6.262: Discrete Stochastic Processes
Lecture - Markov processes 4/8/09

- Definition

- The M/M/1 queue (again)

- Sampled time approximation (again)

- Steady state behavior

Definition of Markov process

As with most processes, we give several definitions, starting with the one that makes the structure most clear.

A Markov process is a combination of a countable state Markov chain $\{X_n; n \geq 1\}$ along with an exponential holding time rv for each state.

Assume the state space is $\{0, 1, \ldots\}$ if countably infinite and $\{1, 2, \ldots, k\}$ if finite.

If $X_n = j$, the holding time U_n for X_n has the distribution function

$$\mathsf{P}\{U_n \leq u \mid X_n = j\} = 1 - \exp(\nu_j u)$$

where $\nu_j > 0$ is the *rate* for state j. Given X_n, U_n is independent of everything else.

A Markov process is then specified by the transition probabilities P_{ij} for the embedded chain and the rates ν_i for each i in the state space.

For any sample point $x_0, u_1, x_1, u_2, \ldots$, the Markov process is in state x_0 until $s_1 = u_1$ when it jumps to x_1. Then at epoch $s_2 = u_1 + u_2$, it jumps to x_2, etc.

$$
\begin{tikzcd}
 & U_1 & & U_2 & & U_3 & & U_4 & \\
X_0 \arrow[ur] \arrow[rr] & & X_1 \arrow[ur] \arrow[rr] & & X_2 \arrow[ur] \arrow[rr] & & X_3 \arrow[ur] \arrow[rr] & & {}
\end{tikzcd}
$$

A dependence diagram can be used to illustrate the dependences here, e.g., U_1 and X_1 are independent given X_0

In a directed tree of dependencies, each rv depends only on its parent. But the direction in the tree is not needed.

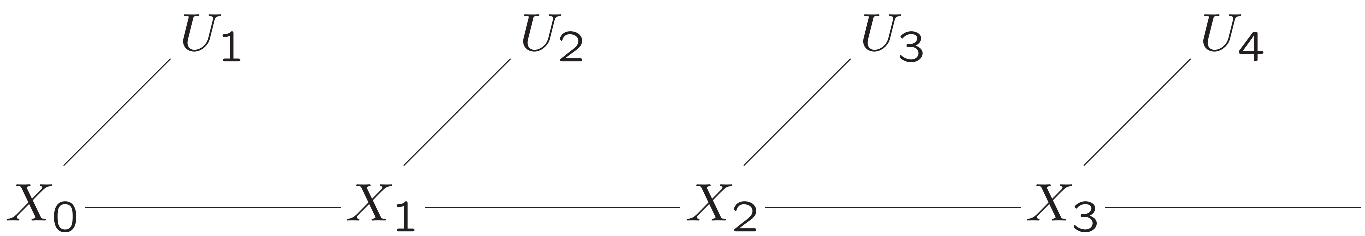

For example,

$$
\begin{aligned}
\mathsf{P}\{X_0X_1X_2U_2\} &= \mathsf{P}\{X_0\}\,\mathsf{P}\{X_1|X_0\}\,\mathsf{P}\{X_2|X_1\}\,\mathsf{P}\{U_2|X_1\} \\
&= \mathsf{P}\{X_1\}\,\mathsf{P}\{X_0|X_1\}\,\mathsf{P}\{X_2|X_1\}\,\mathsf{P}\{U_2|X_1\}
\end{aligned}
$$

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.

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

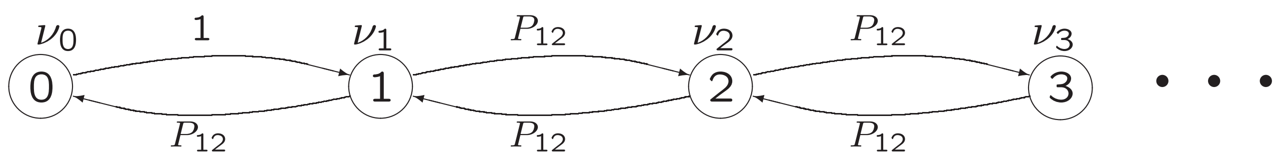

Ultimately, we are usually interested in the state as a function of time, namely the process $\{X(t); t \geq 0\}$, which is usually called the Markov process.

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

Self transitions are usually omitted since they 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 transition time first, then 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 $X_1 = j$?

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

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

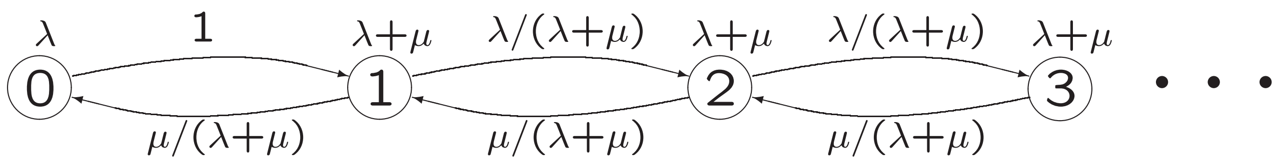

An M/M/1 queue using $[P]$ and ν

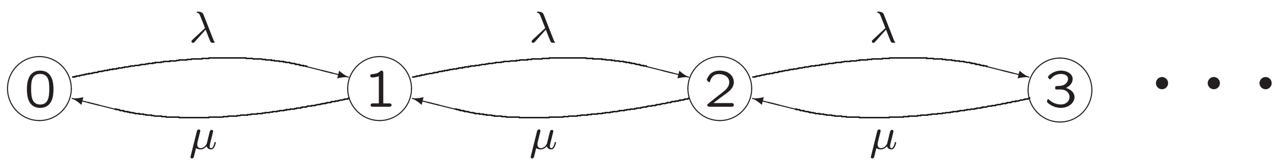

The same M/M/1 queue using $[q]$.

The latter corresponds 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 $p_{ij} = \delta q_{ij}$.

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 $\delta \leq \frac{1}{\max \nu_i}$.

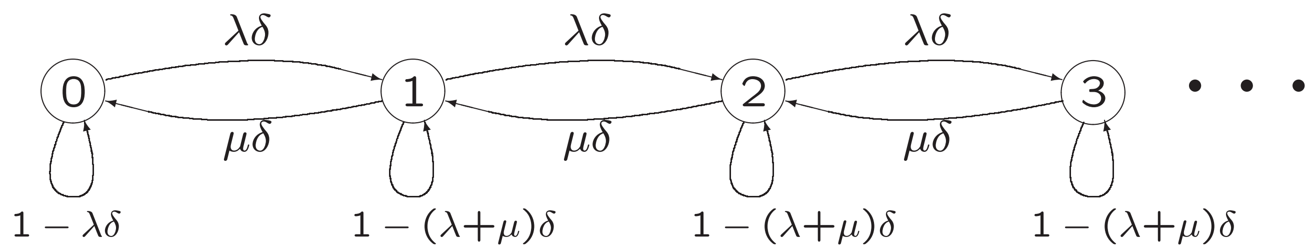

A sampled-time M/M/1 queue.

Note that the steady state probabilities do not depend on δ and are $\pi'_i = (1{-}\rho)\rho^i$ where $\rho = \lambda/\mu$.

In general, the steady state probabilities for a sampled time approximation give the fraction of time in each state exactly.

They are not equal to the steady state probabilities of the embedded chain.

Steady-state behavior of irreducible MP's

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.

MP's can have more highly varied behavior.

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

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

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

$$\lim_{n\to\infty} N_{ij}(n) = \pi_j \quad \textbf{W.P.1}$$

We guess that 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 $\lim_{t\to\infty} M_i(t) = \infty$ W.P.1

Pf: Since $X_0 = i$, U_1 is exponential with rate ν_i and is a rv (is finite W.P.1). U_n has the distribution

$$\mathbf{P}\{U_n \le u\} = \sum_j P_{ij}^n \exp(-\nu_j u)$$

This is a rv and $S_n = U_1 + \cdots + U_n$ is a rv (finite W.P.1.

$\lim_{t\to\infty} \mathbf{P}\{S_n \le t\} = 1$, so $\lim_{t\to\infty} \mathbf{P}\{M_i(t) \ge n\} = 1$ for all n. $M_i(t)$ is nondecreasing so has a limit for each sample point.

Thm: Let $M_{ij}(t)$ be number of transitions into state j by time t given $X_0 = i$. Then $M_{ij}(t)$ is a delayed renewal process if the embedded chain is recurrent.

pf: We have to show that the time of the first visit to j (given $X_0 = i$) is a rv (must occur).

We also have to show that the time from each visit to j until the next is a rv. (Murphy's law).

Use the fact that $M_{ij}(t) = N_{ij}(M_i(t))$.

Define the time average fraction of time in state j (up to time t), in terms of a renewal-reward process that is 1 whenever $X(t) = j$:

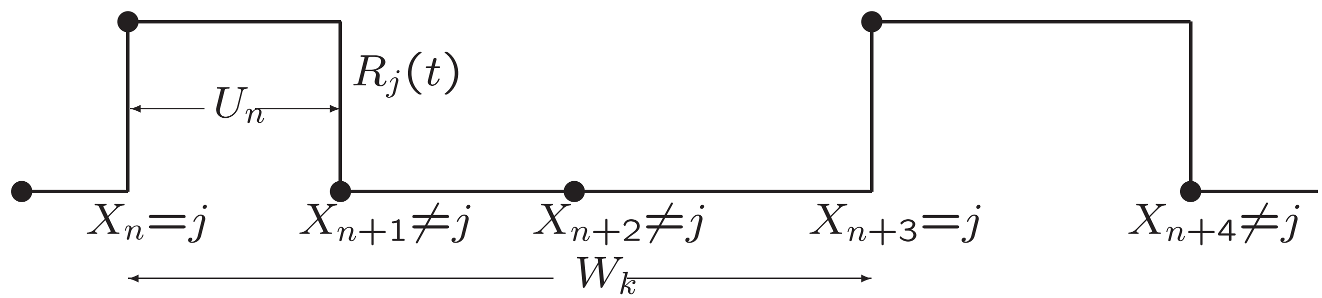

From the (delayed) renewal reward theorem,

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

Thus the time-average state probabilities (W.P.1) have been related to mean recurrence times.

If we could find $\overline{W}(j)$, we would know p_j. Since $M_{ij}(t)$ is a (delayed) renewal process, we have

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

This gives us a nice relation between two things we don't know. Let's relate this to the embedded steady state probabilities.

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

We now know that

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

This tells us that the limiting value of $M_i(t)/t$ is independent of i (remember we have assumed an irreducible recurrent chain).

Now assume a positive recurrent embedded chain and define $M(t)$ as the number of transitions up to time t starting in steady state. Since we stay in steady state, a renewal occurs on each transition. Thus

$$\lim_{t\to\infty} \frac{M(t)}{t} = \frac{1}{\sum_i \pi_i/\nu_i}$$

$$p_j = \frac{\pi_j/\nu_j}{\sum_i \pi_i/\nu_i} \qquad \textbf{W.P.1}$$