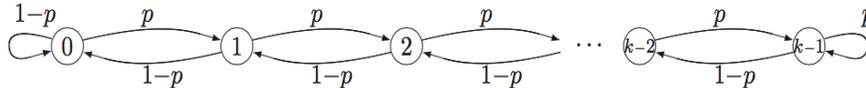


Problem Set 8

Spring 2020

1. Finite Random Walk



- (a) Assume $0 < p < 1$. Find the stationary distribution. *Hint:* Let $q = 1 - p$ and $\rho = \frac{p}{q}$, but be careful when $\rho = 1$.
- (b) Find the limit of π_0 and π_{k-1} , the stationary distribution probability for state 0 and $k - 1$, as $k \rightarrow \infty$.

2. Reversible Markov Chains

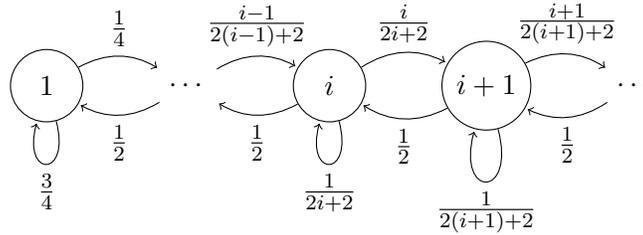
Let $(X_n)_{n \in \mathbb{N}}$ be an irreducible Markov chain on a finite set \mathcal{X} , with stationary distribution π and transition matrix P . The **graph** associated with the Markov chain is formed by taking the transition diagram of the Markov chain, removing the directions on the edges (making the graph undirected), removing self-loops, and removing duplicate edges. Show that if the graph associated with the Markov chain is a tree, then the Markov chain is reversible.

Hint: To solve this problem, try induction on the size of \mathcal{X} :

- (a) Use the fact that every tree has a leaf node x connected to only one neighbor y , and show that detailed balance holds for the edge (x, y) connecting the leaf with its single neighbor.
- (b) Then, argue that if you remove the leaf x from the Markov chain and increase the probability of a self-transition at state y by $P(y, x)$, then the stationary distribution of the original chain (when restricted to $\mathcal{X} \setminus \{x\}$) is the stationary distribution for the new chain, and use this to conclude the inductive proof.

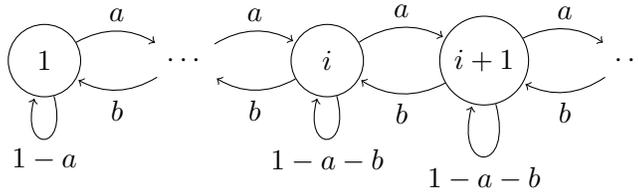
3. Markov Chains with Countably Infinite State Space

- (a) Consider a Markov chain with state space $\mathbb{Z}_{>0}$ and transition probability graph:



Show that this Markov chain has no stationary distribution.

- (b) Consider a Markov chain with state space $\mathbb{Z}_{>0}$ and transition probability graph:



Assume that $0 < a < b$ and $0 < a + b \leq 1$. Show that the probability distribution given by

$$\pi(i) = \left(\frac{a}{b}\right)^{i-1} \left(1 - \frac{a}{b}\right), \text{ for } i \in \mathbb{Z}_{>0},$$

is a stationary distribution of this Markov chain.

4. Product of Rolls of a Die

A fair die with labels (1 to 6) is rolled until the product of the last two rolls is 12. What is the expected number of rolls?

[Hint: You can model this process as a Markov chain with 3 states. Choose your states according to the outcome of last roll. For example, assign one state if it is outcome was 1 or 5 (which is useless if you want the product to be 12). If the outcome was 2,3,4 or 6, it's useful and can be assigned another state. Assign third state to the case when the product last two outcomes was 12.]

5. Poisson Practice

Let $(N(t), t \geq 0)$ be a Poisson process with rate λ . Let T_k denote the time of k -th arrival, for $k \in \mathbb{N}$, and given $0 \leq s < t$, we write $N(s, t) = N(t) - N(s)$. Compute:

- (a) $\mathbb{P}(N(1) + N(2, 4) + N(3, 5) = 0)$.
- (b) $\mathbb{E}(N(1, 3) \mid N(1, 2) = 3)$.
- (c) $\mathbb{E}(T_2 \mid N(2) = 1)$.

6. Illegal U-Turns

Each morning, as you pull out of your driveway, you would like to make a U-turn rather than drive around the block. Unfortunately, U-turns are illegal and police cars drive by according to a Poisson process with rate λ . You decide to make a U-turn once you see that the road has been clear of police cars for $\tau > 0$ units of time. Let N be the number of police cars you see before you make a U-turn.

- (a) Find $\mathbb{E}[N]$.
- (b) Let n be a positive integer ≥ 2 . Find the conditional expectation of the time elapsed between police cars $n - 1$ and n , given that $N \geq n$.
- (c) Find the expected time that you wait until you make a U-turn.

7. Metropolis-Hastings (optional)

This problem proves properties of the **Metropolis-Hastings Algorithm**, which you saw in lab.

Recall that the goal of MH was to draw samples from a distribution $p(x)$. The algorithm assumes we can compute $p(x)$ up to a normalizing constant via $f(x)$, and that we have a proposal distribution $g(x, \cdot)$. The steps are:

- Propose the next state y according to the distribution $g(x, \cdot)$.
- Accept the proposal with probability

$$A(x, y) = \min\left\{1, \frac{f(y) g(y, x)}{f(x) g(x, y)}\right\}.$$

- If the proposal is accepted, then move the chain to y ; otherwise, stay at x .
- (a) The key to showing why Metropolis-Hastings works is to look at the **detailed balance equations**. Suppose we have a finite irreducible Markov chain on a state space \mathcal{X} with transition matrix

P . Show that if there exists a distribution π on \mathcal{X} such that for all $x, y \in \mathcal{X}$,

$$\pi(x)P(x, y) = \pi(y)P(y, x),$$

then π is a stationary distribution of the chain (i.e. $\pi P = \pi$).

- (b) Now return to the Metropolis-Hastings chain. What is $P(x, y)$ in this case? For simplicity, assume $x \neq y$.
- (c) Show $p(x)$, our target distribution, satisfies the detailed balance equations with $P(x, y)$, and therefore is the stationary distribution of the chain.
- (d) If the chain is aperiodic, then the chain will converge to the stationary distribution. If the chain is not aperiodic, we can force it to be aperiodic by considering the **lazy chain**: on each transition, the chain decides not to move with probability $1/2$ (independently of the propose-accept step). Explain why the lazy chain is aperiodic, and explain why the stationary distribution is the same as before.