DISCRETE STOCHASTIC PROCESSES

Lecture 12

Finite State Markov Chains

Reading: Gallager, Chapter 4, pp. 103-120. (Will not emphasize proofs in Sect. 4.4.)

Brief Review

Markov Property, Graph and Matrix Representations, Classes of States,

Transient and Recurrent Classes, Periodic and Aperiodic Classes

6.041 Review – Absorption Probabilities

Matrix Approach

Chapmann - Kolmogorov Equation

Steady-State Behavior: Results from Perron-Frobenius Theory

6.041 Review – Long-Term Frequency Interpretations

Using Renewal Processes to Analyze Markov Chains

FINITE STATE MARKOV CHAINS

Definition: A **Finite State Markov Chain** is an Integer Time Process, $\{X_n; n \geq 0\}$ in which X_n, for each $n \geq 0$ is a random variable with possible values {1, 2, ..., J} with the **Markov Property**

Definition: **The Markov Property**

$$P\left(X_n = j \,\middle|\, X_{n-1} = i, X_{n-2} = h, \ldots, X_0 = m\right) = P\left(X_n = j \,\middle|\, X_{n-1} = i\right) = P_{ij} \quad \text{for all } n, i, j, h, m, \ldots.$$

A Markov chain is completely described by set of transition probabilities P_{ij} plus initial probabilities $P(X_0)$. Sometimes view $\{P_{ij}\}$ graphically, sometimes as matrix.

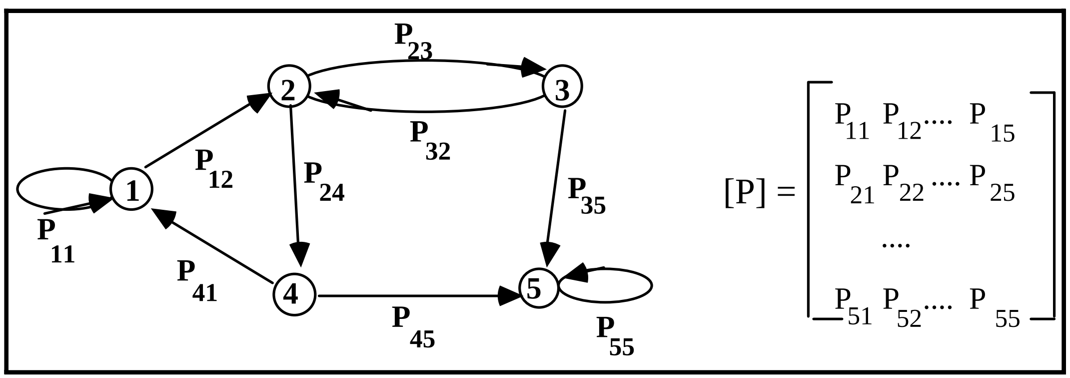

The graph emphasizes the possible and impossible.

Classification of States

Definitions: State j is **accessible** from i ($i \to j$) if path from i to j in graph. States i and j **communicate** if ($i \leftrightarrow j$) if ($i \to j$) and ($j \to i$).

If ($i \leftrightarrow j$) and ($j \leftrightarrow k$) then ($i \leftrightarrow k$). Given a walk from i to j, and a walk from j to k, walk from i can be extended to go to k and $i \to k$. Similarly $k \to i$.

Definition: A **Class** of states is a non-empty set S of states such that ($i \leftrightarrow j$) for each $i \in S$, $j \in S$ and also for $i \in S$, no $j \notin S$ with ($i \leftrightarrow j$).

A class can be though of as a **maximal** set of communicating states.

RECURRENT AND TRANSIENT STATES

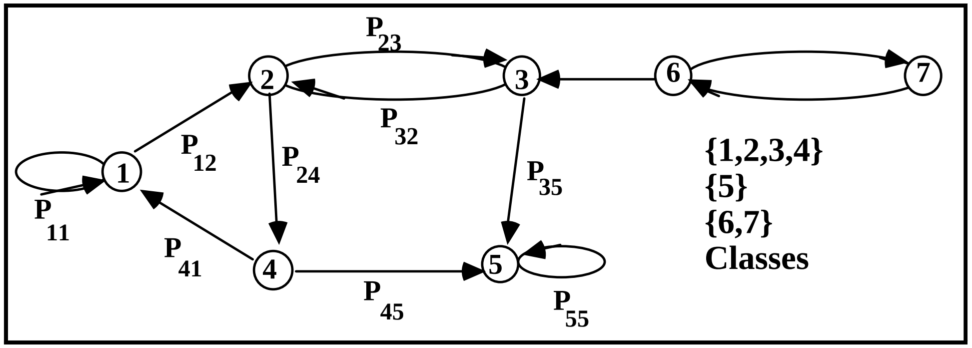

Definitions: A class S is **recurrent** if no $j \notin S$ is accessible from any $i \in S$.

A recurrent class can be thought of as a "trapping" class – once you get in, you never get out.
A class that is not recurrent is **transient**:

{5} recurrent, {1, 2, 3, 4}, {6, 7} transient

Transient means there is an outgoing edge from the class and no possible return.

PERIODIC AND APERIODIC STATES

Definition: The **period**, $d(i)$, of state i is defined as: $d(i) = \text{gcd}\{n : P_{ii}^n > 0\}$

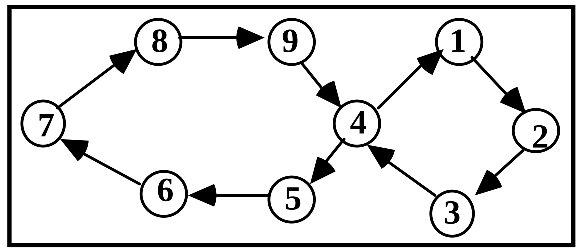

For example, $P_{44}^n > 0$ for n = 4, 6, 8, 10, ... The greatest common divisor is 2, $d(4) = 2$.
For state i = 1, $P_{11}^n > 0$ for n = 4, 8, 10, 12, ... $d(1)$ = 2.

If $d(i)$ = 1, i is defined to be aperiodic; otherwise it is periodic with period $d(i)$.

Theorem: All states in the same class have the same period.

<u>Theorem:</u> If period of a class is $d > 1$, then class has partition $T_1, T_2, \ldots T_d$ and all $i \in T_m$ have transitions only to T_{m+1} (or T_1 if $m = d$).

<u>Definition:</u> A chain is **ergodic** if it has only one class, and that class is aperiodic.

<u>Theorem:</u> An ergodic chain of J states has $P_{ij}^{m} > 0$ for all i, j and all $m \geq J(J-1)$.

6.041 Review – Absorption Probabilities and Expected First Passage Times

These results concern the finite-time behavior of a Markov chain with multiple recurrent classes. A state k is said to be ***absorbing*** if $p_{k,k}= 1$ (and, of course, $p_{k,j} = 0, \forall j \neq k.$) We begin with the simplest case, where each state is either transient or else absorbing.

For any Markov chain where every state is either transient or absorbing, let k be an absorbing state and a_i be the probability of eventually being absorbed into state k starting from state i. Then

$$a_k = 1$$

$$a_j = 0, \text{ all absorbing states } j \neq k$$

$$a_i = \sum_{j=1}^{J} p_{i,j} a_j, \text{ all transient states } i$$

This is a set of simultaneous equations for a_1, --- a_J.

We'd like to apply this to the more useful case in which we are given a Markov chain with multiple recurrent classes and one or more transient classes. Given a transient initial state, we wish to find the probability p_k the state eventually enters (and then necessarily remains in forever after) recurrent class k. Since transitions after the state first enters a recurrent class are irrelevant for answering this question, we can simplify the chain by replacing each recurrent class by a single state, so the modified chain satisfies the description on the previous page.

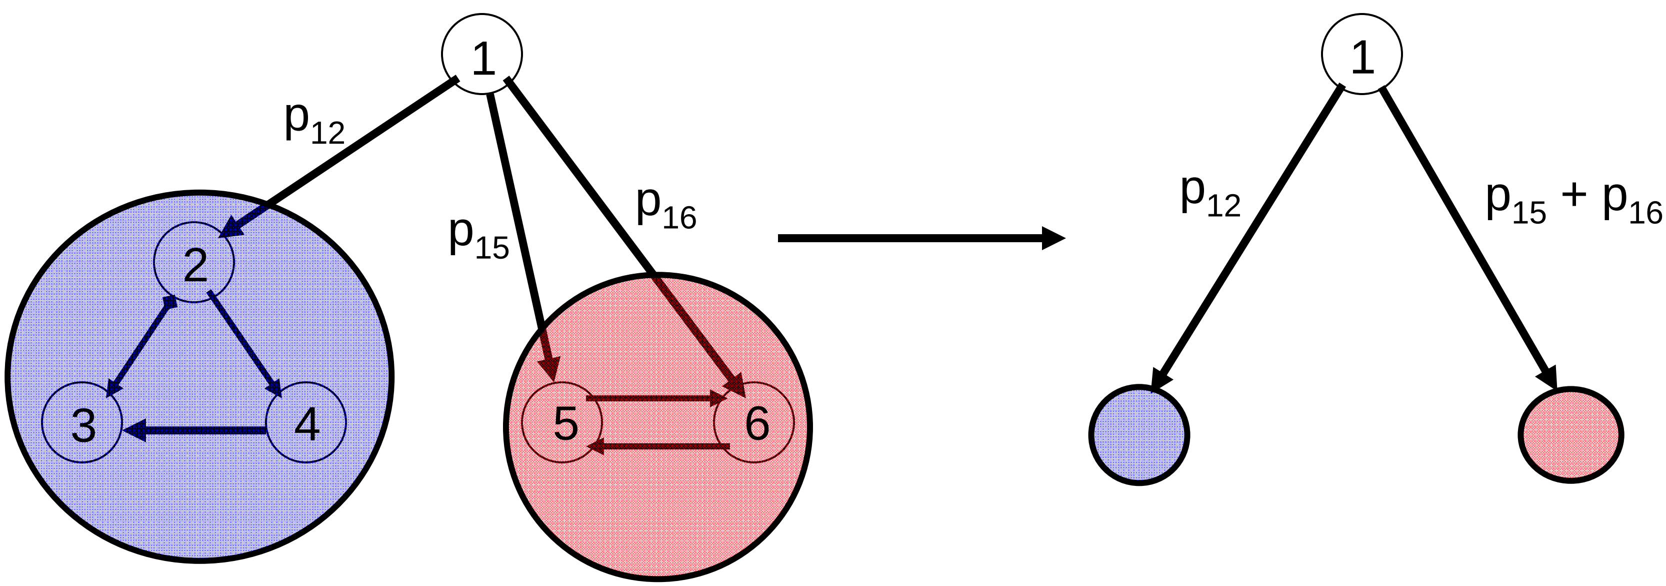

The Chapman - Kolmogoroff Equation

$$P_{ik}^2 = P(X_2 = k \mid X_0 = i) = \sum_{j=1}^{J} P(X_2 = k \mid X_1 = j, X_0 = i) P(X_1 = j \mid X_0 = 1) =$$

(by Markov property)

$$\sum_{j=1}^{J} P(X_2 = k \mid X_1 = j) P(X_1 = j \mid X_0 = i) = \sum_{j=1}^{J} P_{ij} P_{jk}$$

which means that P_{ik}^2 is the *i*, *k* element of $[P]^2$, i.e., the matrix multiplication of $[P]$ times itself. Similarly, P_{ik}^n is the *i*, *k* element of $[P]^n$.

$$P_{ik}^{n+m} = \sum_{j=1}^{J} P_{ij}^n P_{jk}^m$$

(Chapman - Kolmogoroff)

Let $p_k(n)$ be the (unconditional) probability state *k* is occupied at time *n*, and $p_k(0)$ be the prior probability state *k* is occupied at *n*=*0*. Then (from the "Total Probability Theorem")

$$p_k(1) = \sum_j p_j(0) P_{jk}.$$

$$p_k(1) = \sum_j p_j(0) P_{jk}.$$

Writing the unconditional probabilities as a row vector

$$\left(p_1(n),\ p_2(n),\ \ldots, p_J(n) \right) = \vec{p}(n),$$

we see that this can be expressed as a vector-matrix product

$$\vec{p}(1) = \vec{p}(0)[P]$$

$$\vec{p}(n) = \vec{p}(0)[P]^n$$

Long-Term Behavior (Steady-State)

$$\vec{p}(n) = \vec{p}(0)[P]^n$$

Primary interest lies in $[P]^n$ in limit as $n \to \infty$. If P_{ij}^n approaches a function of j, say π_j, as $n \to \infty$, then memory dies out and chain gets into a steady state independent of starting state.

Major Questions about Steady State

1) Does $[P]^n$ converge to a limiting matrix, $[P]^\infty$?

2) If so, are its rows identical ($\lim_{n\to\infty} P_{ij}^n = \lim_{n\to\infty} P_{kj}^n$)?

3) If so, $[P]^\infty = [P]^\infty [P]$, and $\vec{\pi} = \vec{\pi}[P]$. Does $\vec{\pi} = \vec{\pi}[P]$ always have a probability vector solution $\vec{\pi}$?

4) Does $\vec{\pi} = \vec{\pi}[P]$ have a unique probability vector solution $\vec{\pi}$?

5) Does $\vec{\pi} = \vec{\pi}[P]$, $\sum_i \pi_i = 1$ have unique solution?

Major Results about Steady-State

1) **Does $[P]^n$ converge?**

Answer: $[P]^n$ converges iff $[P]$ contains no periodic recurrent class. (A periodic class makes part of $[P]$ change periodically with n.)

2) **If so, are its rows identical?**

Answer: $\lim_{n\to\infty} P_{ij}^n = \lim_{n\to\infty} P_{kj}^n$ for all i, k, j iff $[P]$ has only one recurrent class and that recurrent class is aperiodic. (Chain can have transient classes, but if there is more than one recurrent class, where you end up depends on where you start).

3) **If so, does $\vec{\pi} = \vec{\pi}[P]$ always have a probability vector solution $\vec{\pi}$?**

Definition: A matrix is a **Stochastic Matrix** if it is square, all its elements are non-negative and all its rows sum to 1.

Answer: $\vec{\pi} = \vec{\pi}[P]$ has a probability vector solution for all stochastic matrices [P].

4) **Does $\vec{\pi} = \vec{\pi}P$ have a *unique* probability vector solution ?**

Answer: $\vec{\pi} = \vec{\pi}P$ has a ***unique*** probability vector solution iff the chain has a single recurrent class. (It can have transient classes and the recurrent class can be periodic) (If P has a periodic class, $\vec{\pi}$ averages the rows of P^n over a period, through P^n does not converge if the periodic class is recurrent. But if there is more than one recurrent class, even the averaged rows of P won't be the same.)

5) **Does** $\vec{\pi} = \vec{\pi}P, \sum_{j=1}^{J} \pi_j = 1$ **have a unique solution** $\vec{\pi}$ **in** R^J ?

Answer: Yes, $\vec{\pi} = \vec{\pi}P, \sum_{j=1}^{J} \pi_j = 1$ has a unique solution under the same conditions as (4). (This is helpful because we don't have to worry about all the elements of $\vec{\pi}$ being positive when solving.

These results don't come from conventional linear algebra, but rather from special results about non-negative matrices (Perron theorem, Frobenius theorem & consequences – Sect. 4.4)

Part of result 3 is elementary: If $[P]$ is a Markov chain, it is stochastic, so all rows sum to one, and $[P]\vec{1} = \vec{1}$ (which simply says $\sum_j P_{ij} * 1 = 1$ for all i) .

Thus $[P]$ has 1 as an eigenvalue with a right eigenvector $\vec{1}$.

This means that $[P - I]$ is singular (its rows are linearly dependent). Thus its columns are also linearly independent, so there is a non-zero solution to $\vec{\pi}[P - I] = \vec{0}$ or $\vec{\pi}[P] = \vec{\pi}$.

What is not easy is to show that there must be a real and non-negative solution.

6.041 Review: Long-Term Frequency Interpretations

Let $v_{i,j}(n)$ be the number of visits to state j from initial state i over the time period [0,n]. For a finite Markov chain with a single recurrent class,

$$\lim_{n\to\infty}\frac{v_{i,j}(n)}{n}=\pi_j,\ w.p.1,\ \text{and}\ \lim_{n\to\infty}\frac{E[v_{i,j}(n)]}{n}=\pi_j$$

Under the same assumption, let $q_{i,j,k}(n)$ be the number of transitions from state j to state k over the time period [0,n], starting from state i. Then

$$\lim_{n\to\infty}\frac{q_{i,j,k}(n)}{n}=\pi_j p_{j,k},\ w.p.1,\ \text{and}\ \lim_{n\to\infty}\frac{E[q_{i,j,k}(n)]}{n}=\pi_j p_{j,k}$$

These results were claimed but not proven in 6.041. They can be proved using results from 6.262 for (delayed) renewal processes.

Using Renewal Processes to Think About Finite Markov Chains

Example 1: Mean Recurrence Times

Suppose state 1 of a homogeneous finite Markov chain lies in an aperiodic recurrent class, and suppose the initial state is state 1 at time zero. What is the expected time to return to state 1?

Let $N(n)$ represent the total number of visits to state 1 for times $k = 1, 2, \cdots, n.$ This is an integer time renewal process, since the system *starts over* on every return to state 1.

(More specifically, X_j, the j-th interarrival interval or the time to return after the j-th visit to state 1, is independent of X_k, the time to return after the k-th visit because for $j < k$ the occupancy probabilities for $n > S_{k-1}$ are conditionally independent of those for $n < S_{k-1}$ given that the state at S_{k-1} is state 1. And X_j and X_k are identically distributed because the Markov chain is homogeneous.)

The integer time renewal process $N(n)$ becomes a continuous-time renewal process $N(t)$ if we define

$$N(t) = N(S_n),\ S_n \le t < S_{n+1}, \qquad (1)$$

i.e., by connecting the dots as shown.

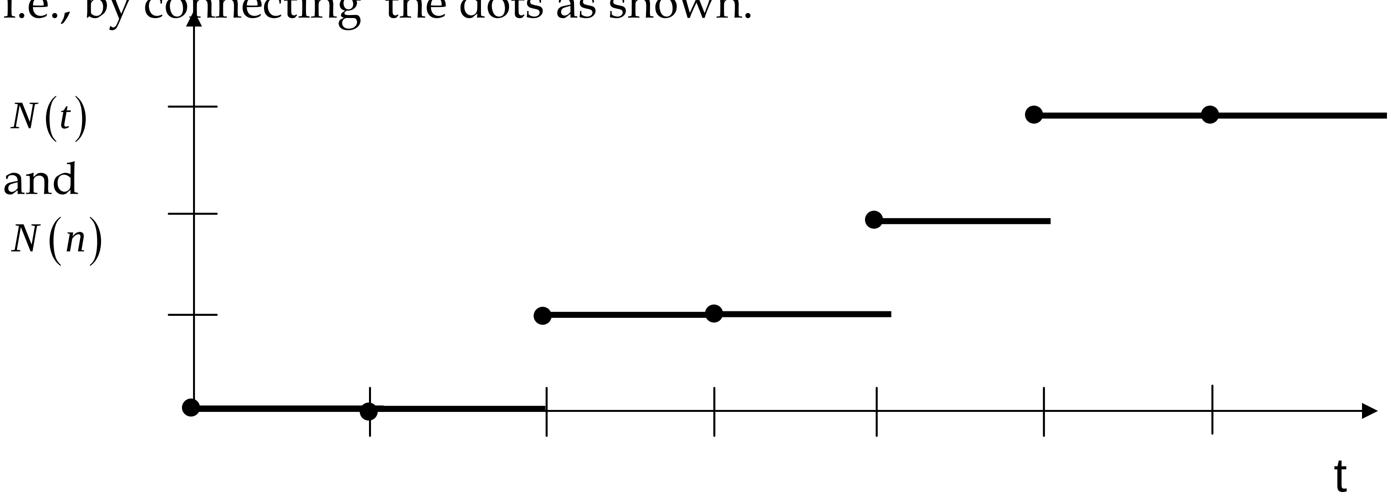

Since $N(t)$ can only jump at integer times, it has an arithmetic distribution. And since state 1 lies in an ***aperiodic*** recurrent class of the Markov chain, the span of the distribution is $d = 1$.

From Blackwell's Theorem

$$\frac{\lim_{m\to\infty} E\left\{N(m)-N(m-1)\right\}}{m-(m-1)=1} = \lim_{m\to\infty} P\left\{\text{renewal at time m}\right\} = 1/\bar{X}, \tag{2}$$

where $\bar{X}$ is the expected interrenewal time of $N(t)$, i.e., the expected time for the Markov chain to return to state 1. Furthermore, since state 1 lies in an aperiodic recurrent class of a finite Markov chain,

$$\lim_{m\to\infty} P\left\{\text{renewal at time m}\right\} = \lim_{m\to\infty} P_{11}^{m} = \pi_1. \tag{3}$$

Comparing (2) and (3), the expected first return time is

$$\bar{X} = 1/\pi_1. \tag{1}$$

Example 2: Mean First Passage Time

Now suppose the question were, instead, the expected time to reach another state k in the same class, $k \neq 1$, from state 1. There are now two events to consider and it isn't initially clear how renewal theory can help with the initial Markov chain. The trick is to collapse the two events into a single succession of events by modifying the transition probabilities so that $P_{k1} = 1,\ P_{kj} = 0,\ j \neq 1$, as shown below for k=4. Now each visit to k is immediately followed by a return to state 1, but the transition probabilities are unaltered until state k is reached again. Start the modified chain in state k at time zero, and let the renewal process (delayed renewal process?) N'(t) counts the total *returns to* state k up to and including time t. (N'(t) does not count the initial state being at k.)

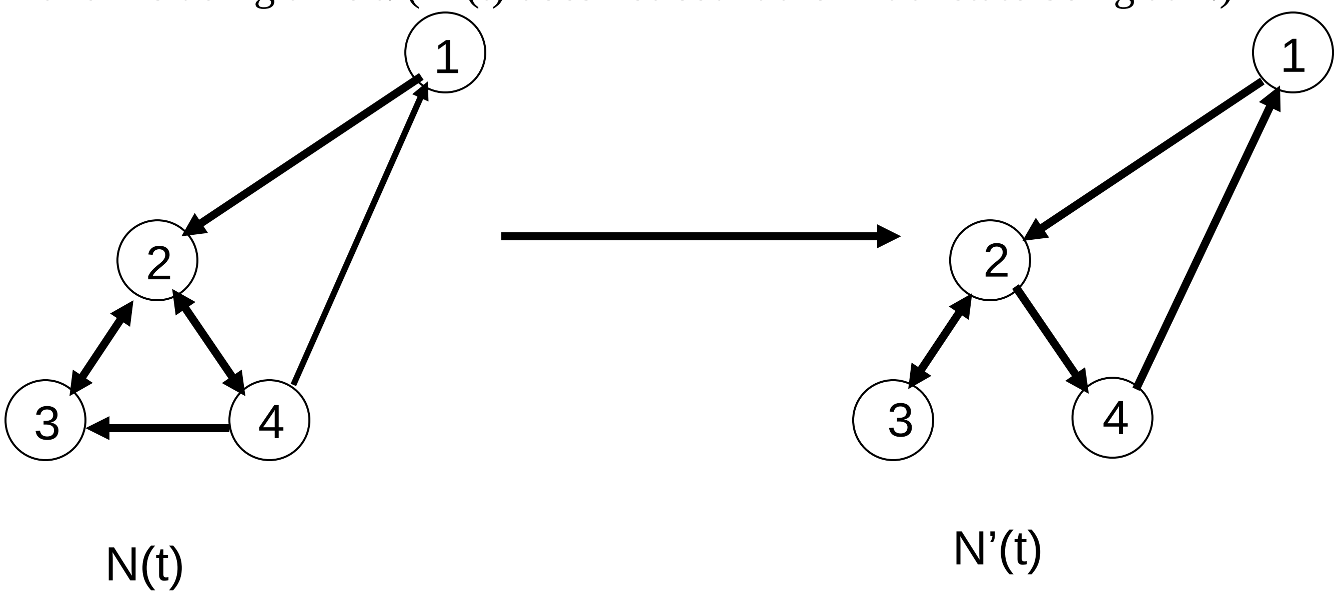

Again by Blackwell's Theorem,

$$\lim_{n\to\infty} P\{\text{renewal at time } n\} = 1/\bar{X}', \qquad (5)$$

where X' is the interrenewal time for $N'(t)$ and

$$X'_j = \{\text{time for j-th passage from state } k \text{ to state } k\} = 1 + \{\text{time for j-th passage from state 1 to state } k\}. \qquad (6)$$

Assuming the class containing states 1 and k in the modified chain remains aperiodic,

$$\lim_{n\to\infty} P\{\text{renewal at time } n\} = \lim_{n\to\infty} P'^{n}_{kk} = \pi'_k \qquad (1)$$

Comparing eqs. (5) and (7),

$$\bar{X}' = 1/\pi'_k, \qquad (2)$$

$$E\{\text{transition time from 1 to } k \text{ in the original chain}\} = (1/\pi'_k) - 1. \qquad (3)$$

Example 3:

With *admirable persistence* we modify the question again, and now ask the **expected number of returns to state 1 before the first passage to** $k.$ We modify the chain the same way as before, by setting $P_{k1}=1,$ the initial state at $n=0$ to $k,$ and let every return to state k count as a renewal, yielding the same renewal process $N'(t)$ as before. It should be clear by now that the continuous time variable t was just an artifice to let us draw on the continuous-time formulation we have been using, and the same arguments hold for $N'(n).$

Now we introduce a ***reward*** of 1 every time the state = 1, i.e., $r_n = 1 \Leftrightarrow X_n = 1$, otherwise $r_n = 0$. This figure helps to visualize the possible pattern of rewards:

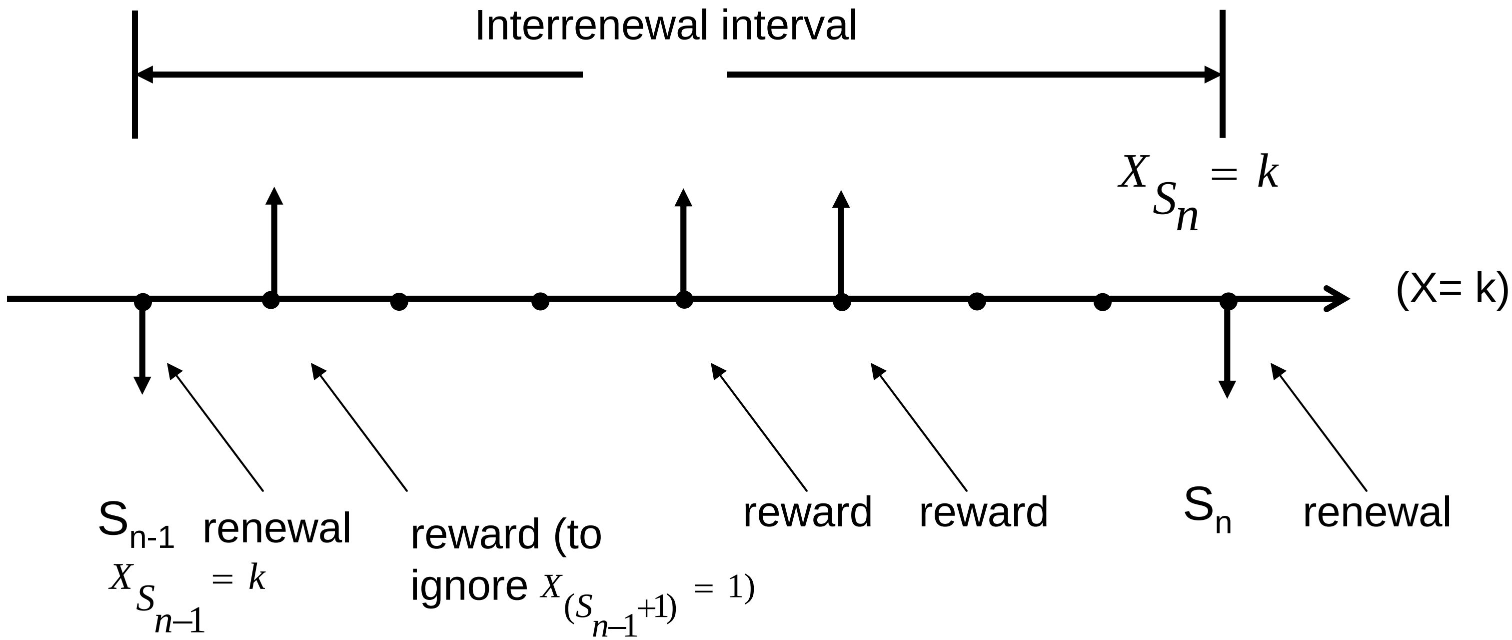

Note that there is never a reward at S_m. There is always a reward at S_m+1, but we don't count it as a return to state 1.

Only those rewards in X_n after $S_{n-1}+1$ will be counted as returns to state 1.

In particular,

$$\{\text{\# returns from state 1 to state 1 in } X_n\} = \{\text{total reward in } X_n\} - 1.$$

The *time average reward* is

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} r(k) = \pi_1'.$$

But from the Renewal Reward Theorem

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} r(k) = \frac{E\{\text{\# rewards in an interrenewal interval}\}}{\bar{X}'},$$

where the numerator on the right is (one greater than) the term we are after.

$$E\{\text{\# returns to 1 before visiting } k\} =$$

$$E\{\text{\# rewards in an interrenewal interval}\} - 1 = \bar{X}'\pi_1' - 1 = \frac{\pi_1'}{\pi_k'} - 1.$$

Note that again we assumed the modified chain remains aperiodic.

Interesting Technical Generalization in this Example

This type of reward is different from the age, residual life and duration examples in Chapter 3, in that the reward is no longer a deterministic function of the interarrival interval X. This reward $r(n)$ is, however, only a function of events during the inter-renewal interval containing n and is therefore legitimate.