Week 11

Limiting probabilities (examples)

Birth and death processes. We consider a birth and death process with birth rates λ_i and death rates μ_i . We assume $\mu_i > 0$ for every $i \ge 1$. If π is a stationary distribution, we must have $\pi Q = 0$, where

$$Q = \begin{pmatrix} -\lambda_0 & \lambda_0 & 0 & 0 & \dots \\ \mu_1 & -\lambda_1 - \mu_1 & \lambda_1 & 0 & \dots \\ 0 & \mu_2 & -\lambda_2 - \mu_2 & \lambda_2 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

This means that we need $-\lambda_0 \pi_0 + \mu_1 \pi_1 = 0$ and, for all $i \ge 1$:

$$\lambda_{i-1}\pi_{i-1} + (-\lambda_i - \mu_i)\pi_i + \mu_{i+1}\pi_{i+1} = 0.$$

From here, we try to express everything in term of π_0 . Using the stationarity equation for i = 0, we find

$$\pi_1 = \frac{\lambda_0}{\mu_1} \pi_0.$$

Next, using the stationarity equation for i = 1, we have

$$\lambda_0 \pi_0 + \mu_2 \pi_2 = (\lambda_1 + \mu_1) \pi_1,$$

so $\mu_2 \pi_2 = \lambda_1 \pi_1$ (since $\lambda_0 \pi_0 = \mu_1 \pi_1$), so $\pi_2 = \frac{\lambda_0 \lambda_1}{\mu_1 \mu_2} \pi_0$. Similarly, using the stationarity equation for i = 2, we get $\lambda_1 \pi_1 + \mu_3 \pi_3 = (\lambda_2 + \mu_2) \pi_2$, so $\mu_3 \pi_3 = \lambda_2 \pi_2$, so $\pi_3 = \frac{\lambda_0 \lambda_1 \lambda_2}{\mu_1 \mu_2 \mu_3} \pi_0$. In general, we have

$$\pi_i = r_i \pi_0$$
, where $r_i = \frac{\lambda_0 \lambda_1 \dots \lambda_{i-1}}{\mu_1 \mu_2 \dots \mu_i}$.

Finally, we must have

$$1 = \sum_{i \ge 0} \pi_i = \left(1 + \sum_{i \ge 1} r_i\right) \pi_0.$$

Hence, if $\sum_{i\geq 1} r_i = +\infty$, then there is no stationary distribution. If the series $\sum_{i\geq 1} r_i$ is finite, we have found the stationary distribution:

$$\pi_0 = \frac{1}{1 + \sum_{i \ge 1} r_i}$$
 and $\pi_n = \frac{r_n}{1 + \sum_{i \ge 1} r_i}$ for $n \ge 1$.

M/M/1 queue. We recall that in the M/M/1 queue, we have a unique server, customers arriving at a certain rate λ , and the service times are i.i.d. $Exp(\mu)$. This is a particular case of birth and death process with $\lambda_n = \lambda$ for $n \ge 0$ and $\mu_n = \mu$ for $n \ge 1$. Hence, we have $r_i = \frac{\lambda^i}{\mu^i}$. If $\lambda < \mu$, then

$$1 + \sum_{i \ge 1} r_i = \sum_{i \ge 0} \left(\frac{\lambda}{\mu}\right)^i = \frac{1}{1 - \lambda/\mu},$$

so the stationary distribution is $\pi_n = \left(\frac{\lambda}{\mu}\right)^n \left(1 - \frac{\lambda}{\mu}\right)$. This means that in the long run, the law of the number of customers is a geometric variable with parameter $1 - \frac{\lambda}{\mu}$. In particular, in the long run, the average number of customers in the system is $\frac{1}{1-\lambda/\mu}$.

On the other hand, if $\lambda \ge \mu$, then $\sum_{i\ge 1} r_i = +\infty$, so there is no stationary distribution, so there are no limiting probability. More precisely, for all i, we have $\mathbb{P}(X(t) = i) \to 0$ when $t \to +\infty$.

Linear growth model. This is another particular case of birth-death process, with $\lambda_n = n\lambda + \theta$ and $\mu_n = n\mu$. We get

$$r_i = \frac{\theta(\theta + \lambda)(\theta + 2\lambda)\dots(\theta + (i-1)\lambda)}{i!\,\mu^i}.$$

To know if the series $\sum_{i\geq 1} r_i$ converge or not, we use the *ratio test*: we compute the limit $\lim_{i\to+\infty} \frac{r_{i+1}}{r_i}$. If this limit is < 1, the series converge. If the limit is > 1, the series diverge. If the limit is 1, we don't know for sure. We have

$$\frac{r_{i+1}}{r_i} = \frac{\lambda_i}{\mu_{i+1}} = \frac{\theta + i\lambda}{i\mu} \xrightarrow[i \to +\infty]{} \frac{\lambda}{\mu}.$$

In particular, once again, if $\lambda < \mu$, the series converge and we have limiting probabilities. If $\lambda > \mu$, the series diverge and the probabilities go to 0. Note that we have already computed $\mathbb{E}[X(t)]$ and seen that it has a finite limit in the first case, and goes to infinity in the second.

5) The embedded Markov chain

If we forget the lengths between jumps in a continuous-time Markov chain X, we get a discrete-time Markov chain Y. Our goal in this paragraph is to understand the link between X and Y.

More precisely, the discrete chain Y obtained by ignoring the time spent between jumps is called the *embedded chain*. Assume that X has a stationary distribution π , and that the discrete-time Markov chain Y has a stationary distribution σ , that is,

$$\sigma_i = \sum_j p_{ji} \sigma_j \text{ for all } i.$$

What is the relation between σ and π ?

Intuition: σ_i is the proportion of the jumps that end up on *i*. Once there, the time spent on *i* is $Exp(v_i)$, so $\frac{1}{v_i}$ in average, so we expect π to be proportional to $\frac{\sigma_i}{v_i}$. Since we must have $\sum_i \pi_i = 1$, we can guess

$$\pi_i = \frac{\sigma_i/v_i}{Z}$$
, where $Z = \sum_j \sigma_j/v_j$.

Check: We now check the π that we guessed is a stationary distribution. First, by definition of Z, we have

$$\sum_{i} \pi_i = \frac{1}{Z} \sum_{i} \frac{\sigma_i}{v_i} = \frac{Z}{Z} = 1.$$

Second, to make sure π is stationary, we need to check

$$v_j \pi_j = \sum_{k \neq j} q_{kj} \pi_k$$
 for all j

For this, we write

$$\sum_{k \neq j} q_{kj} \pi_k = \sum_{k \neq j} v_k p_{kj} \times \frac{1}{Z} \frac{\sigma_k}{v_k} = \frac{1}{Z} \sum_{k \neq j} p_{kj} \sigma_k = \frac{1}{Z} \sigma_j$$

because σ is stationary for the embedded chain. On the other hand, we also have $v_j \pi_j = v_j \times \frac{1}{Z} \frac{\sigma_j}{v_j} = \frac{\sigma_j}{Z}$, which proves that π is stationary for X.

6) Reversibility

In this part, we assume that X is irreducible and has a stationary distribution π . We also assume that X(0) was distributed according to π , so that for all t, the law of X(t) is π .

The reversed process. We now fix some large t and consider the reverse chain, that is $\overline{X}(s) = X(t-s)$ for $0 \le s \le t$. We first claim that, conditionally on $\overline{X} = 0$, the time spent by \overline{X} on state i is $Exp(v_i)$, just like for X. Indeed, we have

$$\mathbb{P}\left(\overline{X}\text{stays on } i \text{ at least until time } s | \overline{X}(0) = i\right)$$

$$= \mathbb{P}\left(X \text{ stays on } i \text{ from time } t - s \text{ to } t | X(t) = i\right)$$

$$= \frac{\mathbb{P}\left(X(s) = t - s\right) \mathbb{P}\left(X \text{ stays on } i \text{ from time } t - s \text{ to } t | X(s) = t - s\right)}{\mathbb{P}\left(X(t - s) = i\right)}$$

$$= \frac{\pi_i e^{-v_i s}}{\pi_i} = e^{-v_i s}.$$

Therefore, the jump rate of \overline{X} on each state is the same as for X. Moreover, the embedded chain of \overline{X} is the time-reversal of the embedded chain of X. By the results from the first chapter of the course, the embedded chain of \overline{X} has transitions

$$\frac{\sigma_j}{\sigma_i} p_{ji}.$$

In particular \overline{X} is the same as X if and only if these are the same as p_{ij} , i.e. if $\sigma_i p_{ij} = \sigma_j p_{ji}$ for all i and j. We recall that $\pi_i = \frac{1}{Z} \frac{\sigma_i}{v_i}$, so $\sigma_i = Z v_i \pi_i$, so this is equivalent to

$$Zv_i\pi_ip_{ij} = Zv_j\pi_jp_{ji},$$

which is equivalent to $\pi_i q_{ij} = \pi_j q_{ji}$.

Definition 1 We say that X is time-reversible if it has a stationary distribution π and

$$\pi_i q_{ij} = \pi_j q_{ji}$$

for all states i and j.

Like in the discrete-time case, the left-hand side can be interpreted as the frequency at which the chain jumps from i to j, and the right-hand side as the frequency at which the chain jumps from j to i.

Birth and death processes.

Proposition 2 Any birth and death process with a stationary distribution is time-reversible.

Instead of a proof, we just give an intuitive argument why this is true: between two transitions from i to i + 1, the process must jump from i + 1 to i exactly once. Hence, in the long run, it will jump as many times from i to i + 1 as from i + 1 to i. Moreover, if |i - j| > 1, it is impossible to jump from i to j or j to i, so both sides will be 0.

The M/M/s queue (i.e. with s servers). We recall the definition of the model: we have s servers. New customers arrive at rate λ . When n customers are present, they leave at rate $n\mu$ if $n \leq s$, and $s\mu$ if $n \geq s$ (i.e. if all servers are busy).

Corollary 3 For the M/M/s queue, if $\lambda < s\mu$, the departures of the clients form a Poisson process with rate λ .

Proof. Let X(t) be the number of customers present at time t. This is a birth and death process, so by the last Proposition this is reversible. We know that by definition of the process, the times where X increases by 1 form a Poisson process with rate λ . When we reverse the process, these times become the times where X(t) decreases by 1, i.e. the times where a client leaves. Hence, the times where a client leaves also form a Poisson process with rate λ .