6.262: Discrete Stochastic Processes 4/13/11

L18: Countable state Markov chains and processes

Outline:

- Review - Reversibility
- Sample-time M/M/1 queue
- Branching processes
- Markov processes with countable state spaces
- The M/M/1 queue

For any Markov chain,

$$\Pr\Big\{X_{n+k}, ...X_{n+1}|X_n, ...X_0\Big\} = \Pr\Big\{X_{n+k}, ..., X_{n+1} \mid X_n\Big\}$$

For any A^+ defined on X_{n+1} up and A^- defined on X_{n-1} down,

$$\Pr\Big\{A^+ \mid X_n, A^-\Big\} = \Pr\Big\{A^+ \mid X_n\Big\}$$

$$\Pr\Big\{A^+, A^- \mid X_n\Big\} = \Pr\Big\{A^+ \mid X_n\Big\}\Pr\Big\{A^- \mid X_n\Big\}.$$

$$\Pr\Big\{A^- \mid X_n, A^+\Big\} = \Pr\Big\{A^- \mid X_n\Big\}.$$

$$\Pr\Big\{X_{n-1} \mid X_n, X_{n+1}, \dots, X_{n+k}\Big\} = \Pr\{X_{n-1} \mid X_n\}.$$

The Markov condition works in both directions, but need steady state in forward chain for homogeneity in backward chain.

For a positive-recurrent Markov chain in steady-state, the backward probabiities are

$$\Pr\{X_{n-1} = j \mid X_n = i\} = P_{ji}\pi_j/\pi_i.$$

Denote $\Pr\{X_{n-1} = j \mid X_n = i\}$ **as the backward transition probabilities. Then**

$$\pi_i P^*_{ij} = \pi_j P_{ji} \; = \; \Pr\{X_n = i, X_{n-1} = j\}\,.$$

Def: A chain is reversible if $P^*_{ij} = P_{ij}$ **for all** i, j**.**

If chain is reversible, then $\pi_i P_{ij} = \pi_j P_{ji}$ **for all** i, j**, i.e., if** $\Pr\{X_n = i, X_{n-1} = j\} = \Pr\{X_n = j, X_{n-1} = i\}$**. In other words, reversibility means that the long-term fraction of** i **to** j **transitions is the same as the long-term fraction of** j **to** i **transitions.**

All positive-recurrent birth-death chains are reversible.

More general example: Suppose the non-zero transitions of a positive-recurrent Markov chain form a tree. Then the number of times a transition is crossed in one direction differs by at most one from the number of transitions in the other direction, so the chain is reversible.

Note that a birth-death chain is a very skinny tree.

The following theorem is a great time-saver and is sometimes called the guessing theorem.

Thm: For a Markov chain $\{P_{ij};\, i, j \geq 0\}$**, if a set of numbers** $\pi_i > 0, \sum_i \pi_i = 1$ **exist such that** $\pi_i P_{ij} = \pi_j P_{ji}$ **for all** i, j**, then the chain is positive-recurrent and reversible and** $\{\pi_i;\, i \geq 0\}$ **is the set of steady-state probabilities.**

Thm: $\pi_i P_{ij} = \pi_j P_{ji}$ for all i, j implies reversibility with $\{\pi_i\}$ steady-state.

Pf: Sum over i to get $\sum_i \pi_i P_{ij} = \pi_j \sum_i P_{ji} = \pi_j$. These (along with $\sum_i \pi_i = 1$ and $\pi_i \geq 0$) are the steady state equations and have a unique, positive solution.

Sanity checks for reversibility: 1) If $P_{ij} > 0$ then $P_{ji} > 0$. 2) If periodic, period is 2. 3) $P_{ij}P_{jk}P_{ki} = P_{ik}P_{kj}P_{ji}$.

Generalization of guessing thm to non-reversible chains: If $\exists\{\pi_i \geq 0;\, i \geq 0\}$ with $\sum_i \pi_i = 1$ and $\exists$ transition probabilites $\{P^*_{ij}\}$ such that $\pi_i P_{ij} = \pi_j P^*_{ji}$ for all i, j, then $\{\pi_i;\, i \geq 0\}$ are steady-state probabilities and $\{P^*_{ij}\}$ are the backward probabilities.

Suppose we sample the state of an M/M/1 queue at some time increment δ so small that we can ignore more than one arrival or departure in an increment. The rate of arrivals is λ and that of departures is $\mu > \lambda$.

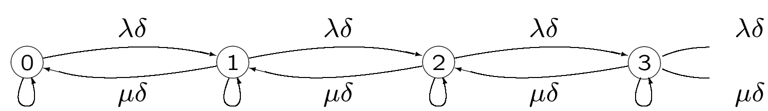

Either from the guessing theorem or the general result for birth/death chains, we see that $\pi_{n-1}\lambda\delta = \pi_n \mu\delta$ so, with $\rho = \lambda/\mu$,

$$\pi_n = \rho\pi_{n-1}; \qquad \pi_n = \rho^n \pi_0; \qquad \pi_n = (1-\rho)\rho^n$$

Curiously, this does not depend on δ (so long as $(\lambda+\mu)\delta \leq 1$), so these are the steady state probabilities as $\delta \to 0$.

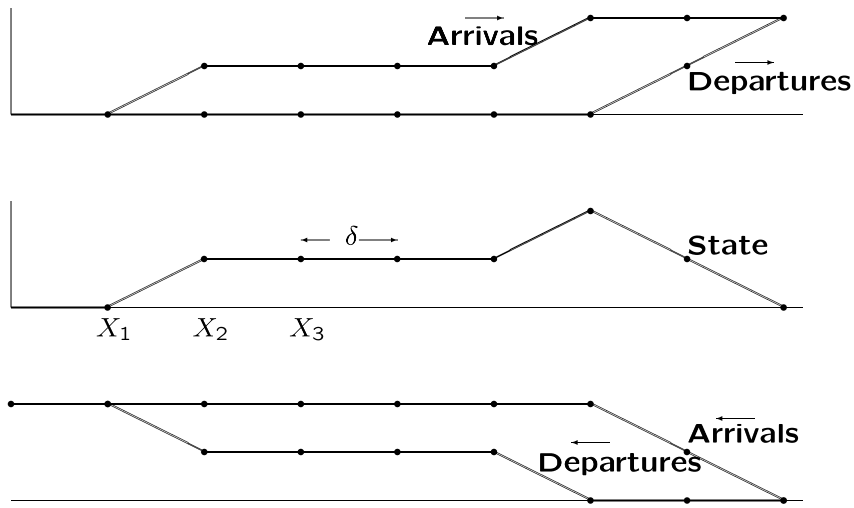

In the original (right-moving) chain, the state increases on arrivals and decreases on departures.

Each sample path corresponds to both a right and left moving chain, each M/M/1

Burke's thm: Given an M/M/1 sample-time Markov chain in steady state, first, the departure process is Bernoulli at rate λ. Second, the state at $n\delta$ is independent of departures prior to $n\delta$.

When we look at a sample path from right to left, each departure becomes an arrival and vice-versa. The right to left Markov chain is M/M/1.

Thus everything we know about the M/M/1 sample-time chain has a corresponding statement with time reversed and arrival-departure switched.

Branching processes

A branching process is a very simple model for studying how organisms procreate or die away. It is a simplified model of photons in a photomultiplier, cancer cells, insects, etc.

Let X_n be the number of elements in generation n. For each element k, $1 \leq k \leq X_n$, let $Y_{k,n}$ be the number of offspring of that element. Then

$$X_{n+1} = \sum_{k=1}^{X_n} Y_{k,n}$$

The nonnegative integer rv's $Y_{k,n}$ are IID over both n and k.

The initial generation X_0 can be an arbitrary positive rv, but is usually taken to be 1.

$$X_{n+1} = \sum_{k=1}^{X_n} Y_{k,n}$$

Examples: If $Y_{k,n}$ is deterministic and $Y_{k,n} = 1$, then $X_n = X_{n-1} = X_0$ for all $n \geq 1$.

If $Y_{k,n} = 2$, then $X_n = 2X_{n-1} = 2^n X_0$ for all $n \geq 1$.

If $\mathsf{p}_Y(0) = 1/2$ and $\mathsf{p}_Y(2) = 1/2$, then $\{X_n; n \geq 0\}$ is a rather peculiar Markov chain. It can grow explosively, or it can die out. If it dies out, it stays dead, so state 0 is a trapping state.

The state 0 is a trapping state in general. The even numbered states all communicate (but, as we will see, are all transient), and each odd numbered state does not communicate with any other state.

Let's find the probability (for the general case) that the process dies out. Let

$$p_Y(k) = p_k \quad \textbf{and } \Pr\{X_n{=}j \mid X_{n-1}{=}i\} = P_{ij}.$$

Let $F_{ij}(n)$ be the probability that state j is reached on or before step n starting from state i. Then

$$F_{ij}(n) = P_{ij} + \sum_{k \neq j} P_{ik} F_{kj}(n-1), n > 1; \quad F_{ij}(1) = P_{ij}.$$

$$\begin{aligned} F_{10}(n) &= p_0 + \sum_{k=1}^{\infty} p_k [F_{10}(n-1)]^k \\ &= \sum_{k=0}^{\infty} p_k [F_{10}(n-1)]^k. \end{aligned}$$

Let $h(z) = \sum_k p_k z^k$. Then $F_{10}(n) = h(F_{10}(n-1)$.

Let $h(z) = \sum_k p_k z^k$. Then $F_{10}(n) = h(F_{10}(n-1)$.

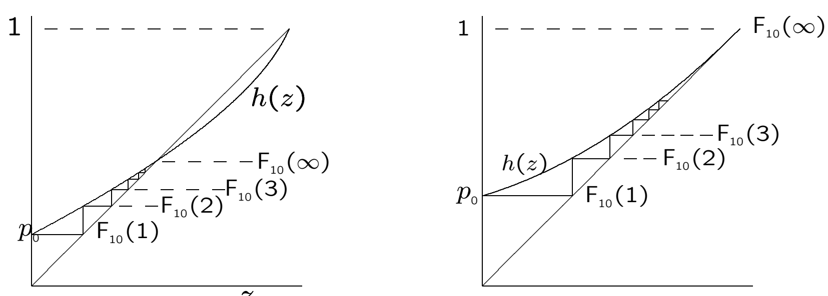

(a) **(b)**

We see that $F_{10}(\infty) < 1$ in case (a) and $F_{10}(\infty) = 1$ in case (b). For case (a), $h'(z)_{z=1} = \overline{Y} > 1$ and in case (b), $h'(z)_{z=1} = \overline{Y} \leq 1$.

For case a), the process explodes (with probability $1 - F_{10}(\infty)$) or dies out (with probability $F_{10}(\infty)$).

Markov processes

A countable-state Markov process can be viewed as an extension of a countable-state Markov chain. Along with each step in the chain, there is an exponential holding time U_i before the next step into state X_i.

The rate of each exponential holding time U_i is determined by X_{i-1} but is otherwise independent of other holding times and other states. The dependence is as illustrated below.

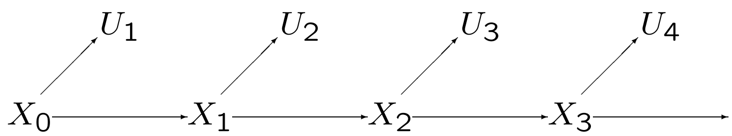

Each rv U_n, conditional on X_{n-1}, is independent of all other states and holding times.

The evolution in time of a Markov process can be visualized by

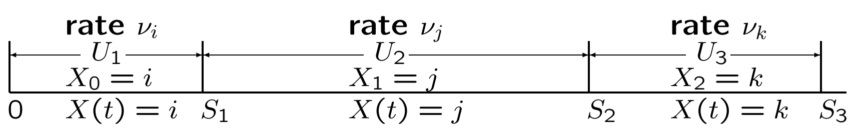

We will usually assume that the embedded Markov chain for a Markov process has no self-transitions, since these are 'hidden' in a sample path of the process.

The Markov process is taken to be $\{X(t); t \geq 0\}$. Thus a sample path of $X_n;\ n \geq 0$ and $\{U_n;\ n \geq 1\}$ specifies $\{X(t);\ t \geq 0\}$ and vice-versa.

$$\Pr\{X(t)=j \mid X(\tau)=i, \{X(s);\ s < \tau\}\} =$$

$$= \Pr\{X(t-\tau)=j \mid X(0)=i\}.$$

The M/M/1 queue

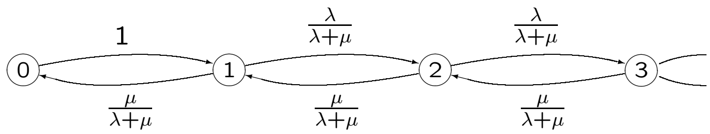

This diagram gives the embedded Markov chain for the M/M/1 Markov process. The process itself can be represented by

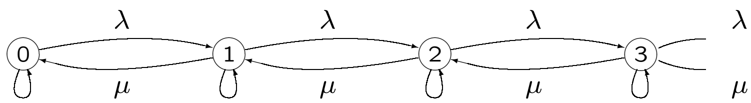

This corresponds to the rate of transitions given a particular state.