6.262: Discrete Stochastic Processes 4/20/11

L19: Countable-state Markov processes

Outline:

- Review - Markov processes
- Sampled-time approximation to MP's
- Renewals for Markov processes
- Steady-state for irreducible MP's

Markov processes

A countable-state Markov process can be defined as an extension of a countable-state Markov chain. Along with each step, say from X_{n-1} to X_n, in the embedded Markov chain, there is an exponential holding time U_n before X_n is entered.

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

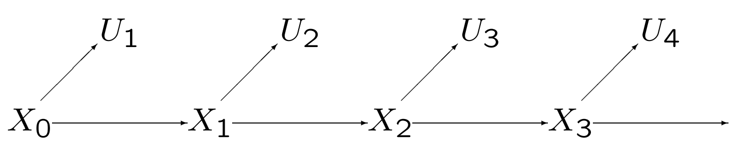

Each rv U_n, conditional on X_{n-1}, is independent of all other states and holding times.

In a directed tree of dependencies, each rv, conditional on its parent, is statistically independent of all earlier rv's. But the direction in the tree is not needed.

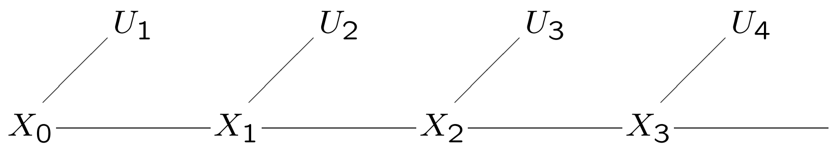

For example,

$$\begin{aligned}\Pr\{X_0X_1X_2U_2\} &= \Pr\{X_0\}\Pr\{X_1|X_0\}\Pr\{X_2|X_1\}\Pr\{U_2|X_1\} \\ &= \Pr\{X_1\}\Pr\{X_0|X_1\}\Pr\{X_2|X_1\}\Pr\{U_2|X_1\}\end{aligned}$$

Conditioning on any node breaks the tree into independent subtrees. Given X_2, (X_0, X_1, U_1, U_2) and (U_3) and (X_3, U_4) are statistically independent.

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

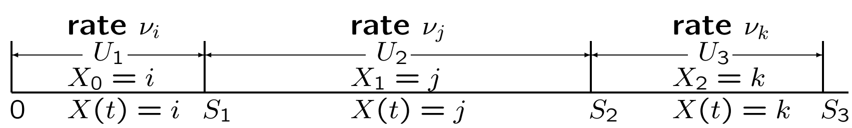

We usually assume that the embedded Markov chain for a Markov process has no self-transitions, since these are hidden in a sample path of the process.

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.

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

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

We can represent a Markov process by a graph for the embedded Markov chain with rates given on the nodes:

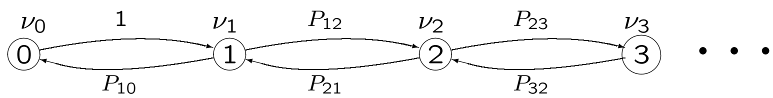

Ultimately, we are usually interested in the state as a function of time, namely the process $\{X(t);\ t \geq 0\}$. This is usually called the Markov process itself.

$$X(t) = X_n \quad \text{for } t \in [S_n, S_{n+1})$$

Self transitions don't change $X(t)$.

We can visualize a transition from one state to another by first choosing the state (via $\{P_{ij}\}$) then choosing the transition time (exponential with ν_i).

Equivalently, choose the transition time first, then the state (they are independent).

Equivalently, visualize a Poisson process for each state pair i, j with a rate $q_{ij} = \nu_i P_{ij}$. On entry to state i, the next state is the j with the next Poisson arrival according to q_{ij}.

What is the conditional distribution of U_1 given $X_0 = i$ and $X_1 = j$?

$$\nu_i = \sum_j q_{ij}; \qquad P_{ij} = q_{ij}/\nu_i : \qquad [q] \text{ specifies } [P], \vec{\nu}.$$

It is often more insightful to use q_{ij} in a Markov process graph.

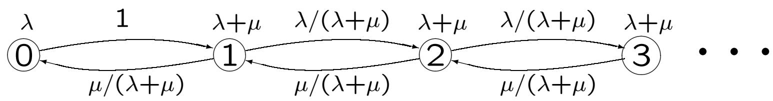

An M/M/1 queue using $[P]$ and $\vec{\nu}$

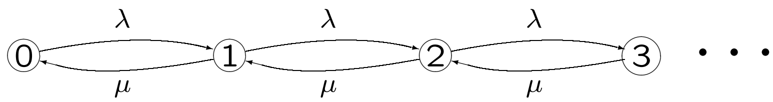

The same M/M/1 queue using $[q]$.

Both these graphs contain the same information. The latter corresponds more closely to our real-world interpretation of an M/M/1 queue.

Sampled-time approximation to MP's

Suppose we quantize time to δ increments and view all Poisson processes in a MP as Bernoulli with $P_{ij}(\delta) = \delta q_{ij}$.

Since shrinking Bernoulli goes to Poisson, we would conjecture that the limiting Markov chain as $\delta \to 0$ goes to a MP in the sense that $X(t) \approx X'(\delta n)$.

It is necessary to put self-transitions into a sampled-time approximation to model increments where nothing happens.

$$P_{ii} = 1 - \delta\nu_i; \qquad P_{ij} = \delta q_{ij} \qquad j \neq i$$

This requires $\delta \leq \frac{1}{\max \nu_i}$ and is only possible when the holding-time rates are bounded.

The embedded-chain model and sampled-time model of an M/M/1 queue:

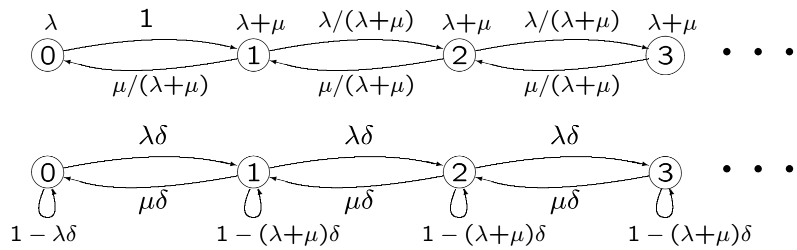

Steady state for the embedded chain, is $\pi_0 = (1-\rho)/2$ and $\pi_i = \frac{1}{2}(1-\rho)^2\rho^{n-1}$ for $i > 1$ where $\rho = \lambda/\mu$. The fraction of transitions going into state i is π_i.

Steady state for sampled-time does not depend on δ and is $\pi'_i = (1-\rho)\rho^i$ where $\rho = \lambda/\mu$. This is the fraction of time in state i.

Renewals for Markov processes

Def: An irreducible MP is a MP for which the embedded Markov chain is irreducible (i.e., all states are in the same class).

We saw that irreducible Markov chains could be transient - the state simply wanders off with high probability, never to return.

We will see that irreducible MP's can have even more bizarre behavior such as infinitely many transitions in a finite time or a transition rate decaying to 0.

Review: An irreducible countable-state Markov chain is positive recurrent iff the steady-state equations,

$$\pi_j = \sum_i \pi_i P_{ij} \text{ for all } j; \;\; \pi_j \geq 0 \text{ for all } j; \;\; \sum_j \pi_j = 1$$

have a solution. If there is a solution, it is unique and $\pi_i > 0$ **for all** i**. Also, the number of visits,** $N_{ij}(n)$**, in the first** n **transitions to** j **given** $X_0 = i$ **satisfies**

$$\lim_{n\to\infty} \frac{1}{n} N_{ij}(n) = \pi_j \quad \textbf{WP1}$$

We guess that for an MP, the fraction of time in state j **should be**

$$p_j = \frac{\pi_j/\nu_j}{\sum_i \pi_i/\nu_i}$$

Thm: Let $M_i(t)$ **be the number of transitions in** $(0, t]$ **for a MP starting in state** i**. Then** $\lim_{t\to\infty} M_i(t) = \infty$ **WP1.**

Essentially, given any state, a transition must occur within finite time. Then another, etc. See text.

Thm: Let $M_{ij}(t)$ **be the number of transitions to** j **in** $(0, t]$ **starting in state** i**. If the embedded chain is recurrent, then** $M_{ij}(t)$ **is a delayed renewal process.**

Essentially, transitions keep occuring so renewals into state j **must keep occuring.**

Steady-state for irreducible MP's

Let $p_j(i)$ be the time-average fraction of time in state j for the delayed RP $\{M_{ij}(t);\, t > 0\}$:

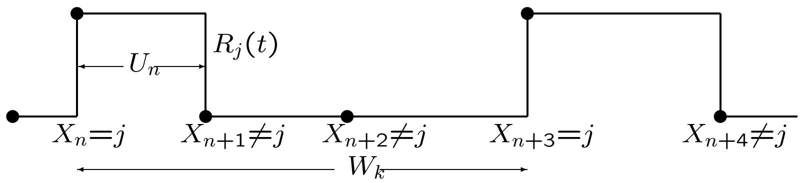

From the (delayed) renewal reward theorem,

$$p_j(i) = \lim_{t\to\infty} \frac{\int_0^t \; R_j(\tau)d\tau}{t} = \frac{\overline{U}(j)}{\overline{W}(j)} = \frac{1}{\nu_j \overline{W}(j)} \qquad \textbf{WP1.}$$

This relates the time-average state probabilities (WP1) to the mean recurrence times. Also $p_j(i)$ is independent of the starting state i.

If we can find $\overline{W}(j)$, we will also know p_j. Since $M_{ij}(t)$ is a (delayed) renewal process, the strong law for renewals says

$$\lim_{t\to\infty} M_{ij}(t)/t = 1/\overline{W}(j) \qquad \textbf{WP1}$$

$$\begin{aligned}\lim_{t\to\infty} \frac{M_{ij}(t)}{M_i(t)} &= \lim_{t\to\infty} \frac{N_{ij}(M_i(t))}{M_i(t)} \\ &= \lim_{n\to\infty} \frac{N_{ij}(n)}{n} = \pi_j \qquad \textbf{WP1}.\end{aligned}$$

$$\begin{aligned}\frac{1}{\overline{W}(j)} &= \lim_{t\to\infty} \frac{M_{ij}(t)}{t} = \lim_{t\to\infty} \frac{M_{ij}(t)}{M_i(t)}\frac{M_i(t)}{t} \\ &= \pi_j \lim_{t\to\infty} \frac{M_i(t)}{t} = p_j \nu_j\end{aligned}$$

This shows that $\lim_t M_i(t)/t$ is independent of i.

$$p_j = \frac{1}{\nu_j \overline{W}(j)} = \frac{\pi_j}{\nu_j} \lim_{t\to\infty} \frac{M_i(t)}{t} \qquad \textbf{WP1.}$$

Thm: If the embedded chain is positive recurrent, then

$$p_j = \frac{\pi_j/\nu_j}{\sum_k \pi_k/\nu_k}; \qquad \lim_{t\to\infty} \frac{M_i(t)}{t} = \frac{1}{\sum_k \pi_k/\nu_k} \quad \textbf{WP1}$$

If $\sum_k \pi_k/\nu_k < \infty$, this is almost obvious except for mathematical details. We can interpret $\lim_t M_i(t)/t$ as the transition rate of the process, and it must have the given value so that $\sum_j p_j = 1$.

It is possible to have $\sum_k \pi_k/\nu_k = \infty$. This suggests that the rate of transitions is 0.

Case where $\sum_k \pi_k/\nu_k = \infty$

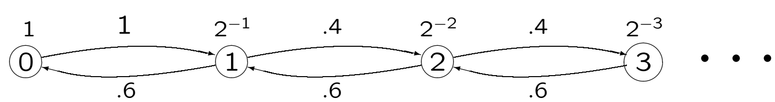

This can be viewed as a queue where the server becomes increasingly rattled and the customers increasingly discouraged as the state increases.

We have $\pi_j = (1-\rho)\rho^j$ for $\rho = 2/3$. Thus

$$\pi_j/\nu_j = 2^j(1-\rho)\rho^j = (1-\rho)(4/3)^j$$

By truncating the chain, it can be verified that the service rate approaches 0 as more states are added.

Again assume the typical case of a positive recurrent embedded chain with $\sum_i \pi_i/\nu_i < \infty$. Then

$$p_j = \frac{\pi_j/\nu_j}{\sum_k \pi_k/\nu_k} \qquad (1)$$

We can solve these directly using the steady-state embedded equations:

$$\begin{aligned} \pi_j &= \textstyle\sum_i \pi_i P_{ij}; \quad \pi_i > 0; \quad \sum_i \pi_i = 1 \\ p_j\nu_j &= \textstyle\sum_i p_i q_{ij}; \quad p_j > 0; \quad \sum_j p_j = 1 \qquad (2) \\ \pi_j &= \frac{p_j\nu_j}{\sum_i p_i\nu_i} \qquad (3) \end{aligned}$$

Thm: If embedded chain is positive recurrent and $\sum_i \pi/\nu_i < \infty$, then (2) has unique solution, $\{p_j\}$ and $\{\pi_j\}$ are related by (1) and (3), and

$$\sum_i \pi_i/\nu_i = (\sum_i p_j\nu_j)^{-1}$$

We can go the opposite way also. If

$$p_j\nu_j = \textstyle\sum_i p_i q_{ij}; \quad p_j > 0; \quad \sum_j p_j = 1$$

and if $\sum_j p_j\nu_j < \infty$, then $\pi_j = p_j\nu_j/(\sum_j p_j\nu_j)$ gives the steady-state equations for the embedded chain and the embedded chain is positive recurrent.

If ν_j is bounded over j, then $\sum_j p_j\nu_j < \infty$. Also the sampled-time chain exists and has the same steady-state solution.

For a birth/death process, we also have $p_i q_{i,1+1} = p_{i+1} q_{i+1,i}$.

If $\sum_j p_j\nu_j = \infty$, then $\pi_j = 0$ for all j and the embedded chain is transient or null-recurrent. In the transient case, there can be infinitely many transitions in finite time, so the notion of steady-state doesn't make much sense.

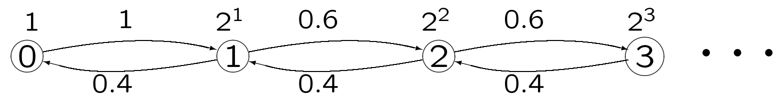

Imbedded chain for hyperactive birth/death

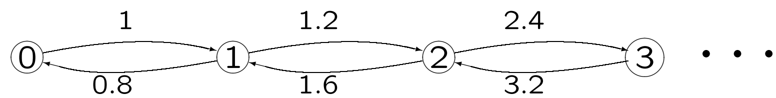

Same process in terms of $\{q_{ij}\}$

There is a nice solution for p_j, but the imbedded chain is transient.

These chains are called irregular. The expected number of transitions per unit time is infinite, and they don't make much sense.

6.262 Discrete Stochastic Processes
Spring 2011