Fundamentals of probability. 6.436/15.085

LECTURE 27
Birth-death processes

27.1. General birth-death processes

An important and a fairly tractable class of infinite continuous time M.c. is a birth-death process. Loosely speaking this is a process which combines the property of a random walk with reflection at zero, studied in the previous lecture and continuous time nature of the transition times. The process is describe as follows. The state space is $\mathbb{Z}_+$. To specify the transition we use the "exponential clock" model. To each state $n \geq 0$ two exponential clocks with rates λ_n, μ_n are attached. Assume the current state is $X(t) = n > 0$. If the clock corresponding to rate λ_n expires first, the chain moves to the state $n+1$ at time $t+U$. Otherwise, it moves to the state $n-1$ at time $t+V$. Formally, let $U \stackrel{d}{=} \text{Exp}(\lambda_n), V \stackrel{d}{=} \text{Exp}(\mu_n)$ be independent. If $U < V$, then $X(t+U) = n+1$. If $V < U$, then $X(t+V) = n-1$. The case $U = V$ occurs with probability zero, so it is ignored. In the special case $n = 0$ we assume that we only have a r.v. U and the chain moves into the state 1 after time U. Using the language of embedded M.c. we can alternatively describe the process as follows. Assume $X(t) = n$. Then a random time $U \stackrel{d}{=} \text{Exp}(\lambda_n + \mu_n)$ the chain moves to a new state which is $n+1$ with probability $\lambda_n/(\lambda_n+\mu_n)$, or $n-1$ with probability $\mu_n/(\lambda_n+\mu_n)$.

One can characterize the dynamics of birth-death processes in terms of a certain family of differential equations. Recall, that in addition to memoryless property, the exponential distribution has the following property: if $U \stackrel{d}{=} \text{Exp}(\lambda)$ then for every t, $\mathbb{P}(U \in [t, t+h)) = \lambda h + o(h)$. From this obtain the following relation

$$\mathbb{P}(X(t+h) = n+m|X(t) = n) = \begin{cases} \lambda_n h + o(h), & \text{if } m = 1; \\ \mu_n h + o(h), & \text{if } m = -1; \\ 1 - \lambda_n h - \mu_n h + o(h), & \text{if } m = 0; \\ o(h), & \text{if } |m| > 1; \end{cases}$$

Fix an arbitrary state k_0 and let $p_n(t) = \mathbb{P}(X(t) = n|X(0) = k_0)$. We have $p_n(t+h) = \sum_j p_j(t)\mathbb{P}(X(t+h) = n|X(t) = j)$. Assume $n > 0$. Then we can rewrite the expression above as

$$\begin{aligned} p_n(t+h) &= p_{n-1}(t)(\lambda_{n-1}h + o(h)) \\ &+ p_{n+1}(t)(\mu_{n+1}h + o(h)) \\ &+ p_n(t)(1 - \lambda_n h - \mu_n h + o(h)) \\ &+ o(h)(\sum_{j \neq n-1,n,n+1} p_j(t)) \end{aligned}$$

Note that $\sum_{j\neq n-1,n,n+1} p_j(t) \le 1$, so the last term in the sum is simply $o(h)$. We conclude

$$\frac{p_n(t+h) - p_n(t)}{h} = \lambda_{n-1}p_{n-1}(t) + \mu_{n+1}p_{n+1}(t) - (\lambda_n + \mu_n)p_n(t) + o(1).$$

By taking the limit $h \to 0$, we obtain a system of differential equations

$$\dot{p}_n(t) = \lambda_{n-1}p_{n-1}(t) + \mu_{n+1}p_{n+1}(t) - (\lambda_n + \mu_n)p_n(t) \tag{27.1}$$

The case $n = 0$ is obtained similarly and simplifies to

$$\dot{p}_0(t) = \mu_1 p_1(t) - \lambda_0 p_0(t). \tag{27.2}$$

The initial condition for this system of equations is $p_{k_0}(0) = 1, p_j(0) = 0, j \neq k_0$.

27.2. Steady state distribution

The system of differential equations we derived gives us a very good "hint" at what should be the stationary distribution, if one exists. Note that since all the states communicate, then one can have only one stationary distribution. We now use this system of differential equations to derive the stationary distribution. Then we will see an alternative way of deriving it using the condition $\pi' G = 0$ stated in the last lecture.

Steady state distribution is, by definition, a distribution which is time invariant. Thus we ask the following question: what should be a distribution π such that if we initialize our M.c. at π at time zero, as opposed a fixed state k_0, we get the same distribution at every time t? Clearly, for this we must have that all derivatives $\dot{p}_n(t)$ vanish and all $p_n(t) = p_k$ are constants. Thus we must have the following system of equations:

$$\lambda_{n-1}p_{n-1} + \mu_{n+1}p_{n+1} - (\lambda_n + \mu_n)p_n = 0, \ n \ge 1 \tag{27.3}$$

$$\mu_1 p_1 - \lambda_0 p_0 = 0. \tag{27.4}$$

On the other hand, if we initialize our M.c. at time zero with distribution $\pi = p = (p_n), n \ge 0$, then at every time t we obtain the same distribution, as the system of functions $p_n(t) = p, \forall t \ge 0$ solves (27.1),(27.2). Recursively solving this system we find that

$$p_n = \frac{\lambda_0\lambda_1 \cdots \lambda_{n-1}}{\mu_1\mu_2\cdots\mu_n}p_0 \tag{27.5}$$

Since we must also have that $\sum_{n\ge0} p_n = 1$, then we must have

$$p_0 = \Big(1 + \sum_{n\ge1}\frac{\lambda_0\lambda_1\cdots\lambda_{n-1}}{\mu_1\mu_2\cdots\mu_n}\Big)^{-1}, \tag{27.6}$$

Clearly, this has a solution iff

$$\sum_{n\geq 1}\frac{\lambda_0\lambda_1\cdots\lambda_{n-1}}{\mu_1\mu_2\cdots\mu_n}<\infty. \tag{27.7}$$

We arrive at the following result.

Theorem 27.8. *A birth-death process with parameters λ_n, μ_n has a stationary distribution if and only if the condition (27.7) holds. In this case the stationary distribution is unique and is given by (27.6),(27.5).*

There is a simpler alternative way of deriving the stationary distribution. Recall the condition $\pi' G$, where $G = (g_{i,j})$ is the generator matrix – the matrix of rates. In this case we have $g_{n,n+1} = \lambda_n, g_{n,n-1} = \mu_n, g_{n,n} = -(\lambda_n + \mu_n), n \geq 1$ and $g_{0,1} = \lambda_0, g_{0,0} = -\lambda_0$. Then the equation $\pi' G = 0$ translates into (27.3),(27.4).

Even a faster way, although the rigorization of this approach is based on the reversibility theory which we did not cover, is as follows. Observe that every time there is a transition from the state n to the state $n+1$ there must be a reverse transition (otherwise the state 0 will not be ever visited after some finite time period). In steady state the probability of transition $n \to n+1$ occurring at a given time interval $[t, t+h]$ is $\pi_n(\lambda_n h + o(h))$. The transition $n+1 \to n$ occurs in the same time interval with probability $\pi_{n+1}(\mu_{n+1}h + o(h))$. Since the two are the same, we obtain $\pi_n\lambda_n = \pi_{n+1}\mu_{n+1}$, which gives us the same solution (27.6),(27.6). These equations are sometimes called *balance equations* for an obvious reason.

27.3. Queueing systems

One of the immediate applications of birth-death processes is queueing theory. Consider the following system, known broadly as M/M/1 queueing system (M/M standing for memoryless arrival, memoryless service distribution). We have a server which serves arriving customer. Serving time for the n-th customer takes a random amount of time which is $\stackrel{d}{=} \text{Exp}(\mu)$. Customers arrive to the system according to the Poisson process with rate λ. Whenever there are no customers in the system, the server idles. Let $L(t)$ denote the number of customers at time t. What is the distribution of $L(t)$? What is the steady state distribution of $L(t)$ if any exists? It turns out that $L(t)$ is described exactly as the birth-death process with rates $\lambda_n = \lambda, \mu_n = \mu$. Indeed if $L(t) = n > 0$ then the next transition will occur either if there is an arrival, or if there is a service completion. Thus the transition occurs when the earlier of the two "exponential clocks" ticks. This happens at a random time $\stackrel{d}{=} \text{Exp}(\lambda+\mu)$ and the transition is to $n+1$ with probability $\lambda/(\lambda+\mu)$ and to $n-1$ with probability $\mu/(\lambda+\mu)$. When $L(t) = 0$ the only change is due to an arrival. So after a random time $\stackrel{d}{=} \text{Exp}(\lambda)$ we have a transition to state 1. We see that indeed we have a birth-death process. Clearly, the condition (27.7) holds if and only if $\rho = \lambda/\mu < 1$ and, in this case, we we obtain the following form of the stationary distribution.

$$p_0 = \Big(1 + \sum_{n\geq 1}\rho^n\Big)^{-1} = 1 - \rho,$$
$$p_n = (1-\rho)\rho^n.$$

There is an intuitive reason why we need condition $\rho < 1$. When $\rho > 1$ the customers arrive at a rate faster than service rate and on average queue builds up at a linear rate. Contrast this with the reflected random walk (essentially the same model except it is discrete time) and the condition $1-p<p$ for existence of steady state. The parameter ρ is called *traffic intensity* and plays a very important role in the theory of queueing systems. For one thing notice that in steady state $p_0=\mathbb{P}(L=0)=1-\rho$. On the other hand, from the general theory of discrete and continuous time M.c. we know that p_0 is the average fraction of time the system spends in this state. Thus $1-\rho$ is the average time the server is idle. Alternatively, ρ is the average time the system is busy. Hence the term "traffic intensity".

Consider now the following variant of a queueing system, known as $M/M/\infty$ system. Here we have infinitely many servers. Each service time is again $\stackrel{d}{=}\text{Exp}(\mu)$ and the arrival occurs according to a Poisson process with rate λ. The difference is that there is no queue any more as, due to infinity of servers, every customer gets instantly to be served. Let $L(t)$ be the number of customers being served at time t. It is not hard to see that this corresponds to a birth-death process with parameters $\lambda_n=\lambda,\mu_n=\mu n$. The arrival parameter explained as before. The service rate being μn is explained as follows: when there are n customers being served the next transition occurs at a time which is minimum of an arrival time till the next customer or the *smallest* of the n service time. The former is $\stackrel{d}{=}\text{Exp}(\lambda)$, the latter $\stackrel{d}{=}\text{Exp}(n\mu)$, hence $\mu_n=\mu n$. Let $\rho=\lambda/\mu$. In this case we find the stationary distribution as

$$p_0=\Big(\sum_{n\ge 0}\frac{\rho^n}{n!}\Big)^{-1}=e^{-\rho},$$

$$p_n=\frac{\rho^n}{n!}e^{-\rho}.$$

In particular, the distribution is a familiar $\text{Pois}(\rho)$. The system always has a steady state distribution, irrespectively of the values λ,μ. This is explained by the fact that we have infinitely many servers and the queue disappears.

27.4. References

- Sections 6.11, [**1**].

BIBLIOGRAPHY

1. G. R. Grimmett and D. R. Stirzaker, *Probability and random processes*, Oxford University Press, 2005.

6.436J / 15.085J Fundamentals of Probability
Fall 2008