Mathematics for Informatics 4a

José Figueroa-O'Farrill

Lecture 20

30 March 2012

The story of the film so far...

- We are studying **continuous-time Markov processes**, particularly those which are (temporally) **homogeneous**
- Examples are the **counting processes** $\{N(t) \mid t \geqslant 0\}$, of which an important special case are the **Poisson processes**, where $N(t)$ is Poisson distributed with mean λt
- **Inter-arrival times** in a Poisson process are exponential, **waiting times** are "gamma" distributed and **time of occurrence** is uniformly distributed
- Continuous-time Markov chains are characterised by
  1. a transition matrix $[p_{ij}]$, which for all i obeys
     - $p_{ii} = 0$
     - $\sum_j p_{ij} = 1$
  2. the exponential transition rates ν_i
- Poisson process: states $\{0, 1, 2, \ldots\}$, $p_{ij} = 0$ for $j \neq i + 1$ and $p_{i,i+1} = 1$, and all states have equal transition rates

Further properties of exponential r.v.s (I)

- Because of the important rôle played by exponential random variables in continuous-time Markov process, we record here some further properties
- In the previous lecture we showed that if a continuous random variable is memoryless, then it is exponential
- In Lecture 13 we showed that the sum of two i.i.d. exponential variables is a "gamma" distribution, and in Lecture 14 we saw this held for any number of i.i.d. exponential variables
- The sum $Z = X + Y$ of two independent exponential variables with different rates is **hypoexponential**:

$$f_X(x) = \lambda e^{-\lambda x} \qquad f_Y(y) = \mu e^{-\mu y}$$

$$\implies f_Z(z) = \frac{\lambda\mu}{\mu - \lambda}\left(e^{-\lambda z} - e^{-\mu z}\right)$$

Further properties of exponential r.v.s (II)

- The sum $Z = X_1 + \cdots + X_n$ of independent exponential variables with different rates is also hypoexponential, but the expression gets increasingly complicated
- However the *minimum* $\min(X_1, \ldots, X_n)$ of independent exponential variables is again exponential with rate equal to the sum of the rates of the X_i
- By induction, it is enough to show prove it for $n = 2$, so let X, Y be independent exponential variables with rates λ, μ
- With $U = \min(X, Y)$, $\mathbb{P}(U \leqslant u) = 1 - \mathbb{P}(U > u)$, but

$$\begin{aligned}\mathbb{P}(U > u) = \mathbb{P}(X > u, Y > u) &= \int_u^\infty \int_u^\infty f(x, y)\,dx\,dy \\ &= \int_u^\infty \int_u^\infty \lambda\mu e^{-\lambda x} e^{-\mu y}\,dx\,dy \\ &= \left(\int_u^\infty \lambda e^{-\lambda x}\,dx\right)\left(\int_u^\infty \mu e^{-\mu y}\,dy\right) = e^{-(\lambda+\mu)u}\end{aligned}$$

Further properties of exponential r.v.s (III)

- The final calculation we will need is $\mathbb{P}(X < Y)$ for X, Y exponential with rates λ, μ
- We calculate it by conditioning on X:

$$\begin{aligned}\mathbb{P}(X < Y) &= \int_0^\infty \mathbb{P}(X < Y \mid X = x) f_X(x) dx \\ &= \int_0^\infty \mathbb{P}(X < Y \mid X = x) \lambda e^{-\lambda x} dx \\ &= \int_0^\infty \mathbb{P}(x < Y) \lambda e^{-\lambda x} dx \\ &= \int_0^\infty e^{-\mu x} \lambda e^{-\lambda x} dx \\ &= \lambda \int_0^\infty e^{-(\lambda+\mu)x} dx \\ &= \frac{\lambda}{\lambda + \mu}\end{aligned}$$

Birth and death processes (I)

- The only allowed transitions in a counting process are those which increase the "population": $n \to n+1$
- They are thus said to be "pure birth" processes
- In a "birth and death" process $\{N(t) \mid t \geqslant 0\}$ we allow transitions $n \to n+1$ (called **births**) and $n \to n-1$ (called **deaths**), but of course $n \geqslant 0$
- Births and deaths are independent and exponentially distributed with rates λ_n and μ_n, respectively, when the population is n
- The parameters $\{\lambda_n \mid n \in \mathbb{N}\}$ and $\{\mu_n \mid n \in \mathbb{N}\}$ are called the **birth rates** and **death rates**, respectively
- A **birth and death process** is a continuous-time Markov process with states $\mathbb{N} = \{0, 1, 2, \dots\}$ for which the allowed transitions are $n \to n+1$ and $n \to n-1$

Birth and death processes (II)

- The transition probabilities are given by $p_{01} = 1$ and

$$p_{n,n+1} = \frac{\lambda_n}{\lambda_n + \mu_n} \qquad p_{n,n-1} = \frac{\mu_n}{\lambda_n + \mu_n} \qquad (n \geqslant 1)$$

- We argue as follows: $p_{n,n+1}$ is the probability that in a population of n a birth occurs before a death, i.e., $\mathbb{P}(B_n < D_n)$, where B_n and D_n are the exponential variables corresponding to a birth and death, respectively, when the population is n.
- Since B_n has rate λ_n and D_n has rate μ_n, the results follows from the earlier discussion
- The transition rates are

$$\nu_0 = \lambda_0 \quad \text{and} \quad \nu_n = \lambda_n + \mu_n \quad (n \geqslant 1)$$

since the time to any transition at population n is $\min(B_n, D_n)$, which is exponential with rate $\lambda_n + \mu_n$

Birth and death processes (III)

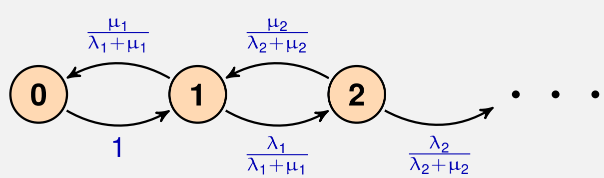

Examples (Pure birth processes)

- **pure birth**: $\mu_n = 0$ for all $n \geqslant 0$
- Poisson: $\mu_n = 0$ and $\lambda_n = \lambda$ for all $n \geqslant 0$
- **Yule**: $\mu_n = 0$ and $\lambda_n = n\lambda$ for all $n \geqslant 0$, corresponding to a Markov process $\{N(t) \mid t \geqslant 0\}$ where $N(t)$ is the size at time t of a population whose members cannot die, and they give birth to new members independently in an exponentially distributed amount of time with rate λ

Example (Linear growth with immigration)

- This is a model in which $\mu_n = n\mu$ and $\lambda_n = n\lambda + \theta$, for $n \geqslant 0$
- Each individual is assumed to give birth at an exponential rate λ
- In addition there is an exponential rate of increase θ of the population due to immigration, so if there are n individuals in the system the total birth rate is $n\lambda + \theta$
- Deaths occur at an exponential rate μ for each member of the population, hence the total death rate for a population of size n is $n\mu$.

A typical question in a birth and death process might be to determine the expectation value $\mathbb{E}(N(t))$ of the size of the population at time t.
Usually one derives a differential equation that $\mathbb{E}(N(t))$ obeys and solves it to determine $\mathbb{E}(N(t))$.

Steady-state distribution

- Recall that regular discrete-time Markov chains have a unique steady-state distribution $\boldsymbol{\pi} = (\pi_n)$, obeying $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{P}$, where $\mathbf{P}$ is the transition matrix which evolves the system one time step.
- In other words, $\boldsymbol{\pi}$ is invariant under (discrete) time translations.
- Some continuous-time Markov chains also have a unique steady-state distribution which is invariant under time translation.
- In other words, $\boldsymbol{\pi} = (\pi_n)$, where $\pi_n(t+s) = \pi_n(t)$, so that is constant in time.
- We will not be concerned with the conditions which guarantee the existence and uniqueness of the steady-state distribution.
- We will assume it exists and is unique and we will show how to find it.

- Let $\{N(t) \mid t \geqslant 0\}$ be a continuous-time Markov chain
- Let $n > 0$ and consider a small time increment δt:
- We compute $\pi_n(t+\delta t) = \mathbb{P}(N(t+\delta t) = n)$ by conditioning on $N(t)$:

$$\begin{aligned}\pi_n(t+\delta t) &= \mathbb{P}(N(t+\delta t) = n \mid N(t) = n)\mathbb{P}(N(t) = n)\\ &\quad + \mathbb{P}(N(t+\delta t) = n \mid N(t) = n+1)\mathbb{P}(N(t) = n+1)\\ &\quad\quad + \mathbb{P}(N(t+\delta t) = n \mid N(t) = n-1)\mathbb{P}(N(t) = n-1)\end{aligned}$$

- Let us focus on one of the conditional probabilities, say, $\mathbb{P}(N(t+\delta t) = n \mid N(t) = n+1)$
- This is the probability that a death occurred in $(t, t+\delta t]$ when the population at time t is $n+1$
- At that population, deaths are exponentially distributed with rate μ_{n+1}, so we want the probability of a death in a time interval of length δt at that rate

- For δt small, this is given by

$$\int_0^{\delta t} \mu_{n+1} e^{-t\mu_{n+1}} dt = 1 - e^{-\mu_{n+1}\delta t} \simeq \mu_{n+1}\delta t$$

- Similarly,

$$\begin{aligned}\mathbb{P}(N(t+\delta t) = n \mid N(t) = n-1) &\simeq \lambda_{n-1}\delta t\\ \mathbb{P}(N(t+\delta t) = n \mid N(t) = n) &\simeq 1 - (\lambda_n + \mu_n)\delta t\end{aligned}$$

- Therefore,

$$\begin{aligned}\pi_n(t+\delta t) = (1 - \delta t(\lambda_n + \mu_n))\pi_n(t) + \mu_{n+1}\delta t \pi_{n+1}(t)\\ + \lambda_{n-1}\delta t \pi_{n-1}(t)\end{aligned}$$

- or, said differently,

$$\frac{\pi_n(t+\delta t) - \pi_n(t)}{\delta t} \simeq \mu_{n+1}\pi_{n+1}(t) + \lambda_{n-1}\pi_{n-1}(t) - (\lambda_n + \mu_n)\pi_n(t)$$

- In the steady state, $\pi_n(t + \delta t) = \pi_n(t)$, whence

$$\mu_{n+1}\pi_{n+1} + \lambda_{n-1}\pi_{n-1} = \lambda_n \pi_n + \mu_n \pi_n \qquad (n \geqslant 1)$$

- **probability flow** $=$ probability $\times$ transition rate
- the above equation is the condition for **zero net flow**

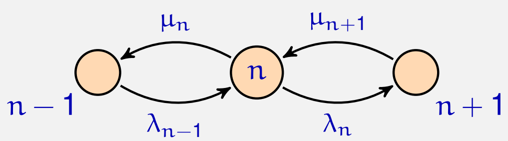

- the "inflow" into state n is $\mu_{n+1}\pi_{n+1} + \lambda_{n-1}\pi_{n-1}$, whereas the "outflow" is $\lambda_n \pi_n + \mu_n \pi_n$
- therefore the steady state is characterised by zero net flow across every state

- We need to pay particular attention to the the zeroth state:

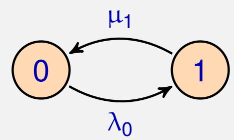

$$\lambda_0 \pi_0 = \mu_1 \pi_1$$

- we rewrite the zero net flow condition for $n \geqslant 1$ as

$$\lambda_{n-1}\pi_{n-1} - \mu_n \pi_n = \lambda_n \pi_n - \mu_{n+1}\pi_{n+1}$$

- which says that the quantity $\lambda_{n-1}\pi_{n-1} - \mu_n \pi_n$ is independent of n
- since it vanishes for $n = 1$, it vanishes for all n, hence the steady state obeys

$$\lambda_n \pi_n = \mu_{n+1}\pi_{n+1} \qquad (n \geqslant 0)$$

- Assuming $\mu_n \neq 0$, we can solve recursively for the π_n in terms of π_0:

$$\pi_1 = \frac{\lambda_0}{\mu_1}\pi_0, \quad \pi_2 = \frac{\lambda_1}{\mu_2}\pi_1 = \frac{\lambda_0\lambda_1}{\mu_1\mu_2}\pi_0, \quad \dots$$

$$\implies \pi_n = \frac{\lambda_0 \cdots \lambda_{n-1}}{\mu_1 \cdots \mu_n}\pi_0$$

- Finally, we solve for π_0 from the normalisation condition $\sum_n \pi_n = 1$, namely

$$\pi_0 \left(1 + \sum_{n \geqslant 1} \frac{\lambda_0 \cdots \lambda_{n-1}}{\mu_1 \cdots \mu_n} \right) = 1$$

- For processes with an infinite number of states, the above series is infinite and convergence is not guaranteed
- Convergence imposes constraints on the birth and death rates for the existence of a steady state

Example (Single server queue)

- Customers arrive at a server according to a Poisson process with rate λ
- Customers are served in exponential time with rate μ
- If the server is idle, customers get served upon arrival, otherwise they join a queue
- The states are labelled by the number $n \in \{0, 1, 2, \dots\}$ of customers in the queue (including anyone being served)
- This is a birth and death process with $\lambda_n = \lambda$ and $\mu_n = \mu$
- If $\lambda > \mu$ customers arrive faster than they are served and the queue keeps growing $\implies$ there is no steady state
- If $\lambda < \mu$, there is a steady state with distribution

$$\pi_n = \frac{\lambda^n}{\mu^n}\pi_0$$

Example (Single server queue — continued)

- The normalisation condition is

$$\pi_0 \sum_{n=0}^{\infty} \frac{\lambda^n}{\mu^n} = 1$$

- As expected, the geometric series converges precisely when $\lambda < \mu$, and

$$\pi_0 \left(\frac{1}{1 - \frac{\lambda}{\mu}} \right) = 1 \implies \pi_0 = 1 - \frac{\lambda}{\mu}$$

- Finally, for all $n \geqslant 1$,

$$\pi_n = \left(1 - \frac{\lambda}{\mu}\right) \left(\frac{\lambda}{\mu}\right)^n$$

Example (Single server queue — continued)

- The steady-state probability generating function is

$$G(s) = \sum_n s^n \pi_n = \sum_{n=0}^{\infty} s^n \frac{\lambda^n}{\mu^n} \left(1 - \frac{\lambda}{\mu}\right) = \frac{1 - \frac{\lambda}{\mu}}{1 - \frac{s\lambda}{\mu}} = \frac{\mu - \lambda}{\mu - s\lambda}$$

  provided that $s < \frac{\mu}{\lambda}$
- The mean length of the queue is the expectation $\mathbb{E}(N)$, given by

$$\mathbb{E}(N) = \sum_n n\pi_n = G'(1) = \frac{\lambda}{\mu - \lambda}$$

  which grows as $\frac{\lambda}{\mu} \to 1$

Summary

- We have discussed **birth and death processes** $\{N(t) \mid t \geqslant 0\}$, with state space $\mathbb{N} = \{0, 1, 2, \dots\}$ and two kinds of transitions:
  1. **birth**: $n \to n + 1$ with rate λ_n
  2. **death**: $n \to n - 1$ with rate μ_n
- **transition probabilities**: $p_{01} = 1$ and

$$p_{n,n+1} = \frac{\lambda_n}{\lambda_n + \mu_n} \qquad p_{n,n-1} = \frac{\mu_n}{\lambda_n + \mu_n} \qquad (n \geqslant 1)$$

- **transition rates**: $\nu_0 = \lambda_0$ and $\nu_n = \lambda_n + \mu_n$ for $n \geqslant 1$
- "Nice" birth and death processes have **steady states** with probabilities (π_n) satisfying the **zero net flow** condition $\lambda_n \pi_n = \mu_{n+1} \pi_{n+1}$ and the normalisation condition $\sum_n \pi_n = 1$