به نام خدا

Performance Evaluation of Computer Systems

Dr. Ali Movaghar

Fall 2014

Performance Modeling and Design of Computer Systems

9- ERGODICITY THEORY

Ergodicity Theory

- probability of being in state j
  - $\pi_j = \lim_{n \to \infty} P_{ij}^n$ is an ensemble average.
- Under what conditions does the limiting distribution exist?
- How does the limiting probability of being in state j, π_j , compare with the long-run time-average fraction of time spent in state j, P_j?
- What can we say about the mean time between visits to state j, and how is this related to π_j ?

Finite-State DTMCs

- Existence of the Limiting Distribution

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

- This chain is periodic; π_j does not exist, although $\lim_{n\to\infty} P_{jj}^{(2n)}$ does exist.

Finite-State DTMCs

- The period of state j is the greatest common divisor (GCD) of the set of integers n, such that $P_{j,j}^{n}$. A state is ***aperiodic*** if it has period 1. A chain is said to be aperiodic if all of its states are aperiodic.

- State j is ***accessible*** from state i if $P_{i,j}^{n}$for some $n > 0$. States i and j ***communicate*** if i is accessible from j and vice versa.

- A Markov chain is ***irreducible*** if all its states communicate with each other.

Finite-State DTMCs

- **Theorem** *Given an aperiodic, irreducible, finite-state DTMC with transition matrix P, as $n \to \infty$, $P_n \to L$ where L is a limiting matrix all of whose rows are the same vector, π. The vector π has all positive components, summing to* 1*.*

- For any aperiodic, irreducible, finite-state Markov chain, the limiting probabilities exist.

Mean Time between Visits to a State

- Consider an *irreducible* finite-state Markov chain with *M* states and transition matrix P.

- Let m_{ij} denote the expected number of time steps needed to first get to state j, given we are currently at state i. Likewise, let m_{ij} denote the expected number of steps between visits to state j.

Mean Time between Visits to a State

- ***Theorem*** *For an irreducible, aperiodic finite-state Markov chain with transition matrix* P

$$m_{jj} = \frac{1}{\pi_j}$$

where m_{ij} *is the mean time between visits to state* j *and* $\pi_j = \lim_{n\to\infty} P_{ij}^n$.

Time Averages

- For a finite-state Markov chain, the limiting distribution $\pi = (\pi_0, \pi_1, \dots, \pi_{M-1})$, when it exists, is equal to the unique stationary distribution.

- The fraction of time that the Markov chain spends in state *j*, P_j *is equal to* π_j.

Infinite-State Markov Chains

- f_j = probability that a chain starting in state *j* ever returns to state *j.*

- A state *j* is either recurrent or transient:
  - If f_j = 1, then *j* is a ***recurrent*** state.
  - If f_j < 1, then *j* is a ***transient*** state.

- Every time we visit state *j* we have probability $1 - f_j$ of never visiting it again. Hence the number of visits is distributed Geometrically with mean $1/(1 - f_j)$.

Infinite-State Markov Chains

- ***Theorem*** *With probability* 1*, the number of visits to a* **recurrent** *state is infinite. With probability* 1*, the number of visits to a* **transient** *state is finite.*

- ***Theorem***
  - $E[\#\ visits\ to\ state\ i\ in\ n\ steps \mid start\ in\ state\ i] = \sum_{n=0}^{s} P_{ii}^{n}$
  - $E[Total\ \#\ visits\ to\ state\ i \mid start\ in\ state\ i] = \sum_{n=0}^{\infty} P_{ii}^{n}$
- ***Theorem***
  - *If state i is recurrent, then* $\sum_{n=0}^{\infty} P_{ii}^{n}=\infty$
  - *If state i is transient, then* $\sum_{n=0}^{\infty} P_{ii}^{n} < \infty$

Infinite-State Markov Chains

- ***Theorem*** *If state i is recurrent and i communicates with j, (i ↔j), then j is recurrent.*
- ***Theorem*** *If state i is transient and i communicates with j, (i ↔ j), then j is transient.*
- ***Theorem*** *For a transient Markov chain:*

$$\lim_{n\to\infty} P_{ij}^n = 0, \quad \forall j.$$

- ***Theorem*** *If for a Markov chain*

$$\pi_j = \lim_{n\to\infty} P_{ij}^n = 0, \quad \forall j,$$

Then

$$\sum_{j=0}^{\infty} \pi_j = 0$$

so the limiting distribution does not exist.

Infinite-State Markov Chains

- **Theorem** For a transient Markov chain the limiting distribution does not exist.

- **Theorem** Given an aperiodic, irreducible chain. Suppose that the limiting probabilities are all zero. That is, $\pi_j = \lim_{n \to \infty} (P_{ij}^n) \quad , \forall j$. Then the stationary distribution does not exist.

Infinite Random Walk Example

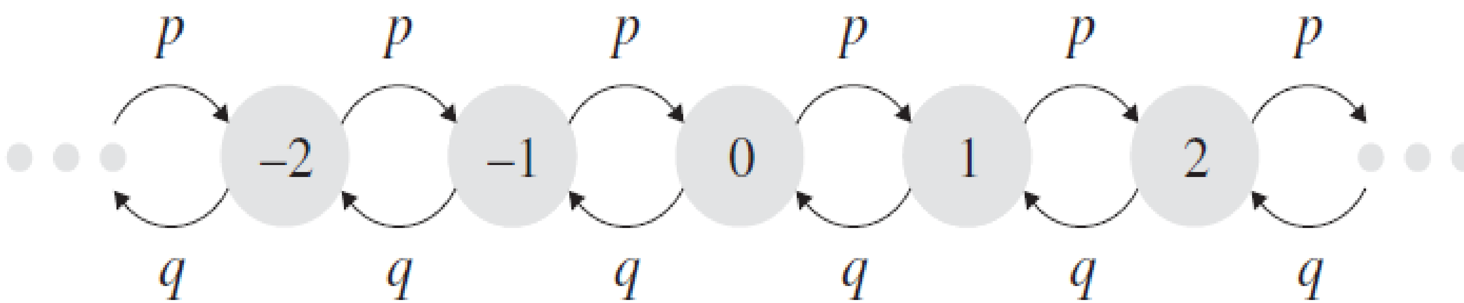

- All states are transient or all are recurrent. To determine whether the chain is recurrent or transient, it suffices to look at state 0.

$$V = \sum_{n=1}^{\infty} P_{00}^{n}$$

Infinite Random Walk Example

- If *V* is finite, then state 0 is transient, Otherwise it is recurrent

- Since one cannot get from 0 to 0 in an odd number of steps, it follows that

$$V = \sum_{n=1}^{\infty} P_{00}^{2n} = \sum_{n=1}^{\infty} \binom{2n}{n} p^n q^n$$

- The equation simplified by using Lavrov's lemma

Infinite Random Walk Example

- **Lavrov's lemma** (due to Misha Lavrov) For n ≥ 1,

$$\frac{4^n}{2n+1} < \binom{2n}{n} < 4^n$$

- **Theorem** The random walk shown in previous slide is recurrent only when $p = 1/2$ and is transient otherwise

Positive Recurrent versus Null Recurrent

- Recurrent Markov chains fall into two types: ***positive recurrent*** and ***null recurrent***. In a positive-recurrent MC, the mean time between recurrences (returning to same state) is finite. In a null-recurrent MC, the mean time between recurrences is infinite.

Positive Recurrent versus Null Recurrent

- **Theorem** If state i is positive recurrent and $i \leftrightarrow j$, then j is positive recurrent. If state i is null recurrent and $i \leftrightarrow j$, then j is null recurrent.

- **Theorem** *F*or the symmetric random walk shown in previous slide with $p = 1/2$,the mean number of time steps between visits to state 0 is infinite.

Ergodic Theorem of Markov Chains

- An ***ergodic*** DTMC is one that has all three desirable properties: aperiodicity, irreducibility, and positive recurrence.

- **Theorem** (Ergodic Theorem of Markov Chains) Given a recurrent, aperiodic, irreducible DTMC, $\pi_j = \lim_{n\to\infty} P_{ij}^n$ exists and

$$\pi_j = \frac{1}{m_{jj}}, \quad \forall j.$$

For a positive recurrent, aperiodic, irreducible DTMC, $\pi_j > 0, \forall\, j > 0$.

Ergodic Theorem of Markov Chains

- **Theorem** For an aperiodic, null-recurrent Markov chain, the limiting probabilities are all zero and the limiting distribution and stationary distribution do not exist.

Ergodic Theorem of Markov Chains

- **Theorem** (Summary Theorem) An irreducible, aperiodic DTMC belongs to one of the following two classes: Either:

i. All the states are transient, or all are null recurrent. In this case $\pi_j = \lim_{n\to\infty} P_{ij}^n = 0, \forall j$, and there does NOT exist a stationary distribution.

ii. All states are positive recurrent. Then the limiting distribution $\vec{\pi} = (\pi_0, \pi_1, \ldots)$,exists, and there is a positive probability of being in each state. Here $\pi_j = \lim_{n\to\infty} P_{ij}^n > 0, \forall j$ is the limiting probability of being in state j. In this case $\vec{\pi}$ is a stationary distribution, and no other stationary distribution exists. Also, π_j is equal to $\frac{1}{m_{ij}}$, where m_{ij} is the mean number of steps between visits to state j.

Time Averages

- **Theorem** For a positive recurrent, irreducible Markov chain, with probability 1,

$$p_j = \lim_{t\to\infty} \frac{N_j(t)}{t} = \frac{1}{m_{jj}} > 0,$$

where m_{jj} is the (ensemble) mean number of time steps between visits to state j and $N_j(t)$ be the number of times that the Markov chain enters state j by time t (t transitions)

Time Averages

- **Corollary** For an ergodic DTMC, with probability 1,

$$p_j = \pi_j = \frac{1}{m_{jj}}$$

Where $p_j = \lim_{t \to \infty} \frac{N_j(t)}{t}$ and $\pi_j = \lim_{n \to \infty} P_{ij}^n$ and m_{jj} is the (ensemble) mean number of time steps between visits to state j.

- **Corollary** For an ergodic DTMC, the limiting probabilities sum to 1 (i.e., $\sum_{j=0}^{j=\infty} \pi_j = 1$).

Time Averages

- **Theorem** (SLLN) Let $X_1, X_2, \dots$ be a sequence of independent, identically distributed random variables each with mean E[X]. Let $S_n = \sum_{i=1}^{n} X_i$. Then with probability 1,

$$\lim_{n \to \infty} \frac{S_n}{n} = \mathbf{E}[X].$$

- A ***renewal process*** is any process for which the times between events are i.i.d. random variables with a distribution *F*.

Time Averages

- **Theorem** (Renewal Theorem) For a renewal process, if E[X] is the mean time between renewals, we have

$$\lim_{t\to\infty} \frac{N(t)}{t} = \frac{1}{\mathbf{E}[X]} \textit{ with probability 1.}$$

Limiting Probabilities Interpreted as Rates

- $\pi_i P_{ij} =$ "rate" of transitions from state i to state j.
- $\sum_j \pi_i P_{ij}$ is the total rate of transitions out of state i, including possibly returning right back to state i.
- $\sum_j \pi_j P_{ji}$ This is the total rate of transitions into state i, from any state, including possibly from state i.

- Total rate leaving state i = Total rate entering state i

$$\pi_i = \sum_j \pi_i P_{ij} = \sum_j \pi_j P_{ji}.$$

Limiting Probabilities Interpreted as Rates

- balance equations

$$\sum_{j \neq i} \pi_i P_{ij} = \sum_{j \neq i} \pi_j P_{ji}$$

Time-Reversibility Theorem

- **Theorem** (Time-reversible DTMC) Given an aperiodic, irreducible Markov chain, if there exist $x_1, x_2, \ldots$ s.t.,$\forall\ i, j$,

$$\sum_i x_i = 1 \quad and \quad x_i P_{ij} = x_j P_{ji},$$

Then

1. $\pi_i = x_i$(the x_i ’s are the limiting probabilities).
2. We say that the Markov chain is time-reversible.

Time-Reversibility Theorem

This leads to the following simpler algorithm for determining the π_j's:

1. First try time-reversibility equations (between pairs of states): $x_i P_{ij} = x_j P_{ji}, \quad \forall i, j$ and $\sum_i x_i = 1$.
2. If you find x_i's that work, that is great! Then we are done: $\pi_i = x_i$.
3. If not, we need to return to the regular stationary (or balance) equations.

Periodic Chains

Lemma In an irreducible DTMC, all states have the same period.

Theorem In an irreducible, positive-recurrent DTMC with period $d < \infty$, the solution π to the stationary equations

$$\vec{\pi} \cdot \mathrm{P} = \vec{\pi} \quad and \quad \sum_i \pi_i = 1$$

exists, is unique, and represents the time-average proportion of time spent in each state.

Periodic Chains

- **Theorem** (Summary Theorem for Periodic Chains) Given an irreducible DTMC with period $d < \infty$, if a stationary distribution π exists for the chain, then the chain must be positive recurrent.

Equivalent representations of limiting probabilities