CTMCs with Costs and Rewards

Natarajan Gautam
Department of Industrial and Systems Engineering
Texas A&M University
235A Zachry, College Station, TX 77843-3131
Email: gautam@tamu.edu
Phone: 979-845-5458
Fax: 979-847-9005

January 13, 2010

Abstract

We consider continuous time Markov chains (CTMCs) where there are costs and rewards associated with each state. For such a CTMC the objective usually is to calculate either the long-run average costs incurred per unit time or the total discounted cost incurred over an infinite horizon. We illustrate the calculations using several examples and conclude with applications to system performance analysis.

Consider a system that can be modeled as a CTMC $\{X(t), t \geq 0\}$ such that $X(t)$ is the state of the system at time t. Note that $X(t)$ could possibly be a vector but for notational convenience we sometimes write $P\{X(t) = j\}$ with the understanding that j in those cases is also a vector. Let S be the state space of this CTMC, i.e. the set of all possible values of $X(t)$ for all t. We use the notation Q to denote the infinitesimal generator matrix of the CTMC with negative values along the diagonals such that each row sums to zero. For terminologies and notation used for CTMCs refer to earlier articles of this encyclopedia. The reader is also encouraged to consult texts such as Kulkarni [1] and Ross [2].

For CTMCs with costs and rewards, an additional piece of notation is necessary. Define C_i as the cost incurred by the system per unit time it spends in state i, for all $i \in S$. Notice that if C_i is negative it is equivalent to the system getting a reward (as opposed to cost) of $-C_i$ per unit time it spends in state i. Here we present only analysis of systems over an infinite horizon. In particular, we consider two infinite horizon cases: average costs and discounted costs. Each case has its own pros and cons that we will subsequently discuss along with some examples. Then we conclude by presenting applications to performance analysis.

Average Costs Case

The objective here is to obtain an expression for the long-run average cost per unit time incurred by the system. For this we need some additional assumptions and notation. Recall that we consider a CTMC $\{X(t), t \geq 0\}$ with state space S, infinitesimal generator matrix Q and costs C_i per unit time the system is in state i for all $i \in S$. We assume that this CTMC is irreducible and positive recurrent (together it is also known as ergodic) with steady-state probabilities p_j, i.e.

$$p_j = \lim_{t \to \infty} P\{X(t) = j\}.$$

The p_j values can be computed by solving for $pQ = 0$ and $\sum_{i \in S} p_i = 1$, where p is a row vector of p_j values.

It is important to notice that for ergodic CTMCs, the steady state probabilities p_j have two other meanings. First, p_j is also the long-run fraction of time the system is in state j, i.e.

$$p_j = \lim_{T \to \infty} \frac{\int_0^T I(X(t) = j)dt}{T}, \tag{1}$$

where $I(\cdot)$ is an indicator function such that $I(X(t) = j) = 1$ if $X(t)$ is j at time t and $I(X(t) = j) = 0$ otherwise. Another definition of p_j is that it is the stationary probability that the CTMC is in state j. That means if the initial state of the CTMC is selected according to p_j, i.e. $P\{X(0) = j\} = p_j$, then at any time t the CTMC would be in state j with probability p_j, i.e. $P\{X(t) = j\} = p_j$ for all $t \geq 0$. Thus the system is considered to be stationary at all $t \geq 0$. Now we are in a position to develop an expression for the long run average cost per unit time for this system.

Let C be the long-run average cost incurred by the system per unit time. By definition, since C is time-averaged, it can be written as

$$C = \lim_{T \to \infty} \frac{\int_0^T C_{X(t)}dt}{T}.$$

We can rewrite the right hand side of the above expression using a sum to get

$$C = \lim_{T \to \infty} \frac{\int_0^T \sum_{i \in S} C_i I(X(t) = i)dt}{T}.$$

Since the CTMC is positive recurrent, we are allowed to write the above expression as

$$C = \sum_{i \in S} \lim_{T \to \infty} \frac{\int_0^T C_i I(X(t) = i)dt}{T}.$$

Using the definition of p_i in Equation (1) and substituting, we get

$$C = \sum_{i \in S} C_i p_i. \tag{2}$$

Discounted Costs Case

Here we consider the case where the costs incurred by the system are discounted at rate α. By definition that means if the system incurs a cost \$ d at time t, its present value at time 0 is $de^{-\alpha t}$. In other words, it is as though a cost of $de^{-\alpha t}$ is incurred at time 0. We call α as the discount factor and it is positive in sign. Having defined the discounted cost, we are ready to develop an expression for the expected total discounted cost over an infinite horizon. Let D_c be the total discounted cost (note this is a random variable) over an infinite horizon incurred by the system. Therefore we have by definition

$$D_c = \int_0^\infty C_{X(t)} e^{-\alpha t} dt.$$

Our objective is to obtain $E[D_c]$. For this we require two notations. Let a_i be the initial probability the CTMC is in state i, i.e.

$$a_i = P\{X(0) = i\}.$$

Notice that if the initial state of the CTMC is known, say j, then $a_j = 1$ and all other a_i's would be 0. Also let d_i denote the expected total discounted cost incurred over the infinite horizon starting in state i, i.e.

$$d_i = E[D_c|X(0) = i].$$

Thus we have

$$E[D_c] = \sum_{i \in S} E[D_c|X(0) = i] P\{X(0) = i\} = \sum_{i \in S} d_i a_i.$$

Note that $E[D_c]$ depends on the initial conditions via the a_i values. Thus we need to be given a_i, and assuming we have that, next we show how to compute d_i for all $i \in S$. Based on the definition we have

$$\begin{aligned} d_i &= E[D_c|X(0) = i] \\ &= E\left[\int_0^\infty C_{X(t)} e^{-\alpha t} dt | X(0) = i\right] \\ &= \int_0^\infty e^{-\alpha t} \sum_{j \in S} C_j P\{X(t) = j | X(0) = i\} dt \end{aligned}$$

provided C_j is bounded. Using the column vector d with elements d_i such that $d = [d_i]$ we can write down a matrix relation

$$d = \left[\int_0^\infty e^{-\alpha t} P(t) dt\right] \overline{C}$$

where $\overline{C}$ is a column vector of C_j values and $P(t)$ is a matrix of $P\{X(t) = j|X(0) = i\}$ for all $i \in S$ and $j \in S$. Using the Laplace Stieltjes Transform (LST) notation for $P(t)$ we get

$$\tilde{P}(s) = \int_0^\infty e^{-st} P(t) dt.$$

Thus

$$d = \tilde{P}(\alpha)\overline{C}.$$

However, since $P(t)$ satisfies $\frac{dP(t)}{dt} = P(t)Q$, by taking the LST we get $s\tilde{P}(s) - P(0) = \tilde{P}(s)Q$. Using the fact that $P(0) = I$ and solving for $\tilde{P}(s)$ we get

$$\tilde{P}(s) = (sI - Q)^{-1}.$$

Therefore we have

$$d = (\alpha I - Q)^{-1}\overline{C}.$$

Notice that the matrix $\alpha I - Q$ is invertible for any $\alpha > 0$. Thus we can compute d_i and hence $E[D_c]$.

Discussion and Examples

It is worthwhile to contrast the discounted costs against the average costs. In the discounted cost case, the CTMC does not have to be ergodic (it could be reducible or irreducible, transient or recurrent). However, the average cost results presented require the CTMC to be ergodic, although there are ways to deal with some non-ergodic cases. One of the demerits of the discounted cost case is that the results depend on the initial conditions, unlike the average cost case. Also the discount factor in many practical situations are hard to determine. Next we present one example for discount cost and one for average cost.

Example 1 *Consider a machine that toggles between being under working condition and under repair. The machine stays working for an exponential amount of time with mean $1/\gamma$ hours and then breaks down. A repair-person arrived instantaneously and spends an exponential amount of time with mean $1/\beta$ hours to repair the machine. A revenue of \$ r per hour is obtained when the machine is working and the repair-person charges \$ b per hour.*

Let $X(t)$ be the state of the machine at time t. Therefore if $X(t) = 0$, then the machine is under repair at time t. Also if $X(t) = 1$, the machine is up at time t. The state space $S = \{0, 1\}$. The generator matrix is

$$Q = \begin{bmatrix} -\beta & \beta \\ \gamma & -\gamma \end{bmatrix}.$$

Consider the discount cost case with discount factor α and given initial probabilities $a_0 = P\{X(0) = 0\}$ and $a_1 = P\{X(0) = 1\} = 1 - a_0$. In fact we could assume that at time 0 the machine could be up or down (i.e. a_0 is 0 or 1 respectively). We seek to obtain an expression for $E[D_c]$, the expected total discounted cost over an infinite horizon incurred by the system in terms of α, a_0, a_1, b, r, β and γ.

Using the discounted cost results we have

$$E[D_c] = E[D_c|X(0) = 0]P\{X(0) = 0\} + E[D_c|X(0) = 1]P\{X(0) = 1\} = d_0 a_0 + d_1 a_1$$

where

$$\begin{bmatrix} d_0 \\ d_1 \end{bmatrix} = (\alpha I - Q)^{-1} \begin{bmatrix} b \\ -r \end{bmatrix} = \begin{bmatrix} \frac{b(\alpha+\gamma)-r\beta}{\alpha(\alpha+\gamma+\beta)} \\ \frac{b\gamma-r(\alpha+\beta)}{\alpha(\alpha+\gamma+\beta)} \end{bmatrix}.$$

Notice that the notations used in the discounted costs derivation d and $\overline{C}$ are $d = [d_0 \ \ d_1]'$ and $\overline{C} = [b \quad -r]'$. The expected total discounted cost over an infinite horizon incurred by the system is therefore

$$E[D_c] = d_0 a_0 + d_1 a_1 = \frac{b\gamma - r\beta + (ba_0 - ra_1)\alpha}{\alpha(\alpha + \gamma + \beta)}.$$

Example 2 *Consider two independent and identical machines each with up times and repair times as described in Example 1. There is only one repair-person who repairs the machines in the order they failed.*

For this example we seek to obtain the long-run average cost per hour incurred by the system. Recall that this is denoted by C which can be obtained using Equation (2). Let $X(t)$ be the number of working machines at time t. The state space $S = \{0, 1, 2\}$. The generator matrix is

$$Q = \begin{bmatrix} -\beta & \beta & 0 \\ \gamma & -\beta - \gamma & \beta \\ 0 & 2\gamma & -2\gamma \end{bmatrix}.$$

The steady-state probabilities p_0, p_1 and p_2 can be obtained by solving for $[p_0 \ p_1 \ p_2]Q = [0 \ 0 \ 0]$ and $p_0 + p_1 + p_2 = 1$ as $p_0 = \frac{2\gamma^2}{2\gamma^2+2\gamma\beta+\beta^2}$, $p_1 = \frac{2\gamma\beta}{2\gamma^2+2\gamma\beta+\beta^2}$, and $p_2 = \frac{\beta^2}{2\gamma^2+2\gamma\beta+\beta^2}$. Also notice that $C_0 = b$, $C_1 = b - r$ and $C_2 = -2r$. Therefore the long-run average cost per hour incurred by the system using Equation (2) is

$$C = \frac{2(\gamma b - \beta r)(\gamma + \beta)}{2\gamma^2 + 2\gamma\beta + \beta^2}.$$

Performance Evaluation

Here we only consider the average costs case. An important realization to make is that the costs C_i do not have to be dollar costs. They are just the rates at which something is incurred when the source is in state i or even any time-averaged measure that depends on the sojourn time in state i. This is extremely useful in performance analysis and evaluation of systems. We illustrate that using examples.

Example 3 *Consider a telephone switch that can handle at most N calls simultaneously. Assume that calls arrive according to a Poisson process with parameter λ to the switch. Any call arriving when there are N other calls in progress receives a busy signal (and hence rejected). Each accepted call lasts for an exponential amount of time with mean $1/\mu$ amount of time.*

Let $X(t)$ be the number of on-going calls in the switch at time t. Clearly, the state space is $S = \{0, 1, \ldots, N\}$. The infinitesimal generator matrix is

$$Q = \begin{bmatrix} -\lambda & \lambda & 0 & 0 & \ldots & 0 \\ \mu & -(\lambda+\mu) & \lambda & 0 & \ldots & 0 \\ 0 & 2\mu & -(\lambda+2\mu) & \lambda & \ldots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \ldots & N\mu & -N\mu \end{bmatrix}.$$

Now, let $(p_0, p_1, \ldots, p_N)$ be the solution to $[p_0\ p_1\ \ldots\ p_N]Q = [0\ 0\ \ldots\ 0]$ and $p_0 + p_1 + \ldots + p_N = 1$. We do not write down the expression for p_i but assume that this can easily be done by the reader.

Say we are interested in finding out the rate of call rejection. Then we have $C_N = \lambda$ (since in state N calls are rejected at rate λ per unit time) and $C_i = 0$ for $i = 0, 1, \ldots, N-1$ (since no calls are rejected when there are less than N in system). Therefore the long-run average number of calls rejected per unit time is $\sum_{i=0}^{N} C_i p_i = \lambda p_N$. We can also obtain the long-run time-averaged number of lines utilized as $\sum_{i=0}^{N} i p_i$.

Example 4 *The unslotted aloha protocol for multi-access communication works as follows. Consider a system where messages arrive according to a Poisson process with parameter λ. As soon as a message arrives, it attempts transmission. The message transmission times are exponentially distributed with mean $1/\mu$ units of time. If no other message tries to transmit during the transmission time of this message, the transmission is successful. If any other message tries to transmit during this transmission, a collision results and all transmissions are terminated instantly. All messages involved in a collision are called backlogged and are forced to retransmit. All backlogged messages wait for an exponential amount of time (with mean $1/\theta$) before starting retransmission.*

Let $X(t)$ denote the number of backlogged messages at time t and $Y(t)$ be a binary variable that denotes whether or not a message is under transmission at time t. Then the process $\{(X(t), Y(t)), t \geq 0\}$ is a CTMC. Let $p = (p_{00}, p_{01}, p_{10}, p_{11}, p_{20}, p_{21}, \ldots)$ be the steady-state probabilities, i.e. the solution to $pQ = 0$ and $\sum_{i,j} p_{ij} = 1$. Say, we are interested in obtaining the system throughput, i.e. the average number of successful transmissions per unit time. The rate of successful transmission is μ whenever $Y(t)$ is 1 and zero when $Y(t)$ is 0. Also when $Y(t)$ is 1, it is not necessary that a transmission would occur, that happens only with probability $\mu/(\mu + i\theta + \lambda)$ when $X(t)$ is i. Therefore the throughput can be computed as $\sum_{i=0}^{\infty} p_{i1}\mu^2/(\mu + i\theta + \lambda)$. Interestingly, there is a much simpler way to obtain the throughput. If the system is stable, then in steady state the throughput must indeed be λ, using a conservation of flow argument. To obtain the time-averaged number of backlogged messages we can compute $\sum_{i=0}^{\infty} i(p_{i0} + p_{i1})$.

References

[1] V.G. Kulkarni. *Modeling and Analysis of Stochastic Systems.* Texts in Statistical Science Series. Chapman and Hall, Ltd., London, 1995.

[2] S.M. Ross. *Introduction to Probability Models.* Academic Press, San Diego, CA, 2003.