EE 126 Probability and Random Processes
March 13, 2018 Continuous-Time Markov Chains I

1 Reading

7.5. Continuous-Time Markov Chains

We model a Markov chain where the time spent between transitions is a continuous random variable. We will call X_n the state right after the nth transition, Y_n the time of the nth transition, and T_n the time between the $(n-1)$st and the nth transition. Furthermore, X_0 is the initial state and $Y_0 = 0$.

If the current state is i, we will assume that

- the time until the next transition is exponentially distributed with a given parameter v_i
- the next state will be j with a given probability p_{ij}

independent of the past history of the process and of the next state/transition. Thus

$$\begin{aligned} P(X_{n+1} = j, T_{n+1} \geq t \mid T_1 = t_1, ..., T_n = t_n, X_0 = i_0, ..., X_n = i) &= P(X_{n+1} = j, T_{n+1} \geq t \mid X_n = i) \\ &= P(X_{n+1} = j \mid X_n = i) P(T_{n+1} \geq t \mid X_n = i) \\ &= p_{ij} e^{-v_i t} \end{aligned}$$

The parameter v_i is called the *transition rate out of state* i. The quantity $q_{ij} = v_i p_{ij}$ is the *transition rate from* i *to* j, i.e. the average number of transitions from i to j per unit time spent at i.

If we have a CTMC with a single recurrent class, then the states j are associated with steady-state probabilities π_j with the properties that $\lim_{t\to\infty} P(X(t) = j \mid X(0) = i) = \pi_j$, $\pi_j = 0$ for transient states, and $\pi_j > 0$ for recurrent states. We can solve for the steady-state probabilities according to the balance equations

$$\begin{aligned} \pi_j \sum_{k \neq j} q_{jk} &= \sum_{k \neq j} \pi_k q_{kj} \quad \text{for } j = 1, ..., m \\ \sum_{k=1}^{m} \pi_k &= 1 \end{aligned}$$

W13.5. Continuous-Time Markov Chains

A continuous-time setting is sometimes simpler because only one event can occur at once. A continuous-time Markov chain is defined by its initial distribution π and rate matrix Q, and exhibits the property that

$$P(X_{t+\epsilon} = j \mid X_t = i, X_u, u < t) = \mathbb{1}\{i = j\} + Q(i, j)\epsilon + o(\epsilon)$$

In other words, the process jumps from i to $j \neq i$ with probability $Q(i, j)\epsilon$ in $\epsilon \ll 1$ time steps. Therefore $Q(i, j)$ is the probability of jumping from i to j per unit of time.

If a CTMC is irreducible, then the states are either all transient, all positive recurrent, or all null recurrent (corresponding to the CTMC being "transient," "positive recurrent," or "null recurrent," respectively). There is no periodicity in continuous time.

If a CTMC is positive recurrent, it has a unique stationary distribution π where $\pi(i)$ is the long-term fraction of time that $X_t = i$. If a CTMC is not positive recurrent, it has no stationary distribution and the fraction of time spent in any state goes to 0.

2 Lecture

Poisson Process

$$P(N_t = k) = \frac{e^{-\lambda t}(\lambda t)^k}{k!}$$

The S_i's are i.i.d. $\sim \text{Exp}(\lambda)$. $\mathbb{E}[S_i] = {}^1\!/\!_\lambda$. $var(S_i) = {}^1\!/\!_{\lambda^2}$.

Example: lightbulbs. Two lightbulbs have independent and exponentially distributed lifetimes T_a and T_b with parameters λ_a and λ_b respectively. We want to determine the distribution of $Z = \min(T_a, T_b)$, i.e. the time for the first lightbulb to burn out:

$$\begin{aligned} P(Z > z) &= P(T_a > z)P(T_b > z) \\ &= e^{-\lambda_a z}e^{-\lambda_b z} \\ &= e^{-(\lambda_a+\lambda_b)z} \implies Z \sim \text{Exp}(\lambda_a + \lambda_b) \end{aligned}$$

This is a merged Poisson process: we are interested in the "first arrival time" of either of the two lightbulbs.

From this perspective, T_a and T_b are the inter-arrival times of the first arrivals of two independent Poisson processes with rates λ_a and λ_b, respectively. If we merge these two processes, the first arrival is at $\min(T_a, T_b)$. But this is the first arrival of the merged Poisson process whose rate is $(\lambda_a + \lambda_b)$! And this arrival time is exponentially distributed.

Erlang-k

The Erlang distribution of the kth order describes T_k, the time of the kth arrival.

$$T_k = S_1 + S_2 + ... + S_k = \text{ sum of } k \text{ i.i.d. Exp}(\lambda) \text{ random variables}$$

which implies that $\mathbb{E}[T_k] = {}^k\!/\!_\lambda$ and $var(T_k) = {}^k\!/\!_{\lambda^2}$. The distribution is

$$\begin{aligned} P(k\text{th arrival is in } (t, t+dt)) &= f_{T_k} \cdot dt \\ &= P(k-1 \text{ arrivals in the interval } (0,t) \text{ and } k\text{th arrival in } (t, t+dt)) \\ &= P(k-1 \text{ arrivals} \in (0,t)) \cdot P(1 \text{ arrival} \in (t, t+dt)) \\ &= \frac{e^{-\lambda t}(\lambda t)^{k-1}}{(k-1)!} \cdot \lambda dt \\ &\implies f_{T_k}(t) = \frac{e^{-\lambda t}\lambda^k t^{k-1}}{(k-1)!} \end{aligned}$$

We can use this to show that

$$\begin{aligned} P(N_t = k) &= \int_0^t P(N_t = k \mid T_k = s) f_{T_k}(s) ds \\ &= \int_0^t e^{-\lambda(t-s)} \frac{\lambda^k s^{k-1} e^{-\lambda s}}{(k-1)!} ds \\ &= \frac{e^{-\lambda t}}{(k-1)!}\lambda^k \int_0^t s^{k-1} ds \\ &= \frac{e^{-\lambda t}(\lambda t)^k}{k!} \end{aligned}$$

RIP

RIP stands for *random incidence "paradox"*. Fix a time t^* and consider the length L of the inter-arrival time containing t^*. (The analogy: we arrive at a bus stop at a certain time, and we're interested in the bus inter-arrival segment that our arrival time intersects.) How is L distributed?

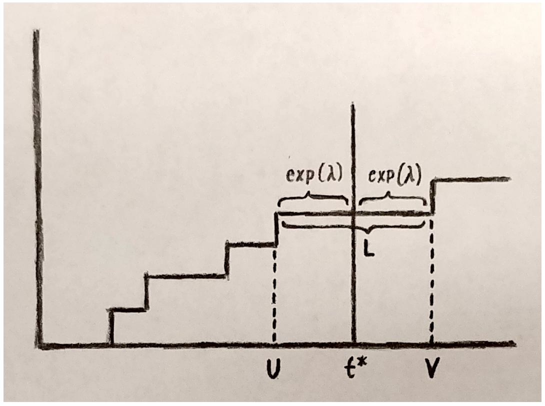

Setup for RIP. By t^*, we assume that the process has been running for an infinite amount of time, and accordingly that an arrival has already occurred. $[U, V]$ **is the inter-arrival interval that contains t^*, meaning $L = V - U$. (Specifically, U is the time of the first arrival prior to t^*; V is the time of the first arrival after t^*.)**

We have

$$\begin{aligned} P(t^* - U > x) &= P(x \text{ seconds have elapsed since the last "success"}) \\ &= P(0 \text{ arrivals in an interval of length } x) \\ &= e^{-\lambda x} \quad \sim \text{Exp}(\lambda) \end{aligned}$$

and

$$P(V - t^*) \sim \text{Exp}(\lambda) \quad \text{by memorylessness.}$$

Since $L = (t^* - U) + (V - t^*)$, we have that $L \sim \text{Exp}(\lambda) + \text{Exp}(\lambda) = \text{Erlang-2}$.

Continuous-Time Markov Chains

Define $\mathcal{X}$ to be a countable set of states. Let π be a probability distribution on $\mathcal{X}$, and Q be the rate matrix. Then a **continuous-time Markov chain** (CTMC), with initial distribution π and rate matrix Q, is a process $\{X_t\}_{t \geq 0}$ s.t.

$$P(X_0 = i) = \pi(i)$$

$$P(X_{t+\epsilon} = j \mid X_t = i, X_u, u \leq t) = P(X_{t+\epsilon} = j \mid X_t = i) \quad \text{for all } t \geq 0 \text{ and } (i,j) \in \mathcal{X}$$

$Q = \{Q(i,j)\}$ for all $i, j \in \mathcal{X}$ such that $Q(i,j) \geq 0$ (for $j \neq i$). We have that $\sum_{j \in \mathcal{X}} Q(i,j) = 0$, i.e. the row sum = 0.

References

[1] D.P. Bertsekas and J.N. Tsitsiklis. *Introduction to Probability.* Athena Scientific books. Athena Scientific, 2002.

[2] Jean Walrand. *Probability in Electrical Engineering and Computer Science: An Application-Driven Course.* Quorum Books, Westport, CT, USA, 2014.