0.6 & 0.2 & 0.2 \\ 0.3 & 0.4 & 0.3 \\ 0.1 & 0.2 & 0.7 \end{pmatrix} .$$

As you might have noticed, we gave redundant information here. You only need the information that the FCB lost its last game. Thus, the Markov chain is currently in the state L and hence, the initial vector is $q_0 = \begin{pmatrix} 0 & 0 & 1 \end{pmatrix}$. The probability distribution q_2 for the game against the FC Zurich is therefore given by

$$q_2 = q_0 \cdot P^2 = (q_0 \cdot P) \cdot P = \begin{pmatrix} 0.1 & 0.2 & 0.7 \end{pmatrix} \cdot \begin{pmatrix} 0.6 & 0.2 & 0.2 \\ 0.3 & 0.4 & 0.3 \\ 0.1 & 0.2 & 0.7 \end{pmatrix}$$
$$= \begin{pmatrix} 0.19 & 0.24 & 0.57 \end{pmatrix} .$$

(Note that q_0 must be a row vector, not a column vector.)

Hint: We exploited the associativity of the matrix multiplication to avoid having to calculate P^2 explicitly. This is usually a good "trick" to avoid extensive and error-prone calculations if no calculator is at hand (as for example in an exam situation ☺).

Given the quotas of the exercise, the expected return for each of the three possibilities (W, T, L) calculates as follows.

$$\begin{aligned} \mathbf{E}[W] &= 0.19 \cdot 3.5 = 0.665 \\ \mathbf{E}[T] &= 0.24 \cdot 4 \quad = 0.96 \\ \mathbf{E}[L] &= 0.57 \cdot 1.5 = 0.855 \end{aligned}$$

Therefore, the best choice is not to bet at all since the expected return is smaller than 1 for every choice. If a "sales representative" of the Swiss gambling mafia were to force you to bet, you would be best off with betting on a tie, though.

c) The new Markov chain model looks like this. In addition to the three states W, T, and L, there is now a new state LL which is reached if the team has lost twice in a row.

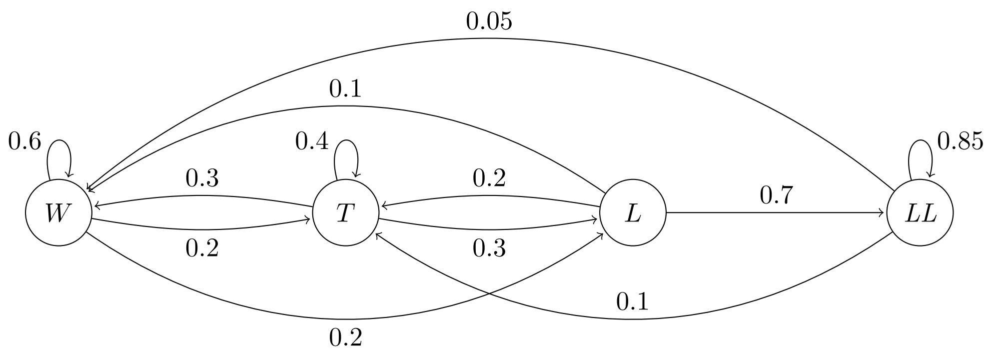

The new transition matrix P is

$$P = \begin{pmatrix} 0.6 & 0.2 & 0.2 & 0 \\ 0.3 & 0.4 & 0.3 & 0 \\ 0.1 & 0.2 & 0 & 0.7 \\ 0.05 & 0.1 & 0 & 0.85 \end{pmatrix} . \quad (1)$$

As the FCB lost its last two games, the Markov chain is currently in the state $q_0 = \begin{pmatrix} 0 & 0 & 0 & 1 \end{pmatrix}$. The probabilities for the game against the FC Zurich can again be computed as follows.

$$q_3 = q_0 \cdot P^2 = (q_0 \cdot P) \cdot P = \begin{pmatrix} 0.05 & 0.1 & 0 & 0.85 \end{pmatrix} \cdot \begin{pmatrix} 0.6 & 0.2 & 0.2 & 0 \\ 0.3 & 0.4 & 0.3 & 0 \\ 0.1 & 0.2 & 0 & 0.7 \\ 0.05 & 0.1 & 0 & 0.85 \end{pmatrix}$$

$$= \begin{pmatrix} 0.1025 & 0.135 & 0.04 & 0.7225 \end{pmatrix}$$

Finally, we can compute the expected profit for each of the three possible bets:

$$\begin{aligned} \mathbf{E}[W] &= 0.1025 \cdot 3.5 &&= 0.35875 \\ \mathbf{E}[T] &= 0.135 \cdot 4 &&= 0.54 \\ \mathbf{E}[L] &= (0.04 + 0.7225) \cdot 1.5 &&= 1.14375 \ . \end{aligned}$$

Now, the best choice is to bet on a loss. Clearly, the addition of the state LL worsens the situation for FCB.

3 PageRank

a) With $v = (1,1,1,1) \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 \end{pmatrix}$ we get a PageRank vector v of $(1,2,2,0)$. We see that website v_4 is quite important for the PageRanks of the other websites, even though it is just a collection of links! So its PageRank should reflect this fact somehow, which it does not at the moment ...

b) We first calculate $d_1 = 1, d_2 = 1, d_3 = 0, d_4 = 3$ and now get a PageRank vector of

$$v = (1,1,1,1) \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \end{pmatrix} = (\frac{1}{3}, \frac{4}{3}, \frac{4}{3}, 0)$$

We now reduced the importance of the link-collection v_4, but still, we do not account for the fact that, according to the calculated PageRank, v_4 is not important at all—nobody recommends it!

c) The results of your iterations should look like this:

(i) $(\frac{1}{3}, \frac{4}{3}, \frac{4}{3}, 0)$

(ii) $(0, \frac{1}{3}, \frac{4}{3}, 0)$

(iii) $(0, 0, \frac{1}{3}, 0)$

(iv) $(0, 0, 0, 0)$

(v) $(0, 0, 0, 0)$

As you can see, when we cannot get to every state from every other state, we run into problems. The $(1 - \alpha) \cdot R$ component in the Google matrix M ensures that M is ergodic, and thus, due to the Ergodic Markov Chain theorem, the process would converge towards the unique stationary distribution. However, continuing these calculations including R might get a bit tedious if you are not a computer—so let us look at the other exercises now :-)

4 Chess Tours

a) Let G be an undirected graph with 64 vertices and m edges that represents knight moves. As per Lemma 3.24 and theorem 3.15, for a simple random walk:

$$h_{u,u} = \frac{2m}{deg(u)}$$

$2m$ can be calculated using $hint_1$ and $hint_2$. Below, we see the degrees of each node in the chess board graph.

2	3	4	4	4	4	3	2
3	4	6	6	6	6	4	3
4	6	8	8	8	8	6	4
4	6	8	8	8	8	6	4
4	6	8	8	8	8	6	4
4	6	8	8	8	8	6	4
3	4	6	6	6	6	4	3
2	3	4	4	4	4	3	2

Using these degrees and the "degree sum theorem", we have:

$$2m = 16 \cdot 8 + 16 \cdot 6 + 20 \cdot 4 + 8 \cdot 3 + 4 \cdot 2 = 336$$

The middle squares of the chess board have a degree of 8. So, the number of moves on average it will take the knight from starting at one of these squares and coming back there is:

$$\frac{1}{\pi_{center\_square}} = \frac{336}{8} = 42$$

b) The degree of a corner square is 2. So, the answer changes to $\frac{336}{2}$, which is 168.

c) Let the 2X2 chess board have the following 4 squares a, b, c, d, with a being the bottom left corner and d being the top right corner. The following Markov chain models the different transition probabilities:

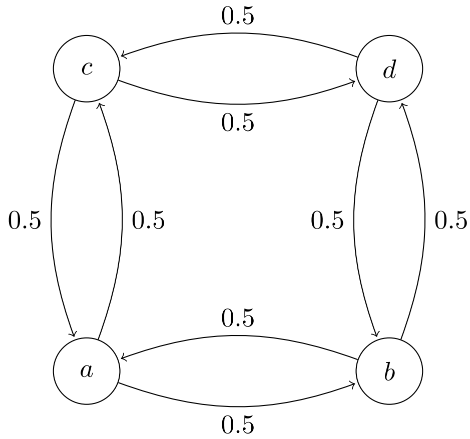

As per Lemma 4.8, expected hitting times are given by:

$$h_{i,j} = 1 + \sum_{k \neq j} p_{i,k} h_{k,j}$$

We have:

$$\begin{aligned} h_{a,d} &= 1 + 0.5 \cdot h_{c,d} + 0.5 \cdot h_{b,d} \\ h_{c,d} &= 1 + 0.5 \cdot h_{a,d} \\ h_{b,d} &= 1 + 0.5 \cdot h_{a,d} \end{aligned}$$

Plugging in $h_{c,d}$ and $h_{b,d}$ in the equation for $h_{a,d}$, and solving for $h_{a,d}$, we get $h_{a,d} = 4$.

5 Queues

Before describing our solutions to the different exercises, we note that the respective systems of linear equations for determining the stationary distributions can of course also be solved in the standard way you learned in your classes. However, sometimes a look into the structure of the respective problem (or into the structure of the respective system of linear equations if you will) can help to solve it faster and somewhat more elegantly, as is the case here.

a) The Markov chain below models the load of the queue:

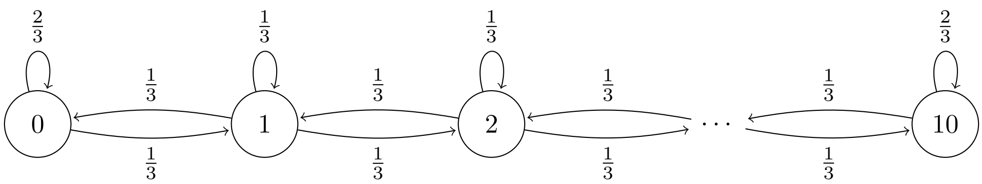

Every loop has a probability of 1/3 except for the first and the last one which have a probability of 2/3 each. Regarding the stationary distribution $\pi = (\pi_0, \pi_1, \ldots, \pi_{10})$, we observe that $\pi_0 = 2/3 \cdot \pi_0 + 1/3 \cdot \pi_1$ which implies $\pi_1 = \pi_0$. By similar arguments, we obtain $\pi_2 = \pi_1, \pi_3 = \pi_2, \ldots, \pi_{10} = \pi_9$. Since the sum of all the π_i is 1, the stationary distribution is $\pi = (1/11, 1/11, \ldots, 1/11)$.

b) The Markov chain now looks like this:

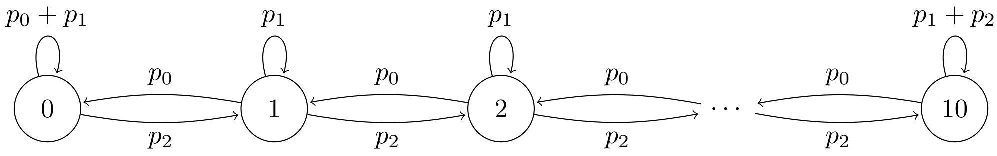

Similarly to the solution to the last question, we can calculate π_1 (and then $\pi_2, \pi_3, \dots$) depending on π_0 (respectively $\pi_1, \pi_2, \dots$). Denoting the probabilites for the arrival of 0, 1 and 2 packets by p_0, p_1 and p_2, respectively, we obtain $\pi_0 = (p_0 + p_1)\pi_0 + p_0\pi_1$ which yields $\pi_1 = p_2/p_0 \cdot \pi_0$ where we use that $p_0 + p_1 + p_2 = 1$. Analogously, we get $\pi_2 = p_2/p_0 \cdot \pi_1$, etc., thus, in general, we have $\pi_{i+1} = k\pi_i$ for all $0 \leq i \leq 9$. Hence, $\pi_i = k^i \pi_0$ for all $0 \leq i \leq 10$ and since the π_i sum up to 1, we obtain

$$\pi_0 = \frac{1}{\sum_{i=0}^{10} k^i} = \frac{k-1}{k^{11}-1} \ ,$$

except if $k = 1$ for which we obtain $\pi_0 = 1/11$. Since all π_i can be calculated from π_0, it is indeed enough to know the factor k between p_0 and p_2 in order to calculate the stationary distribution. In the stationary distribution, the probabilities of any two subsequent nodes in the Markov chain differ exactly by a factor of k.

c) Using the same reasoning as in the last two solutions, we obtain $\pi_0 = \pi_1 = \pi_2 = \dots$. The only choice for π_0 where the (infinite) sum of the p_i does not go to infinity is $\pi_0 = 0$, but in this case the obtained π, while stationary, is not a distribution since the π_i do not sum up to 1. However, there does not necessarily have to be a problem with the stationary distribution in the infinite case: With the arrival probabilities of 2/5, 2/5 and 1/5, we can use the approach in the solution to b) in order to determine the stationary distribution: We obtain

$$k = \frac{p_2}{p_0} = \frac{\frac{1}{5}}{\frac{2}{5}} = \frac{1}{2}$$

which implies

$$\pi_0 = \frac{1}{\sum_{i=0}^{\infty} \left(\frac{1}{2}\right)^i} = \frac{1}{2}$$

and we get

$$\pi_i = \left(\frac{1}{2}\right)^{i+1}$$

in general. As the π_i sum up to 1, this is indeed a stationary distribution.

d) Now the Markov chain looks like depicted below:

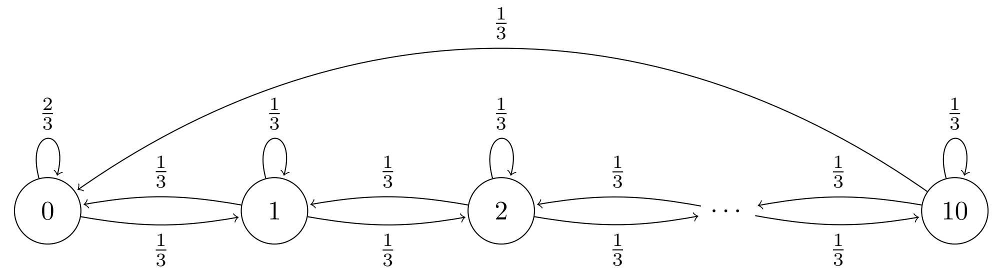

In the previous solutions, we started our reasoning on the left side of the chain since π_0 only depended on (π_0 and) π_1, π_1 depended only on π_0, π_1 and π_2, and so on. As we now have a transition from state 10 to state 0, we start from the other side.

Applying our usual reasoning, we obtain $\pi_9 = 2 \cdot \pi_{10}$. Furthermore we observe that, for each $1 \leq i \leq 9$, it holds that $\pi_i = 1/3 \cdot \pi_{i-1} + 1/3 \cdot \pi_i + 1/3 \cdot \pi_{i+1}$ which implies $\pi_i = 1/2 \cdot \pi_{i-1} + 1/2 \cdot \pi_{i+1}$. In other words, π_i is the arithmetic mean of π_{i-1} and π_{i+1}. We obtain $\pi_i = (11 - i) \cdot \pi_{10}$, for all $0 \leq i \leq 10$. As the π_i sum up to 1, we get $\pi_{10} = 1/66$ which implies $\pi_i = (11 - i)/66$.

e) The Markov chain including retirement looks like this:

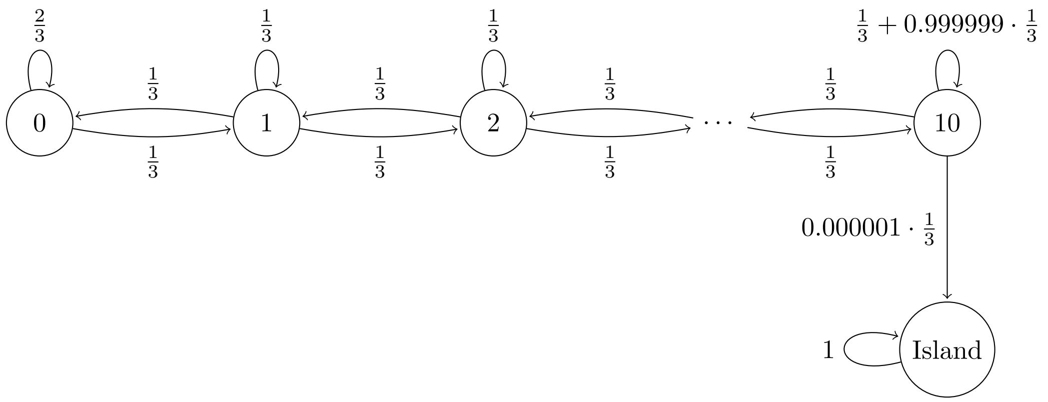

Since the Markov chain is not irreducible anymore, it could be that there is more than one stationary distribution now. However, a closer look reveals that there is only one stationary distribution: Using the argument in the solution to a), we see that $\pi_0 = \pi_1 = \cdots = \pi_{10}$. Moreover, $\pi_{\text{Island}} = \pi_{\text{Island}} + 0.000001 \cdot \pi_{10}$ which implies $\pi_{10} = 0$. Thus, the only stationary distribution is $\pi = (\pi_0, \ldots, \pi_{10}, \pi_{\text{Island}}) = (0, \ldots, 0, 1)$.