Notes on Markov chains, Travel Times, and Opportunistic Routing

Michael J. Neely University of Southern California https://viterbi-web.usc.edu/~mjneely

An updated version of these notes is here: https://viterbi-web.usc.edu/~mjneely/docs/markov-chains-travel-times.pdf

Abstract

This collection of notes provides a brief review of Markov chain theory. It focuses attention on expected travel times and applications to opportunistic routing in wireless networks. The exercises are provided for your own interest.

I. INTRODUCTION TO EXPECTED TRAVEL TIMES



Fig. 1. A simple travel time example with four locations. Travel time between the locations is shown.

We start at home and want to get to school. There are three different paths we can take, depending on where we go first:

- Path one: First go to the coffee shop, then go to school.
- Path two: First go to the store, then go to school.
- Path three: Go directly to school.

Every day we randomly take a path. We want to know the expected time to get to school. To more precisely formulate this problem, represent the four different locations by integers $\{0, 1, 2, 3\}$:

- Location 0: home
- Location 1: coffee shop
- Location 2: store
- Location 3: school

The travel times between each location are shown in the graph of Fig. 1. For simplicity, the travel times are assumed to be constant (not random). For example, from Fig. 1, we see that it takes 8 minutes to go from home to the coffee shop. Assume this 8 minutes includes the time spent ordering coffee. From the coffee shop it takes another 20 minutes to get to school. Thus, the total travel time over path one is 28 minutes. Likewise, total travel time for path two is 20 minutes, and the total travel time for path three is 11 minutes.

Suppose we decide between paths 1, 2, 3 with probabilities p_1, p_2, p_3 , where $p_1 + p_2 + p_3 = 1$ and $p_i \ge 0$ for $i \in \{1, 2, 3\}$. Let T be the random time required to get to school. By the law of total expectation we have:

$$\mathbb{E}[T] = \underbrace{\mathbb{E}[T|\text{go to 1 next}]}_{8+20} p_1 + \underbrace{\mathbb{E}[T|\text{go to 2 next}]}_{15+5} p_2 + \underbrace{\mathbb{E}[T|\text{go to 3 next}]}_{11} p_3$$

For example, if $p_1 = p_2 = p_3 = 1/3$ then $\mathbb{E}[T] = 59/3 \approx 19.667$ minutes.



Fig. 2. An extended example that adds the possibility of a 5 minute self-transition from home to home.

A. An extended problem

That was easy. Let's extend the problem by considering the possibility of deciding to linger at home for 5 minutes with probability p_0 . So, now assume we have probabilities p_0, p_1, p_2, p_3 that are nonnegative and sum to 1. This can be represented by the picture in Fig. 2, where a self-transition from home to home is drawn. Further, assume that every time we self-transition back to home, our memory is somehow erased and we again independently choose between the four options with probabilities p_0, p_1, p_2, p_3 . In particular, we might choose to self-transition to home 32 times before going to the coffee shop, then to school (which would take 32(5) + 8 + 20 minutes). Again let T be the random time required to get to school. By the same idea as before, we have:

$$\mathbb{E}\left[T\right] = \mathbb{E}\left[T|\text{go to 0 next}\right]p_0 + \underbrace{\mathbb{E}\left[T|\text{go to 1 next}\right]}_{8+20}p_1 + \underbrace{\mathbb{E}\left[T|\text{go to 2 next}\right]}_{15+5}p_2 + \underbrace{\mathbb{E}\left[T|\text{go to 3 next}\right]}_{11}p_3$$

The only remaining question is: What is the value of $\mathbb{E}[T|\text{go to } 0 \text{ next}]$? Given that we do a 5-minute self-transition, we know that:

 $\mathbb{E}[T|$ go to 0 next] = 5 + $\mathbb{E}[$ remaining time starting from 0]

Due to our lack of memory, the remaining time starting from 0 is distributed the same as the random time T itself! Thus:

 \mathbb{E} [remaining time starting from 0] = $\mathbb{E}[T]$

Substituting this gives:

$$\mathbb{E}[T] = (5 + \mathbb{E}[T])p_0 + (8 + 20)p_1 + (15 + 5)p_2 + 11p_3$$

The interesting thing is this: We have written an equation for $\mathbb{E}[T]$ in terms of itself. Of course, once we have this equation, there is no remaining challenge. Using basic algebra to solve for $\mathbb{E}[T]$ gives:

$$\mathbb{E}[T] = \frac{5p_0 + 28p_1 + 20p_2 + 11p_3}{1 - p_0}$$

For example, if $p_0 = p_1 = p_2 = p_3 = 1/4$, then $\mathbb{E}[T] = 64/3 \approx 21.333$ minutes. General problems of this type can be solved via the theory of discrete time Markov chains, as described in the remainder of these notes.

II. DISCRETE TIME MARKOV CHAINS

A. Describing a DTMC

Let Z(t) be a discrete time Markov chain (DTMC) defined over time slots $t \in \{0, 1, 2, ...\}$. Recall that a DTMC Z(t) is characterized by:¹

- A state space \mathcal{Z} , specifying the set of all possible states. Thus, $Z(t) \in \mathcal{Z}$ for all $t \in \{0, 1, 2, ...\}$. Here we assume that \mathcal{Z} is a finite or countably infinite set.
- A transition probability matrix $P = (P_{ij})$, where $P_{ij} = Pr[Z(t+1) = j|Z(t) = i]$ for all $i \in \mathbb{Z}$ and $j \in \mathbb{Z}$.

The transition probabilities have the property that, for all $i \in \mathcal{Z}$:

$$\sum_{j\in\mathcal{Z}}P_{ij}=1$$

¹This implicitly assumes the DTMC Z(t) is *time homogeneous*, so that its transition probabilities P_{ij} do not change with time. This is assumed throughout the course unless otherwise stated.

The Markov property of Z(t) guarantees that the future state is conditionally independent of the past, given the current state:

$$Pr[Z(t+1) = j | Z(t) = i, Z(t-1) = z(t-1), \dots, Z(0) = z(0)] = P_i$$

where the above holds for all integers $t \ge 0$, all $i \in \mathbb{Z}$, $j \in \mathbb{Z}$, and all sequences $(z(0), z(1), \ldots, z(t-1)) \in \mathbb{Z}^t$. In order to simulate the DTMC, we must specify either an initial condition Z(0), or a probability distribution on Z(0).

B. Evolution of the probabilities

Suppose that at time slot $t \ge 0$ we know the probabilities Pr[Z(t) = i] for all $i \in \mathbb{Z}$. This can be used to compute the probabilities for slot t + 1. Indeed, by the law of total probability, we have for all $j \in \mathbb{Z}$:

$$Pr[Z(t+1) = j] = \sum_{i \in \mathbb{Z}} Pr[Z(t+1) = j | Z(t) = i] Pr[Z(t) = i]$$

=
$$\sum_{i \in \mathbb{Z}} P_{ij} Pr[Z(t) = i]$$
(1)

where the final equality holds by the Markov chain definition. The equation (1) provides one equation for each $j \in \mathbb{Z}$.

This evolution can be described with vector notation. Let $\vec{\pi}(t) = (\pi_i(t))_{i \in \mathbb{Z}}$ be a row vector that specifies the probability distribution for Z(t), so that for all $i \in \mathbb{Z}$ we have:

$$\pi_i(t) = \Pr[Z(t) = i]$$

The initial probability distribution is given as some probability vector $\vec{\pi}(0)$ (with nonnegative entries that sum to 1). If we want Z(0) to be a certain initial state z_0 (with probability 1), we use a probability vector $\vec{\pi}(0)$ of the form:

$$\vec{\pi}(0) = (0, 0, \dots, 0, 1, 0, \dots, 0)$$

which is a vector with a single "1" in entry z_0 , and with 0s in all other entries.

With this vector notation, the equation (1) is written more concisely by multiplying the row vector $\vec{\pi}(t)$ by the matrix P. That is, for all slots $t \ge 0$:

$$\vec{\pi}(t+1) = \vec{\pi}(t)P \tag{2}$$

Iterating this equation gives:

• $\vec{\pi}(0)$ is given.

•
$$\vec{\pi}(1) = \vec{\pi}(0)P$$
.

- $\vec{\pi}(2) = \vec{\pi}(1)P = \vec{\pi}(0)P^2$.
- $\vec{\pi}(3) = \vec{\pi}(2)P = \vec{\pi}(0)P^3$.
- and on.

Thus, for all slots $t \in \{0, 1, 2, ...\}$ we have:

$$\vec{\pi}(t) = \vec{\pi}(0)P^t \tag{3}$$

Exercise 1. Suppose $\vec{\pi}(0) = \vec{\pi}$, for some probability vector $\vec{\pi}$ that solves the equation $\vec{\pi} = \vec{\pi}P$. This equation is called the stationary equation (also called the global balance equation). Prove that $\vec{\pi}(t) = \vec{\pi}$ for all $t \in \{0, 1, 2, ...\}$. Why do you think the equation $\vec{\pi} = \vec{\pi}P$ is called the "stationary equation"?

Exercise 2. Suppose Z is a finite state space with N states, so P is an $N \times N$ matrix. Recall that multiplying an $N \times N$ matrix by another $N \times N$ matrix requires N^3 multiply-add operations, while multiplying the matrix by a row vector requires N^2 multiply-add operations. The initial probability distribution is $\vec{\pi}(0)$. We want to compute $\vec{\pi}(T)$ for some time T > 0.

a) How many operations are required to compute $\vec{\pi}(T)$ using equation (3)? That is, P^T is computed by T-1 separate multiplications of a square matrix by a square matrix.

b) How many operations are required to compute $\vec{\pi}(T)$ using equation (2)? Under what conditions is this approach simpler? c) Suppose $T = 2^K$ for some positive integer K. Develop a technique for computing $\vec{\pi}(T)$ that, for large K, is faster than the answers in both (a) and (b). Consider your answer for K = 20 and N = 50.

Exercise 3. This exercise is open-ended.

a) Define an N-state Markov chain with a specific $N \times N$ transition probability matrix P (perhaps with N = 5). Use Matlab to plot $\vec{\pi}(t)$ for slots $t \in \{0, 1, ..., T\}$, for some specific initial condition and for some finite integer T that produces a meaningful plot (so that both transient and long term behavior can be seen).

b) Repeat using a different initial condition.

c) Compare to a simulation of the system.

III. EXPECTED TRAVEL TIMES FOR A MARKOV CHAIN

This section considers expected travel times in Markov chains, with applications to opportunistic routing in wireless networks.

A. Computing expected travel times

For simplicity, assume a finite state space $Z = \{1, 2, ..., N\}$ for some positive integer $N \ge 2$. Fix state N as a *destination* state. We are interested in the amount of time required to travel to state N. For each $i \in \{1, ..., N-1\}$, define T_i as the random time required to travel from state i to state N in one particular simulation of the DTMC (assuming initial condition Z(0) = i). Define $T_i = \infty$ if the simulation never reaches state N. Define $T_N = 0$. If Z(0) = i and the DTMC eventually reaches state N, its path from i to N has an integer number of hops. Thus, for all $i \in \{1, 2, ..., N-1\}$, the value T_i is either a positive integer (since it takes at least one slot to travel to the destination) or is equal to ∞ . Fix $i \in \{1, ..., N-1\}$. Then:

$$T_{i} = \begin{cases} 1 + T_{1} & \text{with probability } P_{i1} \\ 1 + \tilde{T}_{2} & \text{with probability } P_{i2} \\ \dots & \dots \\ 1 + \tilde{T}_{N} & \text{with probability } P_{iN} \end{cases}$$

$$\tag{4}$$

where T_j is defined as the random *remaining time* to travel to the destination state N, given that the first hop took us from state i to state j. The random variable \tilde{T}_j is only defined when the path from i to the destination visits state j on its first hop (which happens with probability P_{ij}). Each random variable \tilde{T}_j has the same probability distribution as T_j . There is a subtlety here. Why do we introduce additional notation \tilde{T}_j , instead of just using T_j itself? The reason is that the random variable T_j was defined from a particular instance of a simulation that starts in state j and ends in state N. This particular instance might be different than the simulation associated with \tilde{T}_j . Indeed, the simulation for \tilde{T}_j is a portion of the same simulation that defines T_i . This distinction is particularly important if we travel from state i back to state i on the first hop (which happens with probability P_{ii}). In this case, the correct equation is $T_i = 1 + \tilde{T}_i$ (with probability P_{ii}). In particular, the conditional distribution of T_i , given that the first hop travels back to state i, is different from the *a*-priori distribution of T_i . Without the \tilde{T}_i notation, one might be tempted to write the confusing (and incorrect) equation $T_i = 1 + T_i$ (with probability P_{ii}).

Given that \tilde{T}_j exists, it has the same distribution as the a-priori distribution of T_j . Thus, we can write $\mathbb{E}\left[\tilde{T}_j\right] = \mathbb{E}[T_j]$ for all $j \in \{1, \ldots, N\}$, and so from (4) we have (using $\mathbb{E}\left[\tilde{T}_N\right] = \mathbb{E}[T_N] = 0$):

$$\mathbb{E}[T_i] = 1 + \sum_{j=1}^{N-1} P_{ij} \mathbb{E}[T_j] \quad \forall i \in \{1, \dots, N-1\}$$
(5)

The equation (5) says that the expected time to travel from state *i* to the destination is equal to 1 (for the first hop) plus the expected residual time after the first hop. Indeed, one can define $T_i = 1 + R$, where R is the residual time, and so $\mathbb{E}[T_i] = 1 + \mathbb{E}[R]$. The expected residual time can be computed via the law of total expectation:

$$\mathbb{E}[R] = \sum_{j=1}^{N} \mathbb{E}[R|\text{first hop takes us to state } j] P_{ij} = \sum_{j=1}^{N-1} \mathbb{E}[T_j] P_{ij}$$

Adding 1 to the above expression for $\mathbb{E}[R]$ gives an alternative derivation of equation (5).

Note that (5) specifies N-1 linear equations for the N-1 unknowns $\mathbb{E}[T_1], \mathbb{E}[T_2], \ldots, \mathbb{E}[T_{N-1}]$. The equations (5) can be written in vector form. Let $\mathbb{E}\left[\vec{T}\right] = (\mathbb{E}[T_1], \ldots, \mathbb{E}[T_{N-1}])$ be a column vector. The equations become:

$$\mathbb{E}\left[\vec{T}\right] = \vec{1} + \tilde{P}\mathbb{E}\left[\vec{T}\right]$$
(6)

where \tilde{P} is a truncation of P that removes the last row and column (producing an $(N-1) \times (N-1)$ matrix), and where $\vec{1}$ is an $(N-1) \times 1$ column vector of all 1s. Let I be the $(N-1) \times (N-1)$ identity matrix. If $I - \tilde{P}$ is invertible, the unique solution to (6) is:

$$\mathbb{E}\left[\vec{T}\right] = (I - \tilde{P})^{-1}\vec{1}$$

It can be shown that $(I - \tilde{P})$ is invertible whenever the DTMC has a path of nonzero probability from each state $i \in \{1, \ldots, N-1\}$ to the destination state N. Such a path of nonzero probability from each state to the destination occurs if and only if $\lim_{n\to\infty} \tilde{P}^n = (0)$, which occurs if and only if $\mathbb{E}[T_i] < \infty$ for all $i \in \{1, \ldots, N-1\}$ (see Exercise 5).

Exercise 4. Devise a method for computing travel time variances $Var(T_i)$ for $i \in \{1, ..., N-1\}$.

Exercise 6. There are five computers, labeled $\{1, 2, 3, 4, 5\}$. A single file is randomly transferred between computers according to a transition probability matrix (P_{ij}) . It takes time S_{ij} to shift a file from computer i to computer j over link (i, j). The (P_{ij}) and (S_{ij}) matrices are given below:

$$(P_{ij}) = \begin{bmatrix} 0 & 1/3 & 1/3 & 1/3 & 0 \\ 1/2 & 0 & 1/4 & 1/4 & 0 \\ 1/5 & 0 & 0 & 3/5 & 1/5 \\ 1/4 & 1/4 & 1/4 & 0 & 1/4 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} , \quad (S_{ij}) = \begin{bmatrix} 0 & 10 & 3 & 7 & 0 \\ 4 & 0 & 7 & 7 & 0 \\ 8 & 0 & 0 & 8 & 4 \\ 3 & 3 & 3 & 0 & 3 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

For example, starting at computer 1, the file can take a route $1 \rightarrow 3 \rightarrow 1 \rightarrow 4 \rightarrow 5$, and this route has total time 3+8+7+3 and total probability (1/3)(1/5)(1/3)(1/4).

a) For each $i \in \{1, 2, 3, 4, 5\}$, define $\mathbb{E}[T_i]$ as the expected time to reach computer 5, given the file starts in computer *i*. Note that $\mathbb{E}[T_5] = 0$. Write recursive equations for $\mathbb{E}[T_i]$ using $\mathbb{E}[T_i] = \sum_{j=1}^5 \mathbb{E}[T_i|$ next hop to $j] P_{ij}$. b) Compute $\mathbb{E}[T_1]$.

Exercise 7. Consider a 5-state DTMC with transition probability matrix $P = (P_{ij})$ as given in Exercise 6. We start in state 1. a) What is the probability that we visit state 4 before state 5? Hint: For $i \in \{1, 2, 3, 4, 5\}$ define q_i as the probability we visit state 4 before state 5, given we start in state i. Note that $q_4 = 1$ and $q_5 = 0$. Write linear equations for q_1, q_2, q_3 .

b) As a function of the matrix P, what is the probability that after 10 steps we are still not in state 5?

B. Applications to opportunistic routing

Consider a wireless network with N nodes. The network operates in discrete time with slots $t \in \{0, 1, 2, ...\}$. Assume we have a single packet at node 1 that must be transmitted over the network to eventually reach the destination node N. Each transmission from a node i can be overheard by *multiple other nodes*. Let q_{ij} be the probability that a packet transmitted at node i is successfully received at node j. For simplicity, assume the receptions are independent at each node and across time. Thus, if node 1 transmits, the probability that both nodes 2 and 3 successfully receive is $q_{12}q_{13}$. It is convenient to assume that $q_{ii} = 1$ for all $i \in \{1, ..., N\}$. Thus, if a node i transmits and *no other nodes* successfully receive the transmission, then node i must try again on the next slot.

Assume that all nodes provide ACK/NACK feedback at the end of each transmission (absence of an ACK is considered to be a NACK). At the end of each time slot, the transmitting node observes the successful receivers and makes an *opportunistic routing decision* about who takes the packet next. It can either choose itself (so that it tries again on the next slot) or can pass responsibility to one other node that successfully received on that slot. Once responsibility is shifted, the new node becomes the transmitter, makes a transmission, and then makes similar decisions. Thus, there is one and only one node responsible for the packet at any given time. The packet completes its travel once it reaches the destination node N.

Exercise 8. This exercise is open-ended. There are several types of questions that can be asked in this framework, such as:

a) Suppose the routing decision is made according to the closest-to-destination metric (as in the GeRaF heuristic of [1]). Compute the expected travel time.

b) Give an example network for which the closest-to-destination metric gives infinite average travel time, while another simple algorithm would give finite average travel time.

c) Introduce average energy into the picture.

d) Try to develop an algorithm with optimal average travel time [2][3][4][5] (see also heuristics in [6][7] and optimal treatment of multi-packet flows in [8]).

APPENDIX

Lemma 1. Let A be a square matrix. Let $\{B_n\}_{n=1}^{\infty}$ be a sequence of square matrices (the same size as A) such that $\lim_{n\to\infty} AB_n = I$, where I is the identity matrix. Then A is invertible, B_n has a limit, and $\lim_{n\to\infty} B_n = A^{-1}$.

Proof. Suppose A is not invertible (we reach a contradiction). Then there is a nonzero row vector v such that vA = 0. Thus, for all $n \in \{1, 2, 3, ...\}$:

$$vAB_n = 0B_n = 0$$

Hence, $vAB_n = 0$ for all n. Taking a limit as $n \to \infty$ gives vI = 0, and so v = 0 (a contradiction). Thus, A is invertible. It follows that $B_n = A^{-1}(AB_n)$ for all n. Taking a limit gives $\lim_{n\to\infty} B_n = A^{-1}$.

REFERENCES

- [1] M. Zorzi and R. Rao. Geographic random forwarding (GeRaF) for ad hoc and sensor networks: Multihop performance. *IEEE Trans. on Mobile Computing*, vol. 2, no. 4, Oct.-Dec. 2003.
- [2] C. Lott and D. Teneketzis. Stochastic routing in ad-hoc networks. IEEE Trans. on Automatic Control, vol. 51, no. 1, pp. 52-70, Jan. 2006.
- [3] T. Javidi and D. Teneketzis. Sensitivity analysis for optimal routing in wireless ad hoc networks in presence of error in channel quality estimation. *IEEE Trans. on Automatic Control*, vol. 49, no. 8, Aug. 2004.
- [4] F. Baccelli, B. Blaszczyszyn, and P. Muhlethaler. An aloha protocol for multihop mobile wireless networks. *IEEE Transactions on Information Theory*, 2006.
- [5] R. Laufer, H. Dubois-Ferrière, and L. Kleinrock. Multirate anypath routing in wireless mesh networks. Proc. IEEE INFOCOM, 2009.
- [6] S. Biswas and R. Morris. ExOR: Opportunistic multi-hop routing for wireless networks. Proc. of Sigcomm, 2005.
- [7] P. Larsson and N. Johansson. Multiuser diversity forwarding in multihop packet radio networks. IEEE Proc. WCNC, 2005.
- [8] M. J. Neely and R. Urgaonkar. Optimal backpressure routing in wireless networks with multi-receiver diversity. Ad Hoc Networks (Elsevier), vol. 7, no. 5, pp. 862-881, July 2009.