EE 126 Probability and Random Processes
March 15, 2018 Continuous-Time Markov Chains II

1 Lecture

CTMCs

A CTMC is defined by $\mathcal{X}$ a countable set of states, π an initial probability distribution on $\mathcal{X}$, and Q a rate matrix. The state space is still discrete, but time is now continuous.

$$P(X_0 = i) = \pi(i)$$
$$P(X_{t+s} = j \mid X_t = i, X_u, u < t) = P(X_{t+s} = j \mid X_t = i)$$

for all $s > 0$, $t \geq 0$, and $i, j \in \mathcal{X}$.

Let ϵ be s where s is small. Then

$$P(X_{t+\epsilon} = j \mid X_t = i, X_u, u < t) = \begin{cases} \epsilon Q(i,j) + o(\epsilon) & \text{for } j \neq i \\ 1 - \epsilon \sum_{j \in \mathcal{X}, j \neq i} Q(i,j) + o(\epsilon) & \text{for } j = i \end{cases}$$

$Q = \{Q(i,j)\}\ \forall\ i, j \in \mathcal{X}$ such that

$$Q(i,j) \geq 0\ \forall\ j \neq i$$
$$\sum_j Q(i,j) = 0 \implies \text{row sums of } Q \text{ are equal to } 0$$

Note: Q is a rate matrix, *not* a probability transition matrix.

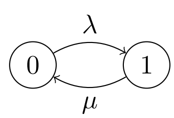

For example, the rate matrix for the above CTMC is $Q = \begin{bmatrix} -\lambda & \lambda \\ \mu & -\mu \end{bmatrix}$. $Q(i,j)$ is drawn next to the arrow from i to j.

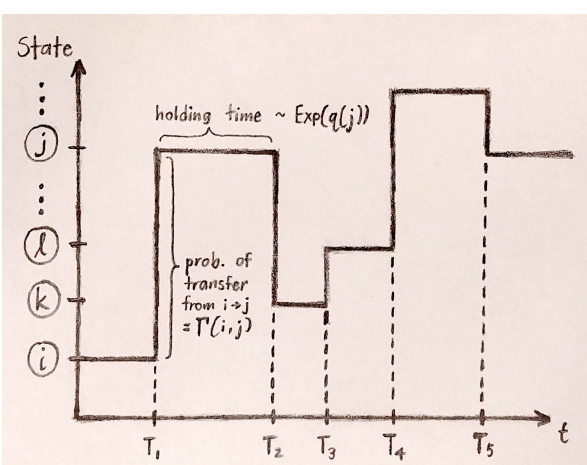

Excursions into the state space. Note that the graph need not be a staircase anymore. We define *holding time* as the time we hold in a particular state before we jump. We define $\Gamma(i,j) = Q(i,j)/q(i)$, where $Q(i,j)$ is the rate at which we transfer from state i to state j (to use the clock analogy, the jumping probability after we have woken up), while $q(i)$ is the rate at which we escape from state i, period.

A CTMC is a generalization of a Poisson process.

$q(i) = -Q(i,i) = \sum_{j \in \mathcal{X}, j \neq i} Q(i,j)$, i.e. the rate at which the Markov chain exits state i.

If the current state is i, then the time to jump from state i is $\text{Exp}(q(i))$ (independent of the past history of the process *and* independent of the next state to which it jumps). If the current state is i, then the next state is j with probability $\Gamma(i,j) = Q(i,j)/q(i)$ (independent of the past history of the process).

To "visualize" how the CTMC evolves,

- Start at time t. Say $X(t) = 0$.
- Choose a random time $\tau = \text{Exp}(q(i))$.
- At time $t + \tau$, roll a die and jump to state j with probability $\Gamma(i,j) = Q(i,j)/q(i)$ for all $j \neq i$.

In summary, we have a clock and we have a die. When the clock rings, we roll a die to decide where to go. *Alternative interpretation*: keep many alarms, one for every other state. As soon as the first alarm rings, jump to that state.

Thus, if $X_t = i$, $P(X_{t+\epsilon} = j)$ is the probability that

- the process jumps in $[t, t+\epsilon] \implies q(i)\epsilon$, **and**
- it jumps to state $j \implies \Gamma(i,j)$

So $P(X_{t+\epsilon} = j \mid X_t = i) = q(i)\epsilon\Gamma(i,j) = q(i)\epsilon\frac{Q(i,j)}{q(i)} = \epsilon Q(i,j)$ to within $o(\epsilon)$.

Example.

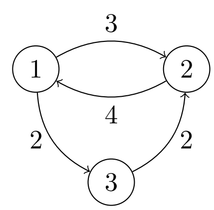

Recall: $Q(i,j) = 2$ signifies that if we're at i, we're jumping with rate 2 to j. To make some calculations,

$$\begin{aligned} q(1) &= Q(1,2) + Q(1,3) = 5 \\ q(3) &= Q(3,2) = 2 \\ \Gamma(1,2) &= Q(1,2)/q(1) = 3/5 \\ \Gamma(1,3) &= Q(1,3)/q(1) = 2/5 \\ \Gamma(3,2) &= 1 \end{aligned}$$

The full Q matrix ends up being $Q = \begin{bmatrix} -5 & 3 & 2 \\ 4 & -4 & 0 \\ 0 & 2 & -2 \end{bmatrix}$.

Example.

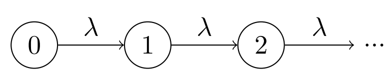

Here, $Q = \begin{bmatrix} -\lambda & \lambda & 0 & 0 & 0 & \cdots \\ 0 & -\lambda & \lambda & 0 & 0 & \cdots \\ 0 & 0 & -\lambda & \lambda & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$.

Example: B & T 7.14. We're working with a machine which can be in one of three states as follows:

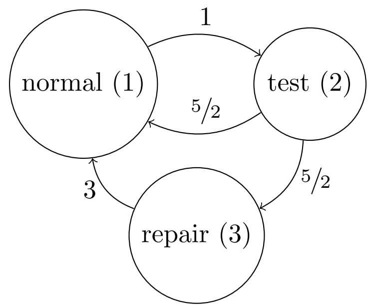

What is the stationary distribution π? (We can go from any state to any other state, so there must exist a stationary distribution.) Furthermore, we only make money when the machine is in the normal state. So what is the fraction of time that the machine will be in the normal state?

The balance equations still revolve around the idea that "flow in" must equal "flow out." However, when it comes to CTMCs, the balance equations are $\pi Q = 0$ (in contrast to DTMCs, where the balance equations are $\pi P = \pi$). Therefore, we have

$$\begin{aligned} \pi(1) \cdot 1 &= \pi(2) \cdot \frac{5}{2} + \pi(3) \cdot 3 \\ \pi(2) \cdot 5 &= \pi(1) \cdot 1 \\ \pi(1) + \pi(2) + \pi(3) &= 1 \end{aligned}$$

or

$$\begin{bmatrix} \pi_1 & \pi_2 & \pi_3 \end{bmatrix} \begin{bmatrix} -1 & 1 & 0 \\ 5/2 & -5 & 5/2 \\ 3 & 0 & 3 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \end{bmatrix}$$

According to Prof. Ramchandran, solving these equations yields $\pi = 1/41 \begin{bmatrix} 30 & 6 & 5 \end{bmatrix}$.
In other words, we're in state 1 (and making money!) about 75% of the time.

If we divide the rates by anything ≥ 5 (because 5 is the highest rate) and add self-loops weighted by 1 minus the sum of all outgoing weights, we can convert the rates into probabilities. By extension, our CTMCs become DTMCs.

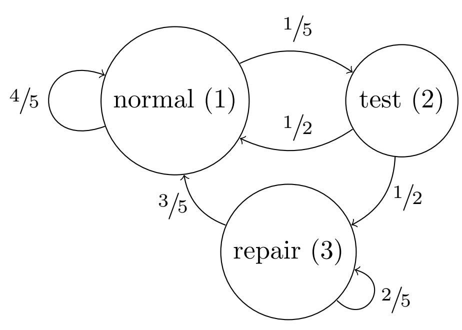

Example: lightbulbs. Consider 20 lightbulbs with independent lifetimes that are exponentially distributed with a mean of 1 month. What is the average time before all of the bulbs burn out? (Time, of course, is continuous here.)

In our Markov chain, each state will represent how many bulbs are up.

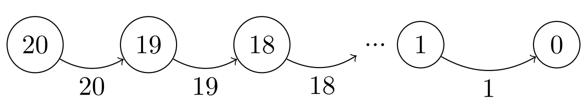

The numbers each describe the rate at which one more lightbulb burns out.
This is defined by a sum of Exp(1)s, one for each of the remaining lightbulbs.

The first step equations are

$$\begin{aligned}
\beta(20) &= \frac{1}{20} + \beta(19) \\
\beta(19) &= \frac{1}{19} + \beta(18) \\
&\vdots \\
\beta(m) &= \frac{1}{m} + \beta(m-1) \quad \text{for } m = 1, ..., 20 \\
&\vdots \\
\beta(0) &= 0
\end{aligned}$$

Again, $\beta(i)$ is the average time to hit 0 if we start at i. $1/m$ is the average time it takes for an alarm to ring – recall that the exponential rate is m. If we solve these equations, we should observe that $\beta(20) \approx \ln 20$.

Say there are 3 states, numbered 1 to 3, to which the state i can transition. If $Q(i,1) = 5$, $Q(i,2) = 6$, and $Q(i,3) = 7$ then we can think of it as setting an Exp(r_i) alarm clock for each state $i \in \{1,2,3\}$ (where r_i = the rate for state i). We will go to the state whose alarm clock goes off first. Until then, we shall remain in i.