Lecture-24: Reversed Processes

1 M/M/1 queue

The M/M/1 queue is the simplest and most studied models of queueing systems. We assume a continuous-time queueing model with following components.

(a) There is a single queue for waiting that can accommodate arbitrarily large number of customers.
(b) Arrivals to the queue occur according to a Poisson process with rate $\lambda > 0$. That is, let A_n be the arrival instant of the nth customer, then the sequence of inter-arrival times ξ is *i.i.d.* exponentially distributed with rate λ.
(c) There is a single server and the service time of nth arrival is denoted by a random variable σ_n. The sequence of service times $\sigma : \Omega \to \mathbb{R}_+^{\mathbb{N}}$ is *i.i.d.* exponentially distributed with rate $\mu > 0$, independent of the Poisson arrival process.
(d) We assume that arrivals join the tail of the queue, and hence begin service in the order that they arrive *first-in-queue-first-out* (FIFO).

Let L_t denote the number of entities in the system at time $t \in \mathbb{R}_+$, where "system" means the queue plus the service area. For example, $L_t = 2$ means that there is one entity in service and one waiting in line.

1.1 Transition rates

Since the inter-arrival and the service times are memoryless, the residual time for next arrival Y_t^A is identically distributed to ξ_1 and independent of past $\mathcal{F}_t$ and residual service time Y_t^S for entity in service is identically distributed to σ_1 and independent of past $\mathcal{F}_t$. We observe that L_t remains unchanged in the time $t + [0, \min\{Y_t^A, Y_t^S\})$ for $L_t \geqslant 1$. In particular, L_t can have a unit increase if $Y_t^A < Y_t^S$ corresponding to an arrival, and a unit decrease for $L_t \geqslant 1$ if $Y_t^A < Y_t^S$ corresponding to a departure. If $L_t = 0$, there can be no service and L_t remains 0 until $t + Y_t^A$, and has a unit increase at time $t + Y_t^A$. It follows that $L : \Omega \to \mathbb{Z}_+^{\mathbb{R}_+}$ is a right continuous process with left limits, and is piece-wise constant. Since inter-arrival and service times are independent, it follows that sojourn time is exponentially distributed with rate $\lambda + \mu$ for $L_t > 0$ and rate λ for $L_t = 0$. The probability of increase in L_t after the end of sojourn time, is unity for $L_t = 0$, and $P\{Y_t^A < Y_t^S\} = \frac{\lambda}{\lambda+\mu}$ for $L_t > 0$. It follows that L is a time homogeneous CTMC, and we can write the corresponding generator matrix as

$$Q(n,m) = \lambda \mathbb{1}_{\{m-n=1\}} + \mu \mathbb{1}_{\{n-m=1, m \geqslant 0\}}.$$

We observe that $Q(n,n) = -(\lambda + \mu)$ for $n \in \mathbb{N}$ and $Q(0,0) = -\lambda$. It follows that $M/M/1$ queue occupancy is an irreducible CTMC.

1.2 Equilibrium distribution and reversibility

Recall that the system load $\rho \triangleq \frac{\mathbb{E}\sigma_1}{\mathbb{E}\xi_1} = \frac{\lambda}{\mu} < 1$ for a stable queue. We can find the invariant distribution $\pi \in \mathcal{M}(\mathbb{Z}_+)$ of time homogeneous CTMC L, by solving the global balance equation $\pi = \pi Q$ which gives

$$\pi_{n-1} Q_{n-1,n} + \pi_{n+1} Q_{n+1,n} = -\pi_n Q_{nn}, \qquad \pi_1 Q_{1,0} = -\pi_0 Q_{00}.$$

Taking the discrete Fourier transform $\Pi(z) = \sum_{n \in \mathbb{Z}_+} z^n \pi_n$ of the distribution π, we get

$$z\lambda \Pi(z) + z^{-1} \mu (\Pi(z) - \pi(0)) = (\lambda + \mu)\Pi(z) - \mu \pi(0).$$

Since $z \neq 1$, we obtain $\Pi(z) = \frac{\pi(0)}{(1-z\rho)} = \pi_0 \sum_{n \in \mathbb{Z}_+} \rho^n z_n$, and it follows that $\pi_n = \pi_0 \rho^n$ for each $n \in \mathbb{Z}_+$. Since $\sum_{n \in \mathbb{Z}_+} \pi(n) = 1$, we obtain $\pi_0 = 1 - \rho$ for $\rho < 1$.

Example 1.1 ($M/M/1$ queue). From the generator matrix for the number of entities $L : \Omega \to \mathbb{Z}_+^{\mathbb{R}}$ in an M/M/1 queue, we observe that it is a birth-death process. Hence, this time homogeneous CTMC is time-reversible at stationarity, with the equilibrium distribution $\pi \in \mathcal{M}(\mathbb{Z}_+)$ satisfying the detailed balance equations $\pi_n \lambda = \pi_{n+1}\mu$ for each $n \in \mathbb{Z}_+$. This yields $\pi_{n+1} = \rho \pi_n$ for the system load $\rho = \frac{\mathbb{E}\sigma_1}{\mathbb{E}\xi_1} = \frac{\lambda}{\mu}$. Since $\sum_{n \in \mathbb{Z}_+} \pi_n = 1$, we must have $\rho < 1$, such that $\pi_n = (1-\rho)\rho^n$ for each $n \in \mathbb{Z}_+$. In other words, if $\lambda < \mu$, then the equilibrium distribution of the number of customers in the system is geometric with parameter $\rho = \frac{\lambda}{\mu}$. We say that the M/M/1 queue is in the *stable* regime when $\rho < 1$.

Corollary 1.2. *The number of customers in a stable M/M/1 queueing system at equilibrium is a reversible Markov process.*

Theorem 1.3 (Burke). *Departures from a stable M/M/1 queue are Poisson with same rate as the arrivals.*

Exercise 1.4. Directly characterize the departure process from a stable $M/M/1$ queue at stationarity.

2 Reversed Processes

Definition 2.1. Let $X : \Omega \to \mathcal{X}^T$ be a stochastic process with index set T being an additive ordered group such as $\mathbb{R}$ or $\mathbb{Z}$. Then, $\hat{X}^\tau : \Omega \to \mathcal{X}^T$ defined as $\hat{X}^\tau_t \triangleq X_{\tau - t}$ for all $t \in T$ is the **time-reversed process** of X for some $\tau \in T$.

Remark 1. Note that a reversed process, doesn't have to have the identical distribution to the original process. For a reversible process X, the reversed process would have identical distribution.

Lemma 2.2. *If $X : \Omega \to \mathcal{X}^T$ is a Markov process, then the reversed process $\hat{X}^\tau$ is also Markov for any $\tau \in T$.*

Proof. Let $\mathcal{F}_\bullet$ be the natural filtration of the process X. From the Markov property of process X, for any future event $F \in \sigma(X_u : u > t)$, past event $H \in \sigma(X_s : s < t)$, states $x, y \in \mathcal{X}$, and times $u, s > 0$, we have

$$P(F \mid \{X_t = y\} \cap H) = P(F \mid \{X_t = y\}).$$

Markov property of the reversed process follows from the observation, that

$$P(H \mid \{X_t = y\} \cap F) = \frac{P(H \cap \{X_t = y\}) P(F \mid H \cap \{X_t = y\})}{P\{X_t = y\} P(F \mid \{X_t = y\})} = P(H \mid \{X_t = y\}).$$

□

Remark 2. Even if the forward process X is time-homogeneous, the reversed process need not be time-homogeneous. For a non-stationary time-homogeneous Markov process, the reversed process is Markov but not necessarily time-homogeneous.

Theorem 2.3. *If $X : \Omega \to \mathcal{X}^{\mathbb{R}}$ is an irreducible, positive recurrent, stationary, and time-homogeneous Markov process with transition kernel $P : \mathbb{R}_+ \to \mathcal{M}(\mathcal{X})^{\mathcal{X}}$ and invariant distribution $\pi \in \mathcal{M}(\mathcal{X})$, then the reversed process $\hat{X}^\tau : \Omega \to \mathcal{X}^{\mathbb{R}}$ is also irreducible, positive recurrent, stationary, and time-homogeneous Markov with the same invariant distribution π nd a transition kernel $\hat{P} : \mathbb{R}_+ \to \mathcal{M}(\mathcal{X})^{\mathcal{X}}$ defined as*

$$\hat{P}_{xy}(t) \triangleq \frac{\pi_y}{\pi_x} P_{yx}(t) \text{ for all } t \in \mathbb{R}_+, \text{ and states } x, y \in \mathcal{X}. \tag{1}$$

Further, for any finite sequence of states $x \in \mathcal{X}^n$, finite sequence of times $t \in \mathbb{R}_+^n$ such that $t_1 < t_2 < \cdots < t_n$, and any shift $\tau \in \mathbb{R}$, we have $P_\pi\Big(\cap_{i=1}^n \{X_{t_i} = x_i\}\Big) = \hat{P}_\pi\Big(\cap_{i=1}^n \Big\{\hat{X}^\tau_{t_i} = x_i\Big\}\Big)$.

Proof. We observe that $\hat{X}^\tau_{t_i} = X_{\tau - t_i}$ for all $i \in [n]$.

Step 1: We verify that $\hat{P}$ defined in (1) is a probability transition kernel. This follows from the fact that (a) $\hat{P}_{xy}(t) \geqslant 0$ for all $t \in \mathbb{R}_+$, and (b) $\sum_{y \in \mathcal{X}} \hat{P}_{xy}(t) = \frac{1}{\pi_x} \sum_{y \in \mathcal{X}} \pi_y P_{yx}(t) = 1$.

Step 2: We verify that π is an invariant distribution for $\hat{P}$, since $\sum_{x\in\mathcal{X}} \pi_x \hat{P}_{xy}(t) = \pi_y \sum_{x\in\mathcal{X}} P_{yx}(t) = \pi_y$, for all states $y \in \mathcal{X}$.

Step 3: We next verify that $\hat{P}$ defined in (1) is the probability transition kernel for the reversed process $\hat{X}^\tau$. Since the forward process is stationary and time-homogeneous, we can write the probability transition kernel for the reversed process as

$$P(\{\hat{X}^\tau_{t+s} = y\} \mid \{\hat{X}^\tau_s = x\}) = \frac{P_\pi\{X_{\tau-t-s} = y, X_{\tau-s} = x\}}{P_\pi\{X_{\tau-s} = x\}} = \frac{\pi_y P_{yx}(t)}{\pi_x}.$$

This implies that the reversed process is time-homogeneous and has the probability transition kernel $\hat{P}$.

Step 4: From Step 2 and Step 3, it follows that π is the invariant distribution for the reversed process. From the positive recurrence of forward process, it follows that $\pi_x > 0$ for all states $x \in \mathcal{X}$. Hence, this shows the positive recurrence of the reversed process as well. From the stationarity of forward process and definition of reversed process, it follows that the marginal distribution for the reversed process at any time t is π.

Step 5: We next verify the irreducibility of $\hat{P}$. Let $x, y \in \mathcal{X}$. From irreducibility of P, there exists $t \in \mathbb{R}_+$ such that $P_{yx}(t) > 0$. From positive recurrence of P, we have $\pi_x > 0$ for each state $x \in \mathcal{X}$. From the definition of $\hat{P}$ in (1), it follows that $\hat{P}_{xy}(t) > 0$ and hence $x \to y$. Since the choice of states $x, y \in \mathcal{X}$ was arbitrary, the irreducibility of the reversed process follows.

Step 6: Finally, we verify that $P_\pi\Big(\cap_{i=1}^n \{X_{t_i} = x_i\}\Big) = \hat{P}_\pi\Big(\cap_{i=1}^n \Big\{\hat{X}^\tau_{t_i} = x_i\Big\}\Big)$. From the Markov property of the underlying processes and definition of $\hat{P}$, we can write $P_\pi\Big(\cap_{i=1}^n \{X_{t_i} = x_i\}\Big)$ as

$$\pi_{x_1} \prod_{i=1}^{n-1} P_{x_i x_{i+1}}(t_{i+1} - t_i) = \pi_{x_n} \prod_{i=1}^{n-1} \hat{P}_{x_{i+1} x_i}((\tau - t_i) - (\tau - t_{i+1})) = \hat{P}_\pi\Big(\cap_{i=1}^n \Big\{\hat{X}^\tau_{t_i} = x_i\Big\}\Big).$$

Step 7: From the stationarity of the forward process X, we see that the joint distributions of $(X_{t_1}, \dots, X_{t_n})$ and $(X_{s+t_1}, \dots, X_{s+t_n})$ are identical for all $s \in T$ and any finite $n \in \mathbb{N}$. It follows that $\hat{X}^\tau$ is also stationary, since $(\hat{X}^\tau_{t_n}, \dots, \hat{X}^\tau_{t_1})$ and $(\hat{X}^\tau_{s+t_n}, \dots, \hat{X}^\tau_{s+t_1})$ have the identical distribution.

□

Remark 3. There is a subtle difference between reversed process and reversible process. Reversed process has a different evolution probabilities than the forward process, whereas for the reversible process the evolution probabilities are identical.

2.1 Reversed Markov Chain

Corollary 2.4. *If $X : \Omega \to \mathcal{X}^\mathbb{Z}$ is an irreducible, positive recurrent, stationary, time-homogeneous Markov chain with transition matrix $P \in \mathcal{M}(\mathcal{X})^\mathcal{X}$ and invariant distribution $\pi \in \mathcal{M}(\mathcal{X})$, then the reversed chain $\hat{X}^\tau : \Omega \to \mathcal{X}^\mathbb{Z}$ is an irreducible, positive recurrent, stationary, time-homogeneous Markov chain with the same invariant distribution π, and transition matrix $\hat{P}$ defined as $\hat{P}_{xy} \triangleq \frac{\pi_y}{\pi_x} P_{yx}$, for all states $x, y \in \mathcal{X}$.*

Corollary 2.5. *Consider an irreducible Markov chain $X : \Omega \to \mathcal{X}^\mathbb{Z}$ with transition matrix $P \in \mathcal{M}(\mathcal{X})^\mathcal{X}$. If one can find a positive distribution $\alpha \in \mathcal{M}(\mathcal{X})$ and other transition matrix $P^* \in \mathcal{M}(\mathcal{X})^\mathcal{X}$ that satisfies the detailed balance equation*

$$\alpha_x P_{xy} = \alpha_y P^*_{yx} \tag{2}$$

for all states $x, y \in \mathcal{X}$, then P^* is the transition matrix for the reversed chain and α is the invariant distribution for both chains.

Proof. Summing both sides of the detailed balance equation (2) over x, we obtain that α is the invariant distribution of the forward chain. Since $P^*_{yx} = \frac{\alpha_x P_{xy}}{\alpha_y}$, it follows that $P^* \in \mathcal{M}(\mathcal{X})^\mathcal{X}$ is the transition matrix of the the reversed chain and α is the invariant distribution of the reversed chain. □

2.2 Reversed Markov Process

Corollary 2.6. *If $X : \Omega \to \mathcal{X}^\mathbb{R}$ is an irreducible, positive recurrent, stationary, time-homogeneous Markov process with generator matrix Q and invariant distribution π, then the reversed process $\hat{X}^\tau : \Omega \to \mathcal{X}^\mathbb{R}$ is also an irreducible, positive recurrent, stationary, time-homogeneous Markov process with same invariant distribution π and generator matrix $\hat{Q}$ defined as $\hat{Q}_{xy} \triangleq \frac{\pi_y}{\pi_x} Q_{yx}$ for all states $x, y \in \mathcal{X}$.*

Corollary 2.7. *Let $Q \in \mathbb{R}^{\mathcal{X}\times\mathcal{X}}$ denote the rate matrix for an irreducible Markov process. If we can find $Q^* \in \mathbb{R}^{\mathcal{X}\times\mathcal{X}}$ and a positive distribution $\pi \in \mathcal{M}(\mathcal{X})$ such that for $y \neq x \in \mathcal{X}$, we have*

$$\pi_x Q_{xy} = \pi_y Q^*_{yx}, \qquad \text{and} \qquad \sum_{y\neq x} Q_{xy} = \sum_{y\neq x} Q^*_{xy},$$

then Q^* is the rate matrix for the reversed Markov process and π is the invariant distribution for both processes.

3 Applications of Reversed Processes

3.1 Truncated Markov Processes

Definition 3.1. For a Markov process $X : \Omega \to \mathcal{X}^{\mathbb{R}}$, and a subset $A \subseteq \mathcal{X}$ the boundary of A is defined as

$$\partial A \triangleq \left\{ y \notin A : Q_{xy} > 0, \text{ for some } x \in A \right\}.$$

Example 3.2. Consider a birth-death process. Let $A = \{3,4\}$. Then, $\partial A = \{2,5\}$.

Definition 3.3. Consider a transition rate matrix $Q \in \mathbb{R}^{\mathcal{X}\times\mathcal{X}}$ on the countable state space $\mathcal{X}$. Given a nonempty subset $A \subseteq \mathcal{X}$, the truncation of Q to A is the transition rate matrix $Q^A \in \mathbb{R}^{A\times A}$, where for all states $x, y \in A$

$$Q^A_{xy} \triangleq \begin{cases} Q_{xy}, & y \neq x, \\ -\sum_{z\in A\setminus\{x\}} Q_{xz}, & y = x. \end{cases}$$

Proposition 3.4. *Suppose $X : \Omega \to \mathcal{X}^{\mathbb{R}}$ is an irreducible, time-reversible Markov process on the countable state space $\mathcal{X}$, with generator matrix $Q \in \mathbb{R}^{\mathcal{X}\times\mathcal{X}}$ and positive invariant distribution $\pi \in \mathcal{M}(\mathcal{X})$. Suppose the truncated Markov process X^A to a set of states $A \subseteq \mathcal{X}$ with generator matrix Q^A is irreducible. Then, $X^A : \Omega \to A^{\mathbb{R}}$ at stationarity is time-reversible, with positive invariant distribution $\pi^A \in \mathcal{M}(A)$ defined for all $y \in A$, as*

$$\pi^A_y \triangleq \frac{\pi_y}{\sum_{x\in A} \pi_x}.$$

Proof. It is clear that π^A is a distribution on state space A. We must show the reversibility with this distribution π^A. That is, we must show $\pi^A_x Q_{xy} = \pi^A_y Q_{yx}$ for all states $x, y \in A$. However, this is true since the original chain is time reversible. □

Example 3.5 (Limiting waiting room: M/M/1/K). Consider a variant of the $M/M/1$ queueing system with load $\rho \triangleq \frac{\lambda}{\mu}$ that has a finite buffer capacity of at most K customers. Thus, customers that arrive when there are already K customers present are *rejected*. It follows that the CTMC for this system is simply the $M/M/1$ CTMC truncated to the state space $\{0, 1, \dots, K\}$, and so it must be time-reversible with invariant distribution

$$\pi_i = \frac{\rho^i}{\sum_{j=0}^{k} \rho^j}, \quad 0 \leqslant i \leqslant k.$$

Example 3.6 (Two queues with joint waiting room). Consider two independent $M/M/1$ queues with arrival and service rates λ_i and μ_i respectively for $i \in [2]$. Then, the joint distribution of two queues is

$$\pi(n_1, n_2) = (1-\rho_1)\rho_1^{n_1}(1-\rho_2)\rho_2^{n_2}, \quad n_1, n_2 \in \mathbb{Z}_+.$$

Suppose both the queues are sharing a common waiting room, where if arriving customer finds R waiting customer then it leaves. Defining $A \triangleq \left\{ n \in \mathbb{Z}^2_+ : n_1 + n_2 \leqslant R \right\}$, we observe that the joint Markov process is restricted to the set of states A, and the invariant distribution for the truncated Markov process is

$$\pi(n_1, n_2) = \frac{\rho_1^{n_1}\rho_2^{n_2}}{\sum_{(m_1,m_2)\in A} \rho_1^{m_1}\rho_2^{m_2}}, \quad (n_1, n_2) \in A.$$