6.262: Discrete Stochastic Processes 2/16/11

Lecture 5: Poisson combining and splitting etc.

Outline:

- Review of Poisson processes
- Combining independent Poisson processes
- Splitting a Poisson process
- Non-homogeneous Poisson processes
- Conditional arrival densities

Review of Poisson processes

A Poisson process is an arrival process with IID exponentially-distributed interarrival times.

It can be represented by its arrival epochs, $S_1, S_2, \dots$, or by its interarrival times, $X_1, X_2, \dots$ or by its counting process, $\{N(t);\, t > 0\}$, where

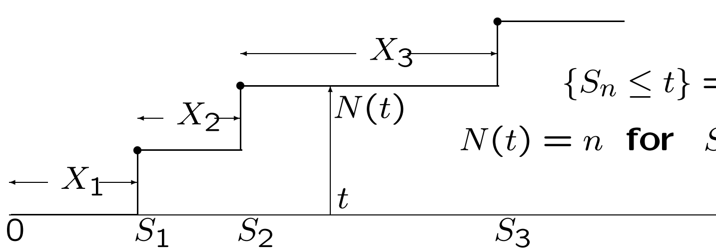

$$\{S_n \le t\} = \{N(t) \ge n\}$$

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

The interarrival times X_i of a Poisson process are memoryless, i.e., for $x, t > 0$,

$$\Pr\{X_i > t + x \mid X_i > t\} = \Pr\{X_i > x\} = \exp(-\lambda x)$$

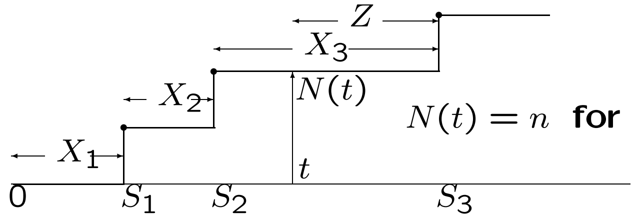

$$N(t) = n \text{ for } S_n \le t < S_{n+1}$$

Given $N(t) = n$ and $S_n = \tau$, we have $X_{n+1} > t - \tau$. The interval $Z = X_{n+1} - (t - \tau)$.

$$\Pr\Big\{Z > z \mid N(t), S_{N(t)}\Big\} = \exp(-\lambda z)$$

Z is independent of $\{N(\tau);\, \tau \le t\}$; Z is the first interarrival of the Poisson process $\{N(t') - N(t);\, t' > t\}$; $\{N(t') - N(t);\, t' > t\}$ is independent of $\{N(\tau);\, \tau \le t\}$.

For any set of times, $0 < t_1 < t_2 < \cdots t_k$, the Poisson process increments, $\{N(t);\, 0 < t \le t_1\}, \{\widetilde{N}(t_1, t);\, t_1 < t \le t_2\}, \dots, \{\widetilde{N}(t_{k-1}, t);\, t_{k-1} < t \le t_k\}$ are stationary and independent Poisson counting processes (over their given intervals). Also,

$$\mathsf{p}_{N(t)}(n) = \frac{(\lambda t)^n \exp(-\lambda t)}{n!}$$

This is a function only of the mean λt. By the stationary and ind. inc. property, we know that $N(t_1)$ and $\widetilde{N}(t_1, t)$ are independent. They are also Poisson and their sum, $N(t)$, is Poisson. In general, sums of independent Poisson rv's are Poisson, with the means adding.

Alternate definitions of Poisson process: (i.e., alternate conditions which suffice to show that an arrival process is Poisson).

Thm: If an arrival process has the stationary and independent increment properties and if $N(t)$ has the Poisson PMF for given λ and all $t > 0$, then the process is Poisson.

Thm: If an arrival process has the stationary and independent increment properties and satisfies the following incremental condition, then the process is Poisson.

$$\Pr\Big\{\widetilde{N}(t, t+\delta) = n\Big\} = \begin{cases} 1 - \lambda\delta + \mathbf{o}(\delta) & \textbf{for} \quad n = 0 \\ \lambda\delta + \mathbf{o}(\delta) & \textbf{for} \quad n = 1 \\ \mathbf{o}(\delta) & \textbf{for} \quad n \geq 2 \end{cases}$$

Combining independent Poisson processes

Two Poisson processes $\{N_1(t);\, t>0\}$ and $\{N_2(t);\, t>0\}$ are independent if for all $t_1, \dots, t_n$, the rv's $N_1(t_1)$, $\dots, N_1(t_n)$ are independent of $N_2(t_1), \dots,, N_2(t_n)$.

Thm: if $\{N_1(t);\, t>0\}$ and $\{N_2(t);\, t>0\}$ are independent Poisson processes of rates λ_1 and λ_2 and $N(t) = N_1(t) + N_2(t)$ for all $t > 0$, then $\{N(t);\, t>0\}$ is a Poisson process of rate $\lambda = \lambda_1 + \lambda_2$.

The idea is that in any increment $(t, t{+}\delta]$,

$$\widetilde{N}(t, t{+}\delta) = \widetilde{N}_1(t, t{+}\delta) + \widetilde{N}_2(t, t{+}\delta)$$

$$\begin{aligned} \mathsf{p}_{\widetilde{N}(t,t+\delta)}(1) &= \mathsf{p}_{\widetilde{N}_1(t,t+\delta)}(1)\mathsf{p}_{\widetilde{N}_2(t,t+\delta)}(0) \\ &+ \mathsf{p}_{\widetilde{N}_1(t,t+\delta)}(0)\mathsf{p}_{\widetilde{N}_2(t,t+\delta)}(1) \\ &= [\delta\lambda_1 + \mathbf{o}(\delta)][1 - \delta\lambda_2 + \mathbf{o}(\delta)] \\ &+ [1 - \delta\lambda_1 + \mathbf{o}(\delta)][\delta\lambda_2 + \mathbf{o}(\delta)] \\ &= \delta(\lambda_1 + \lambda_2) + \mathbf{o}(\delta) \end{aligned}$$

Similarly, $\widetilde{N}(t, t{+}\delta) = 0$ if both $\widetilde{N}_1(t, t{+}\delta) = 0$ and $\widetilde{N}_2(t, t{+}\delta) = 0$. Thus

$$\Pr\left\{\widetilde{N}(t, t{+}\delta) = 0\right\} = [1 - (\lambda_1 + \lambda_2)\delta + \mathbf{o}(\delta)]$$

It is much cleaner analytically to use the Poisson distribution directly. Since $\widetilde{N}_1(t, t{+}\delta)$ and $\widetilde{N}_2(t, t{+}\delta)$ are independent and Poisson,

$$\widetilde{N}(t, t{+}\delta) = \widetilde{N}_1(t, t{+}\delta) + \widetilde{N}_2(t, t{+}\delta)$$

is a Poisson rv with mean $\lambda\delta$.

The sum of many small independent arrival processes tends to be close to Poisson even if the small processes are not. In a sense, the independence between the processes overcomes the dependence between successive arrivals in each process.

Splitting a Poisson process

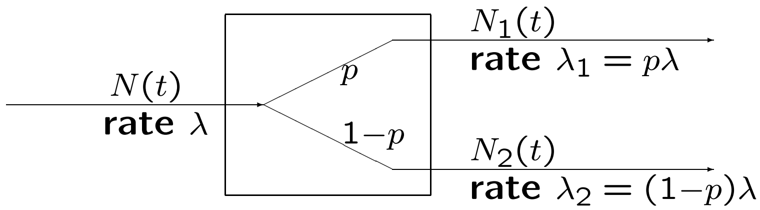

Each arrival is switched to $\{N_1(t); t > 0\}$ with probability p and otherwise goes to $\{N_2(t); t > 0\}$. View the switch as a Bernoulli process independent of $\{N(t); t > 0\}$. A p biased coin is flipped independently at each arrival.

Each new process clearly has the stationary and independent increment property and each satisfies the small increment property. Thus each is Poisson.

The small increment property doesn't make it clear that the split processes are independent. For independence, both processes must sometimes have arrivals in the same small increment. Independence is hidden in the o(δ) terms. See text.

Combining and splitting are often done together. First one views separate independent Poisson processes as a combined process. Then it is split again with binary choices between processes.

Example: Consider a last-come first-serve queue with Poisson arrivals, rate λ and independent exponential services, rate μ. A new arrival starts service immediately, but is interrupted if a new arrival occurs before service completion.

View services as a Poisson process. We can either ignore this process when there is nothing to serve, or visualize a low priority task that is served at rate μ when there is nothing else to do. The arrival process plus the service process is Poisson, rate $\lambda+\mu$.

The probability an arrival completes service before being interrupted is $\mu/(\lambda+\mu)$.

Given that you are interrupted, what is the probability of no further interruption? Two services (the interruping job and you) must finish before the next interruption, so Answer: $\frac{\mu^2}{(\lambda+\mu)^2}$.

Non-homogeneous Poisson processes

Consider optical transmission, where an optical stream of photons is modulated by variable power. The photon stream is reasonably modelled as a Poisson process, and the modulation converts the steady photon rate into a variable rate, say $\lambda(t)$.

We model the number of photons in any interval $(t, t+\delta]$ as a Poisson random variable whose rate parameter over $(t, t+\delta]$ is the average photon rate over $(t, t+\delta]$ times δ.

In the small increment model, we have

$$\Pr\Big\{\widetilde{N}(t, t+\delta) = n\Big\} = \begin{cases} 1 - \delta\lambda(t) + o(\delta) & \text{for} \quad n = 0 \\ \lambda(t)\delta + o(\delta) & \text{for} \quad n = 1 \\ o(\delta) & \text{for} \quad n \geq 2 \end{cases}$$

We can use this small increment model to see that the number of arrivals in each increment is Poisson with

$$\Pr\left\{\widetilde{N}(t,\tau) = n\right\} = \frac{[\widetilde{m}(t,\tau)]^n \exp[-\widetilde{m}(t,\tau)]}{n!}$$

where

$$\widetilde{m}(t,\tau) = \int_t^\tau \lambda(u)\,du.$$

Combining and splitting non-homogeneous processes still works as in the homogeneous case, but the independent exponential interarrivals doesn't work.

We now return to homogeneous Poisson processes.

Conditional arrival densities

There are many interesting and useful results about the increment $(0,t]$ of a Poisson process conditional on $N(t)$. First condition on $N(t) = 1$.

$$\begin{aligned}
\mathsf{f}_{S_1|N(t)}(s_1|1) &= \lim_{\delta\to 0} \frac{\mathsf{p}_{N(s_1)}(0)\ \mathsf{p}_{\widetilde{N}(s_1,s_1+\delta)}(1)\ \mathsf{p}_{\widetilde{N}(s_1+\delta,t)}(0)}{\delta \mathsf{p}_{N(t)}(1)} \\
&= \frac{e^{-\lambda s_1}\lambda\delta e^{-\lambda\delta} e^{-\lambda(t-s_1-\delta)}}{\delta\lambda t e^{-\lambda t}} \quad = \quad \frac{1}{t}
\end{aligned}$$

The important point is that this does not depend on s_1, i.e., it is uniform over $(0,t]$.

Next consider $N(t) = 2$.

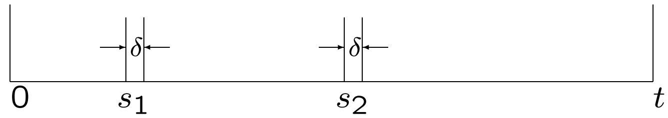

$$
\begin{aligned}
\mathsf{f}(s_1 s_2|2) &= \lim_{\delta} \frac{e^{-\lambda s_1}\lambda\delta e^{-\lambda\delta} e^{-\lambda(s_2-s_1-\delta)}\lambda\delta e^{-\lambda\delta} e^{-\lambda(t-s_2-\delta)}}{\delta^2 \mathsf{p}_{N(t)}(2)} \\
&= \frac{2}{t^2}
\end{aligned}
$$

Again, this does not depend on s_1, s_2, or λ for $0 < s_1 < s_2 < t$, i.e., it is uniform over the given region of s_1, s_2.

We can do the same thing for $N(t) = n$ for arbitrary n. Note that the exponents above always sum to λt.

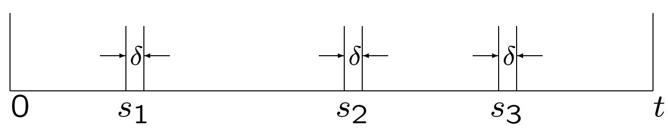

$$
\begin{aligned}
\mathsf{f}_{\vec{S}^{(n)}|N(t)}(\vec{s}^{(n)}|n) &= \lim_{\delta\to 0} \frac{(\delta\lambda)^n \exp(-\lambda t)}{\delta^n \mathsf{p}_{N(t)}(n)} \\
&= \frac{n!}{t^n}
\end{aligned}
$$

This is 'uniform' over $0 < s_1 < \cdots < s_n < t$.

This is a uniform n dimensional probability density over the volume $t^n/n!$ corresponding to the constraint region $0 < s_1 < \cdots < s_n < t$.

How did this derivation 'know' that the volume of $s_1, \ldots, s_n$ over $0 < s_1 < \cdots < s_n < t$ is $n!/t^n$?

To see why $n!$ appears in this uniform density, let $U_1, \dots, U_n$ be n IID rv's, each uniform over $(0, t]$. Let $S_1, \dots, S_n$ be the order statistics for $\vec{U}^n$, i.e.,

$$S_1 = \min(U_1, \dots U_n), \dots, S_k = k^{th} \textbf{ smallest}, \dots$$

The region of volume t^n where the density of $\vec{U}^n$ is nonzero partitions into $n!$ regions, one in which $u_1 < u_2 < \cdots < u_n$ and one for each other ordering of $u_1, \dots, u_n$. From symmetry, each volume is the same, and thus each is $t^n/n!$.

The region where $\vec{S}^n$ is nonzero is one of these partitions, and thus also has volume $t^n/n!$.

Since $\vec{S}^n$ has the same density, whether it is the conditional density of n arrival epochs given $N(t) = n$ or the order statistics of n uniform rv's, we can use results about either for the other.

As an example of using order statistics, consider finding the distribution function of S_1 conditional on $N(t) = n$. Viewing S_1 as the minimum of $U_1, \dots, U_n$, we have

$$\begin{aligned}
\mathsf{Pr}\left\{\min_{1 \le i \le n} U_i > s_1\right\} &= \textstyle\prod_{i=1}^n \mathsf{Pr}\{U_i > s\} \\
&= \textstyle\prod_{i=1}^n \left(1 - \frac{s_1}{t}\right) = \left(1 - \frac{s_1}{t}\right)^n \\
\mathsf{F}^{\mathsf{c}}_{S_1|N(t)}(s_1|n) &= \left(1 - \frac{s_1}{t}\right)^n \qquad \textbf{for } s_1 \le t
\end{aligned}$$

$\mathsf{E}[S_1 \mid N(t) = n]$ **can be found by integration,**

$$\mathsf{E}[S_1 \mid N(t) = n] = \frac{t}{n+1}$$

Return to $N(t) = 2$ and look at $\mathsf{f}_{X_1X_2|N(t)}(x_1, x_2|2)$.

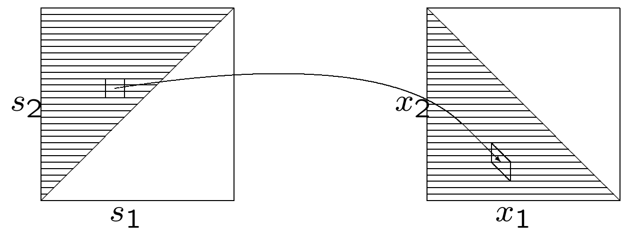

Note that the area in the s_1, s_2 space where $0 < s_1 < s_2 < t$ is $t^2/2$, explaining why the density is $2/t^2$.

Note that a δ^2 box in the s_1, s_2 space maps into a parallelepided of the same area in the x_1, x_2 space. The area in the x_1, x_2 space where $0 < x_1 + x_2 < t$ is again $t^2/2$ and $\mathsf{f}(x_1, x_2|N(t){=}2) = 2/t^2$.

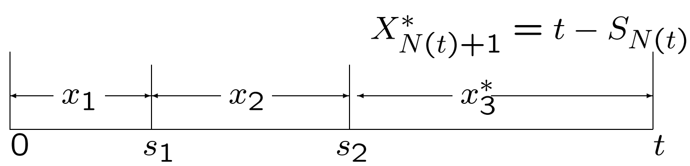

$\mathsf{f}(x_1, x_2|N(t)) = 2/t^2$ **for** $x_1, x_2, x_3^* > 0;\ x_1{+}x_2{+}x_3^* = t$
From symmetry, any 2 of the variables X_1, X_2, X_3^* can replace X_1, X_2 above.

Also, since $X_1 = S_1$, X_1 has density and expected value

$$\begin{aligned} \mathsf{F}^{\mathsf{c}}_{X_1|N(t)}(x_1|2) &= \left(1 - \frac{x_1}{t}\right)^2 \\ \mathsf{E}\left[X_1 \mid N(t){=}2\right] &= \frac{t}{3}. \end{aligned}$$

we see that X_2 and X_3^* can be substituted for X_1 in the above formulas. As a sanity check, note that $\mathsf{E}\left[X_1{+}X_2{+}X_3^*\right] = t$

This extends to $N(t) = n$ for arbitrary n.

$$\begin{aligned} \mathsf{f}_{X_1|N(t)}(x_1|n) &= \left(1 - \frac{x_1}{t}\right)^n \\ \mathsf{E}\left[X_1 \mid N(t){=}n\right] &= \frac{t}{n+1}. \end{aligned}$$

This relation also applies to $X^*_{n+1} = t - S_{N(t)}$

$$\begin{aligned} \mathsf{f}_{X^*_{n+1}|N(t)}(x|n) &= \left(1 - \frac{x}{t}\right)^n \\ \mathsf{E}\left[X^*_{n+1} \mid N(t){=}n\right] &= \frac{t}{n+1}. \end{aligned}$$

$$\begin{aligned} \mathsf{E}\left[X^*_{N(t)+1}\right] &= \sum_{n=0}^{\infty} \frac{t}{n{+}1} \, \frac{(\lambda t)^n e^{-\lambda t}}{n!} \\ &= \sum_{n=0}^{\infty} \frac{(\lambda t)^{n+1} e^{-\lambda t}}{\lambda (n{+}1)!} \;=\; \frac{1 - e^{-\lambda t}}{\lambda} \end{aligned}$$

Paradox: The mean interarrival time for a Poisson process is $1/\lambda$. But the mean time from any given t to the next arrival is $1/\lambda$ and the mean time back to the previous arrival is $(1/\lambda)(1{-}e^{-\lambda t})$. Thus the mean length of the interval containing t is $(1/\lambda)(2 - e^{-\lambda t})$.

This paradox will become clearer when we study renewals. A temporary half-intuitive explanation is to first choose a sample path for a Poisson process and then choose a uniform random value for t over some large interval far from 0. The larger inter-arrival intervals occupy proportionally more of the overall interval than the smaller, so t is biased to lie in one of those larger intervals.