Network Analysis:

The Poisson Process

Exponential Distribution

Overview

- Goal of network performance analysis
- Poisson process
- Exponential distribution

Network Performance Analysis

- For any cloud system or network, jobs arrive randomly
- We want to compute
  - Mean arrival rates
  - Mean service rates
  - Network throughput
- Enable a revenue model
  - Revenue = income/hour – cost/hour

Network Metrics

- Count arriving customers
  - Poisson distribution
- Estimate network throughput
  - Exponential distribution

Statistical Modeling of Networks

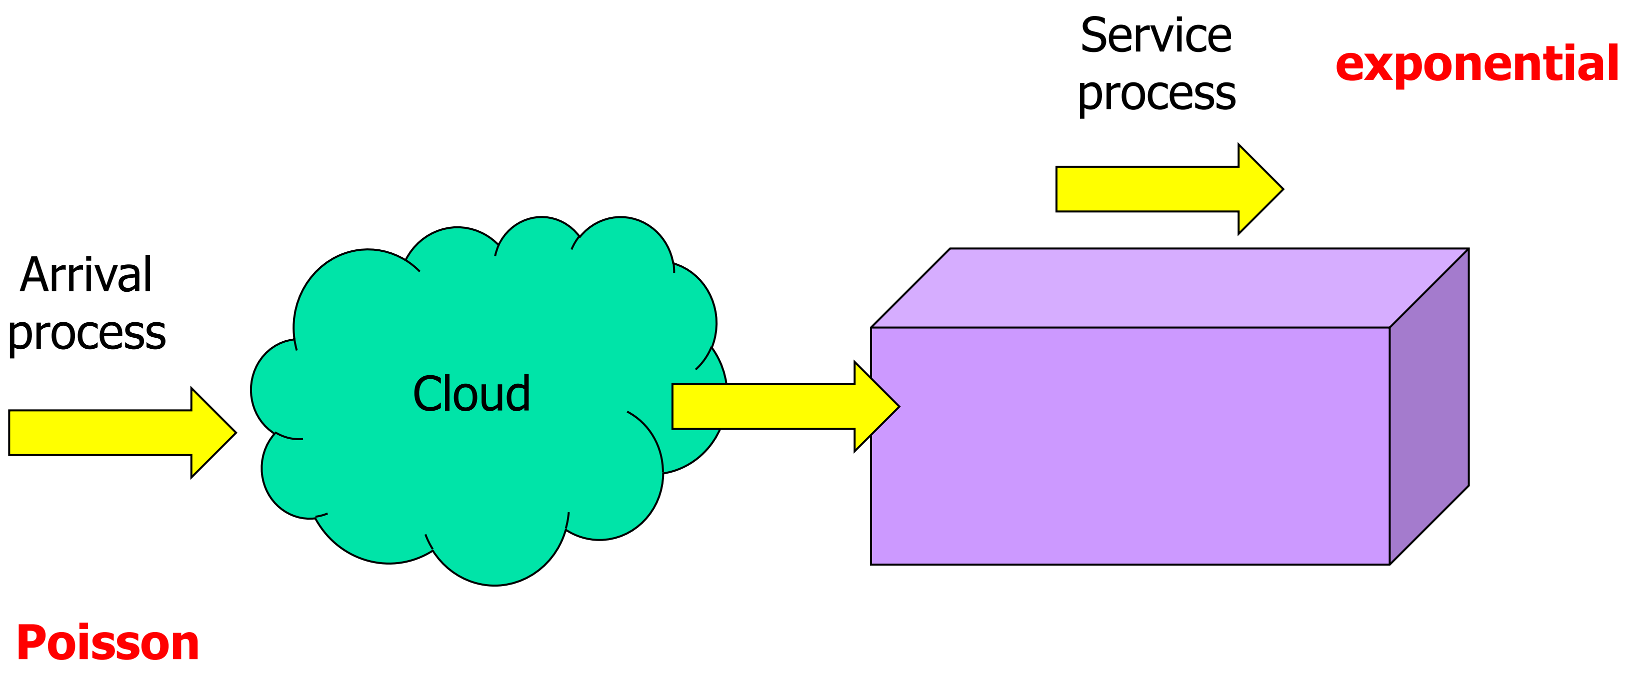

Definition

What is A Poisson Process?

The Poisson Process is a ***counting process*** that counts the number of occurrences of some specific event through time

Examples:

- Number of requests entering a cloud system
- Number of calls received at a telephone exchange
- Number of customers arriving to a counter

The Poisson Process

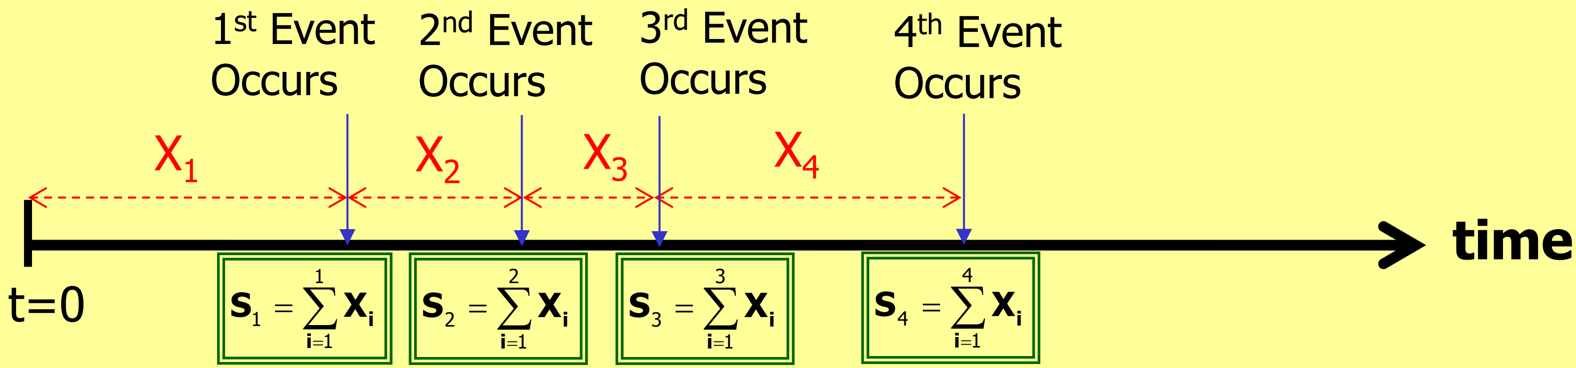

- X_1, X_2, … represent a sequence of +ve independent random variables with identical distribution
- X_n depicts the time elapsed between the $(n-1)^{th}$ event and n^{th} event occurrences
- S_n depicts a random variable for the time at which the n^{th} event occurs
- Define N(t) as the number of events that have occurred up to some arbitrary time t.

The counting process { N(t), t>0 } is called **a Poisson process** if the inter-occurrence times X_1, X_2, … follow the **exponential distribution**

The Poisson Process: Example

For some reason, you decide everyday at 3:00 PM to go to the bus stop and count the number of buses that arrive. You record the number of buses that have passed after 10 minutes

Sunday → $N(t=10\ min) = 2$

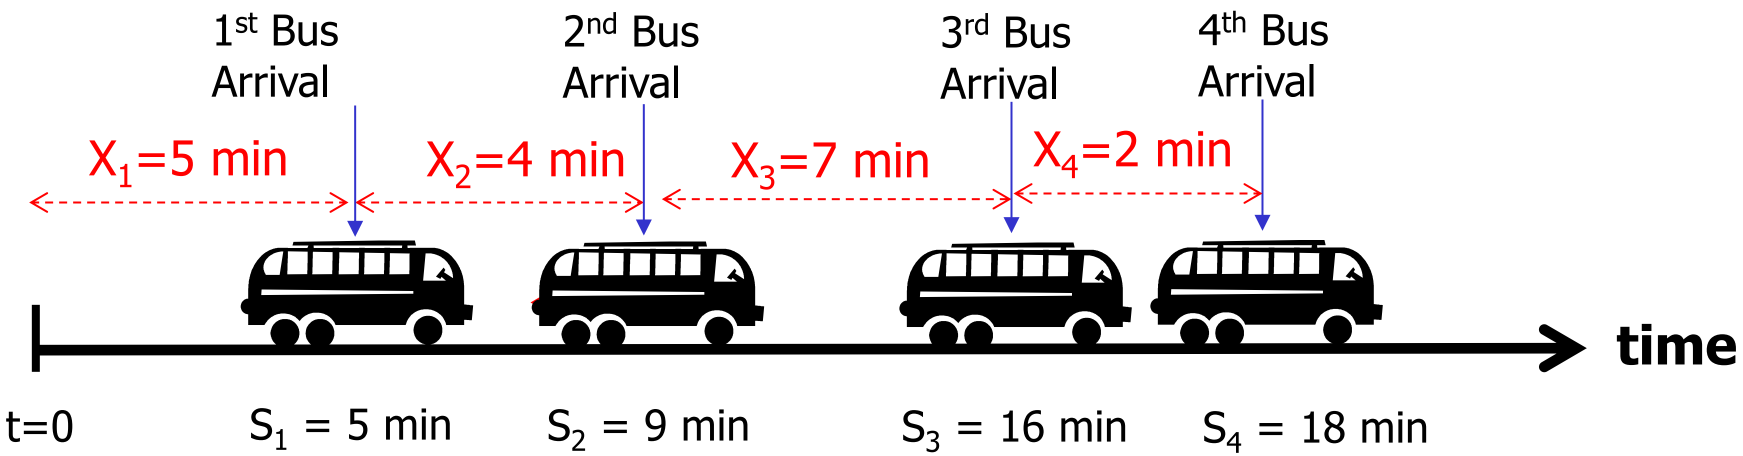

The Poisson Process: Example

For some reason, you decide everyday at 3:00 PM to go to the bus stop and count the number of buses that arrive. You record the number of buses that have passed after 10 minutes

Monday → $N(t=10\ \text{min}) = 4$

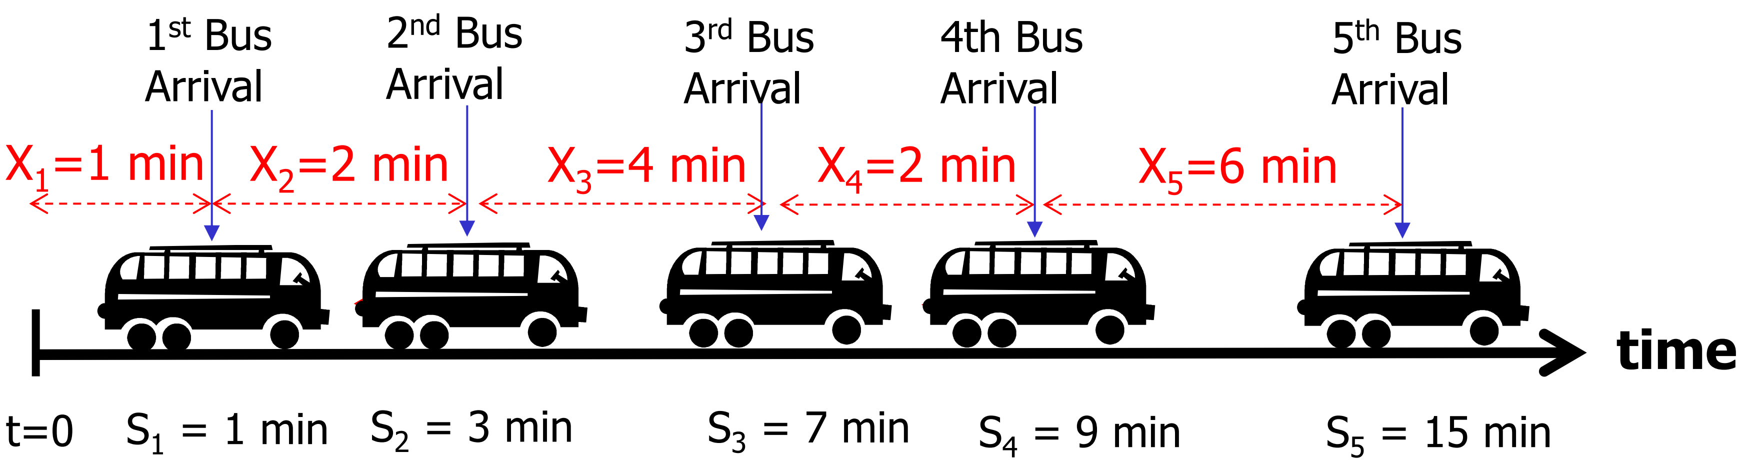

The Poisson Process: Example

For some reason, you decide everyday at 3:00 PM to go to the bus stop and count the number of buses that arrive. You record the number of buses that have passed after 10 minutes

Tuesday → $N\,(t{=}10\text{ min}) = 1$

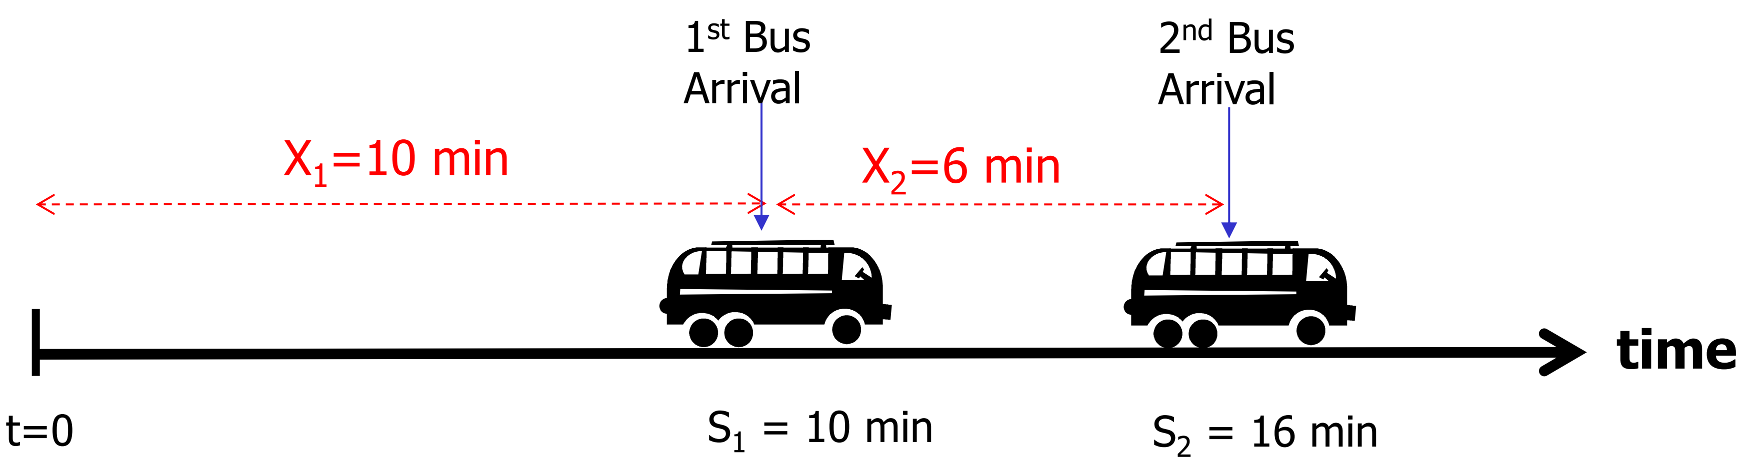

The Poisson Process: Example

- Sunday → $N(t=10\ min) = 2$
- Monday → $N(t=10\ min) = 4$
- Tuesday → $N(t=10\ min) = 1$
- Wednesday → $N(t=10\ min) = 5$
- Thursday → $N(t=10\ min) = 0$
- ⋮ → ⋮

Given that X_i follow an exponential distribution then $N(t=10)$ follows a Poisson Distribution

Poisson Distribution [Discrete]

- Poisson distribution describes many random processes quite well and is mathematically quite simple.
  - where $\alpha > 0$, pdf and cdf are:

$$p(x) = \begin{cases} \dfrac{e^{-\alpha}\alpha^{x}}{x!}, & x = 0,1,\ldots \\ 0, & \text{otherwise} \end{cases}$$

$$F(x) = \sum_{i=0}^{x} \frac{e^{-\alpha}\alpha^{i}}{i!}$$

  - $\alpha = \lambda t$ *with rate* λ
  - $E(X) = \alpha = V(X)$

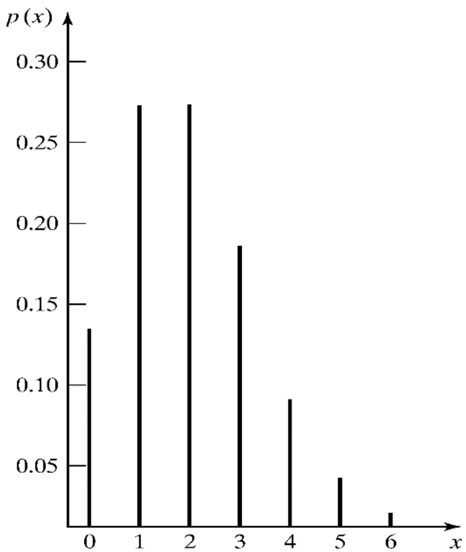

(a)

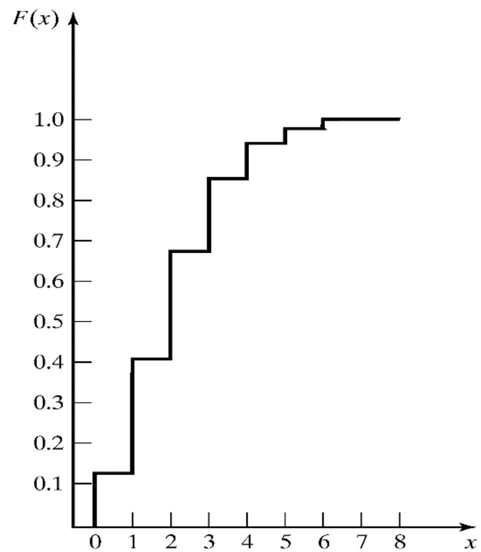

(b)

Poisson Distribution [Discrete]

- Example: A computer repair person is "beeped" each time there is a call for service. The number of beeps per hour ~ Poisson (α = 2 per hour).

  - The probability of three beeps in the next hour:

    $p(3) = e^{-2}2^3/3! = 0.18$

    also, $p(3) = F(3) – F(2) = 0.857-0.677=0.18$

  - The probability of two or more beeps in a 1-hour period:

    $p(2 \text{ or more}) = 1 – p(0) – p(1)$

    $= 1 – F(1)$

    $= 0.594$

Poisson Process

- Definition: *N(t)* is a counting function that represents the number of events occurred in *[0,t].*
- A counting process {*N(t), t>=0*} is a Poisson process with mean rate λ if:
  - Arrivals occur one at a time
  - {*N(t), t>=0*} has stationary increments
  - {*N(t), t>=0*} has independent increments
- Properties

$$P[N(t)=n]=\frac{e^{-\lambda t}(\lambda t)^n}{n!}, \qquad \text{for } t\geq 0 \text{ and } n=0,1,2,\ldots$$

  - Equal mean and variance: $E[N(t)] = V[N(t)] = \lambda t$
  - Stationary increment: The number of arrivals in time *s* to *t* is also Poisson-distributed with mean $\lambda(t-s)$

Poison Process (2)

- Note that N(t) in the previous slide has the Poisson distribution with parameter $\alpha = \lambda t$.

- This accounts for the mean equaling the variance.

- An alternative definition of a Poisson process:
  - if the interarrival times are **distributed exponentially** and independently, then the number of arrivals by time t, say N(t), meets the three Poisson assumptions and is therefore a Poisson process.

The Exponential Distribution

The exponential distribution describes a *continuous random variable*

Cumulative Distribution Function (CDF)

$$F_{X_n}(x) = \Pr\left[X_n \leq x\right] = 1 - e^{-\lambda x}$$

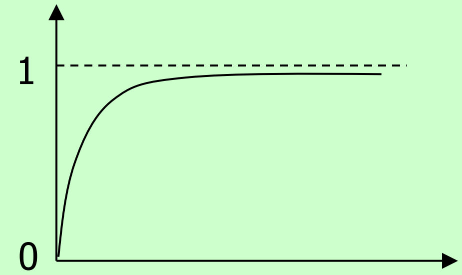

Probability Density Function (PDF)

$$f_{X_n}(x) = \frac{dF_{X_n}(x)}{dx} = \lambda e^{-\lambda x}$$

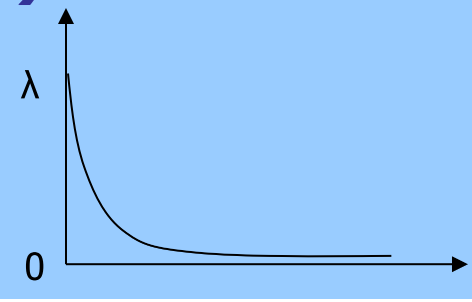

Time between random events / time till first random event ?

If a Poisson process has constant average rate ν, the mean after a time t is

$$\lambda = \nu t.$$

What is the probability distribution for the time to the first event?

$\Rightarrow$ Exponential distribution

Poisson - *Discrete* distribution: P(number of events)

Exponential - *Continuous* distribution: P(time till first event)

Exponential distribution

The continuous random variable Y has the Exponential distribution, with ***constant*** rate parameter ν if:

$$f(y) = \begin{cases} \nu e^{-\nu y}, & y > 0 \\ 0, & y < 0 \end{cases}$$

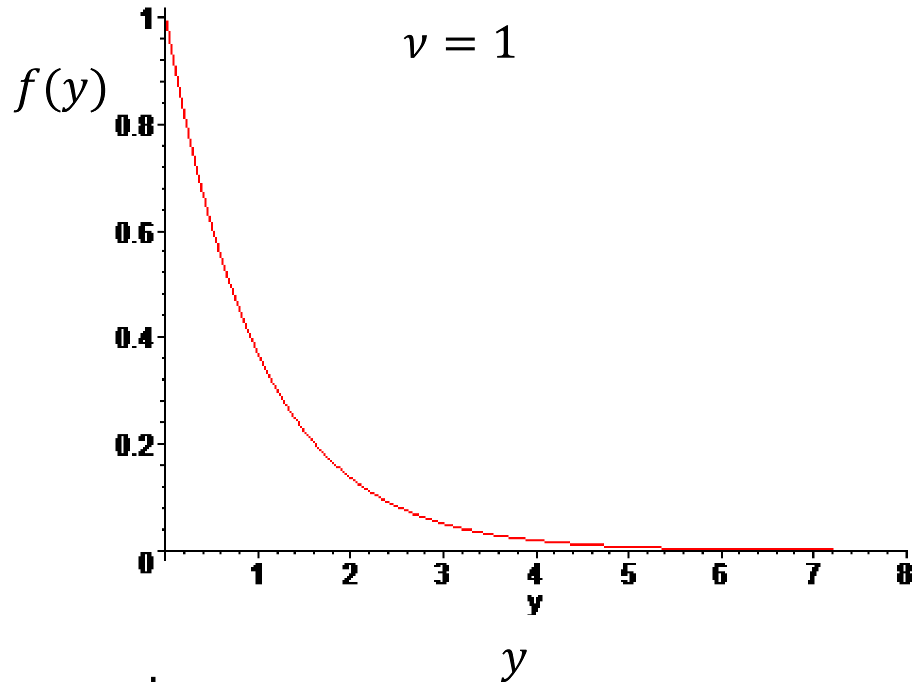

Occurrence

1) Time until the failure of a part.

2) Separation between randomly occurring events

- Assuming the probability of the events is constant in time: $\nu = \text{const}$

Relation to Poisson distribution

Poisson process has constant average rate ν, the mean after a time t is $\lambda = \nu t$.

The probability of no-occurrences in time t is

$$P(k=0) = \frac{e^{-\lambda}\lambda^k}{k!} = e^{-\lambda} = e^{-\nu t}.$$

If $f(t)$ is the pdf for the first occurrence, then the probability of no occurrences is

$$P(\text{no occurrence by } t) = 1 - P(\text{first occurrence has happened by } t)$$

$$= 1 - \int_0^t f(t)dt$$

$$\Rightarrow 1 - \int_0^t f(t)dt = e^{-\nu t} \qquad \Rightarrow \int_0^t f(t)\,dt = 1 - e^{-\nu t}$$

Solve by differentiating both sides respect to t assuming constant ν,

$$\frac{d}{dt}\int_0^t f(t)\,dt = \frac{d}{dt}\left(1 - e^{-\nu t}\right)$$

$$\Rightarrow f(t) = \nu e^{-\nu t}$$

The time until the first occurrence (and between subsequent occurrences) has the Exponential distribution, parameter ν.

Example

On average lightening kills three people each year, $\lambda = 3$. So the rate is $\nu = 3/\text{year}$.

Assuming strikes occur randomly at any time during the year so ν is constant, time from today until the next fatality has pdf (using t in years)

$$f(t) = \nu e^{-\nu t} = 3\, e^{-3t}$$

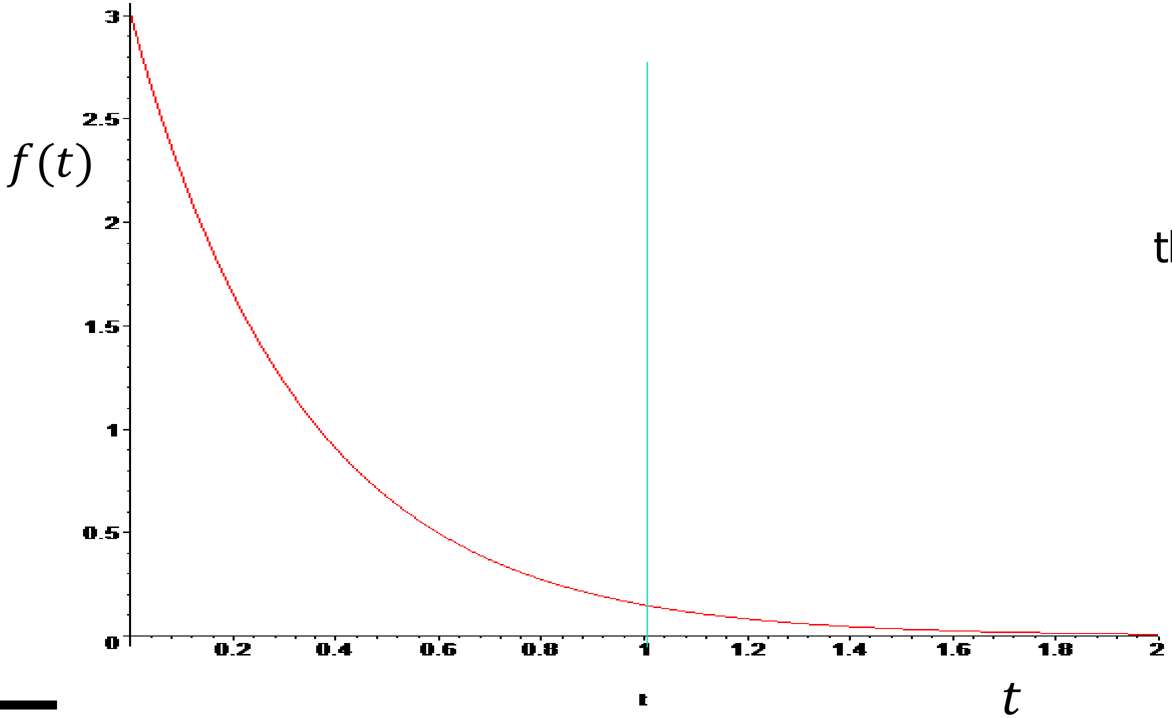

E.g. Probability the time till the next death is less than one year?

$$\int_0^1 f(t)dt = \int_0^1 3\, e^{-3t} dt$$

$$= \left[\frac{3e^{-3t}}{-3}\right]_0^1$$

$$= -e^{-3} + 1 \approx 0.95$$

Example: Reliability

The time till failure of an electronic component has an Exponential distribution and it is known that 10% of components have failed by 1000 hours.

(a) What is the probability that a component is still working after 5000 hours?

(b) Find the mean and standard deviation of the time till failure.

Answer

Let Y = time till failure in hours; $f(y) = \nu e^{-\nu y}$.

(a) First we need to find ν $\quad P(Y \le 1000) = \int_0^{1000} \nu e^{-\nu y}$

$$= [-e^{-\nu y}]_0^{1000} \quad = 1 - e^{-1000\nu}$$

$P(Y \le 1000) = 0.1 \Rightarrow \quad 1 - e^{-1000\nu} = 0.1$

$$\Rightarrow e^{-1000\nu} = 0.9$$

$$\Rightarrow -1000\nu = \ln 0.9 = -0.10536 \quad \Rightarrow \nu \approx 1.05 \times 10^{-4}$$

Continuous Random Variables

- Continuous RV
  - The values that the random variable can take are continuous
- Examples:
  - The failure time of a system
  - The value of a circuit resistance
- CDF F(X): cumulative distribution function
- The density function *f(X)* is given by the derivative of the cumulative distribution function
  - $f(X) = F'(X)$
- Example:

"The failure time of a system is exponentially distributed"

Exponential Distribution

- The cumulative distribution function: $F(X \leq t) = 1 - e^{-\lambda t}$
- $F(X>t) = 1-F(X\leq t) = e^{-\lambda t}$
- The exponential density function is $f(x) = \lambda e^{-\lambda x}$ if $x \geq 0$
  - The parameter λ is constant

$$f(x) = \lambda e^{-\lambda x}$$

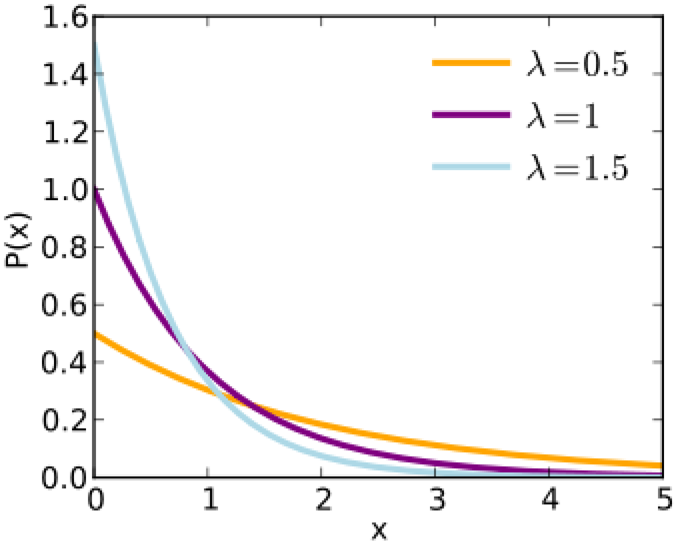

Exponential CDF

- The CDF is shown below:

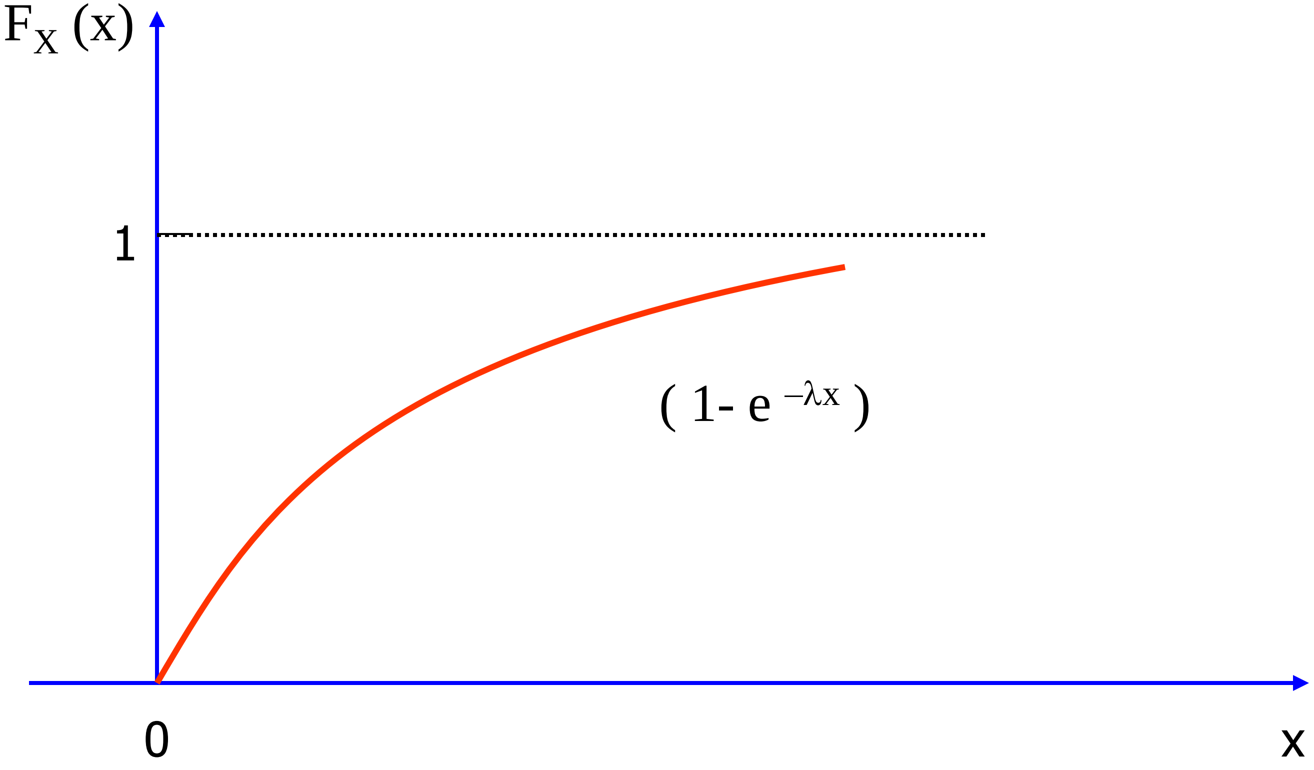

Interpretation- Exponential Distribution

- Exponential distribution occurs in reliability work over and over again, in the way used as the distribution of the time to failure for a great number of electronic-system parts
- The parameter λ is constant and is usually called the **failure rate** (with the units fraction failures per hour)
- The cumulative distribution function: $F(t) = 1 - e^{-\lambda t}$
- The success probability (probability of no failure): $1 - F(t) = e^{-\lambda t}$
- expected value (Mean Time Between Failures): 1/λ (MTBF)
- The most commonly used distribution in reliability and performance modeling

Exponential Distribution CDF

- ***Problem 1:*** *The transmission time X of messages in a communication system obeys the exponential law with parameter* λ *,that is* $P[X>x] = e^{-\lambda x}$ $x > o$

  Find the cdf of X. Find $P[T<X \leq 2T]$*, where* $T=1/\lambda$*.*

Solution: *The cdf of X is* $F_X(x) = P[X \leq x] = 1 - P[X > x]$:

$$f_X(x) = \begin{cases} 0 & x < 0 \\ 1 - e^{-\lambda x} & x \geq 0 \end{cases}$$

Continued…

Exponential Example

- *Compute $P[T < X \leq 2T]$*

$$P[T < X \leq 2T] = (1-e^{-2}) - (1 - e^{-1}) = e^{-1} - e^{-2} = .233$$

$F_X(x)$ is continuous for all x. Note also that its derivative exits everywhere except at x=0:

$$F'_X(x) = \begin{cases} 0 & x < 0 \\ \lambda e^{-\lambda x} & x > 0 \end{cases}$$

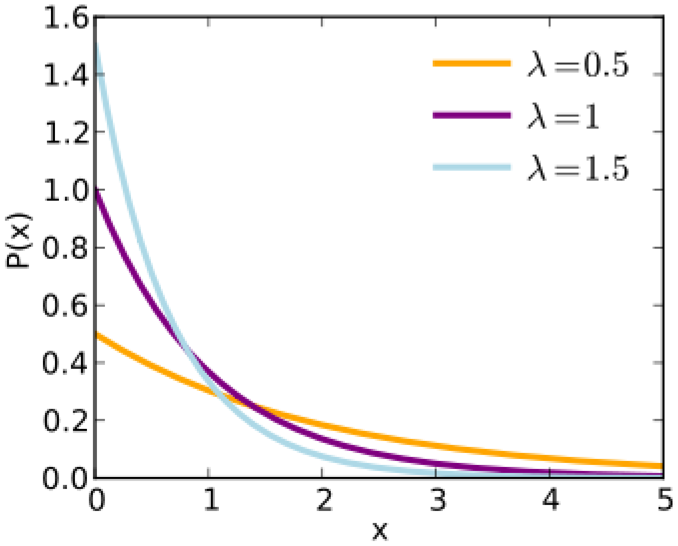

Continued…

Example: Al's Website

- The time between arrivals of orders at Al's Website follows an exponential probability distribution with a mean time between arrivals of 3 minutes.
- Al would like to know the probability that the time between two successive arrivals will be 2 minutes or less.

$$P(x \leq 2) = 1 - 2.71828^{-2/3} = 1 - .5134 = .4866$$

Example: Al's Website

- Graph of the Probability Density Function

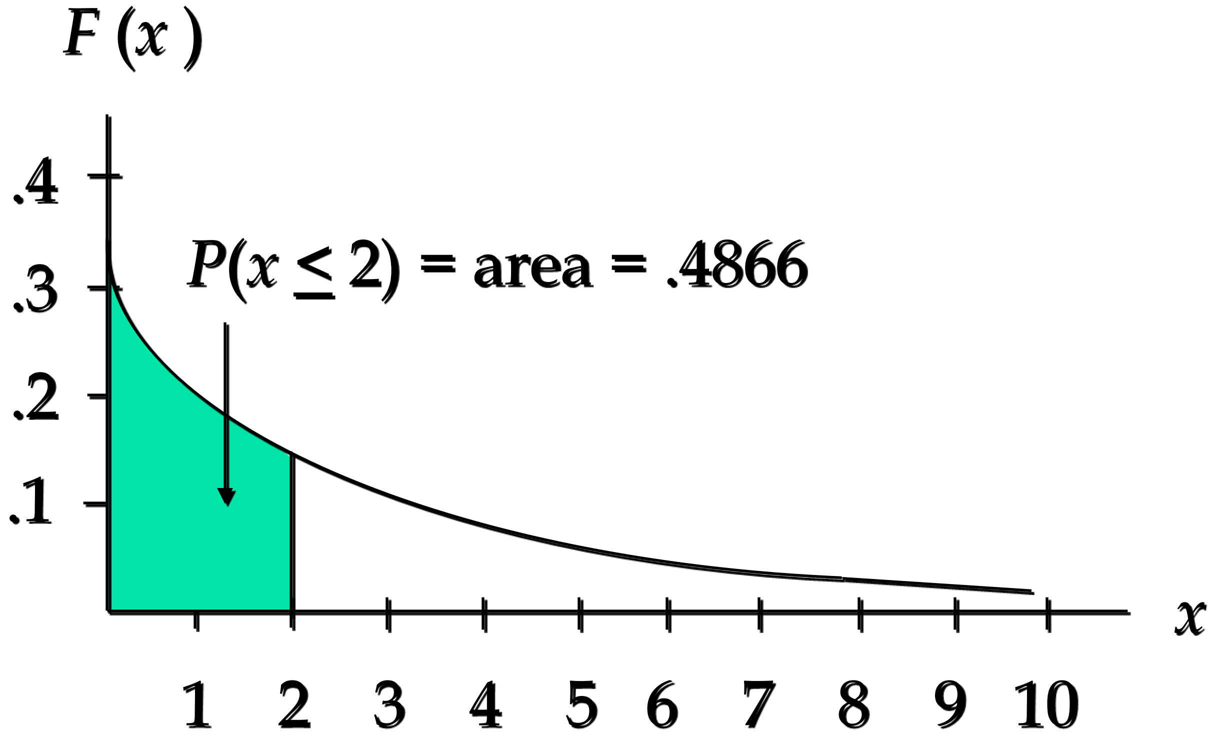

Memory-less property

A random variable X is said to be without memory, or *memory-less*, if

$P\{X>s+t|X>t\} = P\{X>s\}$ for all $s, t \geq 0$

Interpretation

If an item is alive at time t, then the distribution of the remaining amount of time that it survives is the same as the original lifetime distribution.

The item does not remember that it has already been in use for a time t

More on Memoryless Property

$P\{X>s+t|X>t\} = P\{X>s\}$ for all s, $t \geq 0$

$$\frac{P\{X>s+t, X>t\}}{P\{X>t\}} = P\{X>s\}$$

$$P\{X>s+t\} = P\{X>s\}P\{X>t\}$$

When X is exponentially distributed, it follows that $e^{-\lambda(s+t)} = e^{-\lambda s}e^{-\lambda t}$.

Hence, exponentially distributed random variable are memoryless.

Example

- Suppose the amount of time one spends in a bank is exponentially distributed with mean 10 minutes, that is, $\lambda=1/10$.
  - What is the probability that a customer will spend more than fifteen minutes in the bank?
  - What is the probability that a customer will spend more than fifteen minutes in the bank given that s/he is still in the bank after ten minutes?

Solution

$$P\{X > 15\} = e^{-15\lambda}$$

$$P\{X > 10\} = e^{-10\lambda}$$

$$P\{X > 5\} = e^{-5\lambda}$$

$$P\{X > 15 \mid X > 10\} =$$

It turns out that the exponentially distribution is the unique distribution possessing the memory-less property

The Poisson Distribution

$$\Pr\left[\mathbf{N}(\mathbf{t}) \geq \mathbf{k}\right] = \Pr\left[\mathbf{S_k} \leq \mathbf{t}\right]$$

$$\Pr\left[\mathbf{N}(\mathbf{t}) \geq \mathbf{k}+1\right] = \Pr\left[\mathbf{S_{k+1}} \leq \mathbf{t}\right]$$

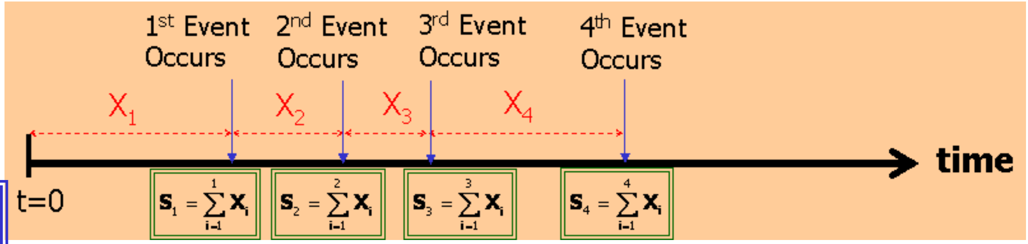

$$\Pr\left[\mathbf{N}(\mathbf{t}) = \mathbf{k}\right] = \Pr\left[\mathbf{N}(\mathbf{t}) \geq \mathbf{k}\right] - \Pr\left[\mathbf{N}(\mathbf{t}) \geq \mathbf{k}+1\right]$$

$$\Pr\left[\mathbf{N}(\mathbf{t}) = \mathbf{k}\right] = \left[1 - \sum_{\mathbf{i}=0}^{\mathbf{k}-1} \mathbf{e}^{-\lambda \mathbf{t}} \frac{(\lambda \mathbf{t})^{\mathbf{i}}}{\mathbf{i}!}\right] - \left[1 - \sum_{\mathbf{i}=0}^{\mathbf{k}} \mathbf{e}^{-\lambda \mathbf{t}} \frac{(\lambda \mathbf{t})^{\mathbf{i}}}{\mathbf{i}!}\right]$$

$$\Pr\left[\mathbf{N}(\mathbf{t}) = \mathbf{k}\right] = \mathbf{e}^{-\lambda \mathbf{t}} \frac{(\lambda \mathbf{t})^{\mathbf{k}}}{\mathbf{k}!}$$

Probability Mass Function for the Poisson Distribution

Mean of the Poisson Distribution

$$E[N(t)] = \sum_{k=0}^{\infty} k \Pr[N(t) = k] = \sum_{k=0}^{\infty} k e^{-\lambda t} \frac{(\lambda t)^k}{k!} = \sum_{k=1}^{\infty} k e^{-\lambda t} \frac{(\lambda t)^k}{k!}$$

$$E[N(t)] = e^{-\lambda t} (\lambda t) \sum_{k=1}^{\infty} \frac{(\lambda t)^{k-1}}{(k-1)!}$$

$$E[N(t)] = e^{-\lambda t} (\lambda t) \sum_{k=0}^{\infty} \frac{(\lambda t)^k}{K!} = (\lambda t)$$

where $\sum_{k=0}^{\infty} \frac{y^k}{k!} = e^y$

On average the time between two consecutive events is $1/\lambda$

This means that the event occurrence rate is λ

Consequently in time t, the expected number of events is **λt**

Variance of the Poisson Distribution

$$E\left[N(t)^2\right] = \sum_{k=0}^{\infty} k^2 \Pr\left[N(t)=k\right] = \sum_{k=0}^{\infty} k^2 e^{-\lambda t} \frac{(\lambda t)^k}{k!} = \sum_{k=1}^{\infty} k^2 e^{-\lambda t} \frac{(\lambda t)^k}{k!}$$

$$E\left[N(t)^2\right] = e^{-\lambda t}(\lambda t)\sum_{k=1}^{\infty}(k-1+1)\frac{(\lambda t)^{k-1}}{(k-1)!}$$

$$E\left[N(t)^2\right] = e^{-\lambda t}(\lambda t)\left[\sum_{k=1}^{\infty}(k-1)\frac{(\lambda t)^{k-1}}{(k-1)!} + \sum_{k=1}^{\infty}\frac{(\lambda t)^{k-1}}{(k-1)!}\right]$$

$$E\left[N(t)^2\right] = e^{-\lambda t}(\lambda t)\left[(\lambda t)\sum_{k=2}^{\infty}\frac{(\lambda t)^{k-2}}{(k-2)!} + \sum_{k=1}^{\infty}\frac{(\lambda t)^{k-1}}{(k-1)!}\right]$$

$$E\left[N(t)^2\right] = e^{-\lambda t}(\lambda t)\left[(\lambda t)e^{\lambda t} + e^{\lambda t}\right] = (\lambda t)^2 + (\lambda t) \Rightarrow \mathrm{Var}\left[N(t)\right] = (\lambda t)$$

Moment Generating Function of Poisson Distribution

The Moment Generating Function of any PMF for a discrete random variable may be used to deduce different parameters and characteristics of the distribution

$$G(Z) = E\left[Z^N\right] = \sum_{k=0}^{\infty} Z^k \Pr\left[N = k\right] = \sum_{k=0}^{\infty} Z^k e^{-\lambda t} \frac{(\lambda t)^k}{k!} = e^{-\lambda t} \sum_{k=0}^{\infty} \frac{(Z\lambda t)^k}{k!} = e^{-\lambda t(1-Z)}$$

What could G(Z) be used for

$$E\left[N(t)\right] = \frac{dG(Z)}{dZ}\Big|_{Z=1} = (\lambda t) e^{-\lambda t(1-Z)}\Big|_{Z=1} = (\lambda t)$$

$$E\left[N(t)^2\right] = \frac{d^2G(Z)}{dZ^2}\Big|_{Z=1} + \frac{dG(Z)}{dZ}\Big|_{Z=1} = (\lambda t)^2 + (\lambda t)$$

Remaining Time of Exponential Distributions

X_{k+1} is the time interval between the k^{th} and $k+1^{th}$ arrivals

Condition:

T units have passed and the $k+1^{th}$ event has not occurred yet

Question:

Given that X^*_{k+1} is the remaining time until the $k+1^{th}$ event occurs

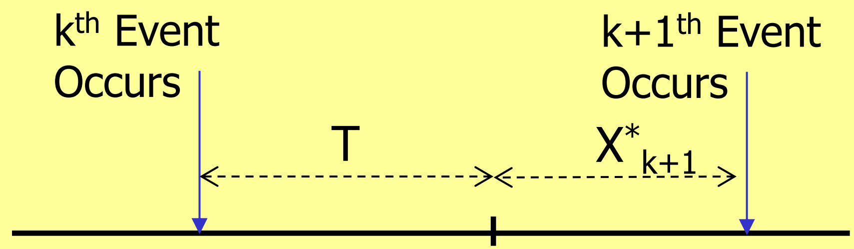

$$S_k = \sum_{i=1}^{k} X_i$$

What is $\Pr[X^*_{k+1} \leq x]$

$$\Pr\left[X^*_{k+1} \leq x\right] = \Pr\left[X_{k+1} \leq T + x \,\middle|\, X_{k+1} > T\right]$$

$$\Pr\left[X^*_{k+1} \leq x\right] = \frac{\Pr\left[T < X_{k+1} \leq T + x\right]}{\Pr\left[X_{k+1} > T\right]}$$

$$\Pr\left[X^*_{k+1} \leq x\right] = \frac{\Pr\left[X_{k+1} \leq T + x\right] - \Pr\left[X_{k+1} < T\right]}{1 - \Pr\left[X_{k+1} < T\right]}$$

Remaining Time of Exponential Distributions

$$\Pr\left[X^*_{k+1} \le x\right] = \frac{\Pr\left[X_{k+1} \le T + x\right] - \Pr\left[X_{k+1} < T\right]}{1 - \Pr\left[X_{k+1} < T\right]}$$

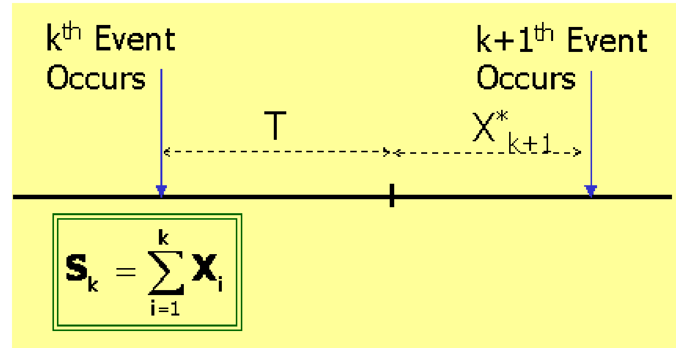

X_{k+1} follows an exponential Distribution, i.e., $\Pr[X_k+1 \le t] = 1 - e^{-\lambda t}$

$$\Pr\left[X^*_{k+1} \le x\right] = \frac{\left(1 - e^{-\lambda(T+x)}\right) - \left(1 - e^{-\lambda(T)}\right)}{1 - \left(1 - e^{-\lambda(T)}\right)}$$

$$\Pr\left[X^*_{k+1} \le x\right] = \frac{e^{-\lambda(T)} - e^{-\lambda(T+x)}}{e^{-\lambda(T)}} = \boxed{1 - e^{-\lambda(x)}}$$

The remaining time X^*_{k+1} follows an exponential distribution with the ***same mean 1/λ*** as that of the inter-arrival time X_{k+1}

The Memoryless Property of Exponential Distributions

The Memoryless Property:

The waiting time until the next arrival has *the same exponential distribution* as the original inter-arrival time regardless of long ago the last arrival occurred

Memoryless Property of Exponential Distribution and the Poisson Process

$$\Pr\left[N(u+s) - N(u) = k\right] = \frac{(\lambda s)^k e^{-\lambda s}}{k!}$$

In the Poisson process, the number of arrivals within any time interval s follows a Poisson distribution with mean λs

Splitting and Pooling [Poisson Process]

- Splitting:
  - Suppose each event of a Poisson process can be classified as Type I, with probability *p* and Type II, with probability *1-p.*
  - *N(t)* = *N1(t)* + *N2(t)*, where *N1(t)* and *N2(t)* are both Poisson processes with rates λp and $\lambda(1-p)$

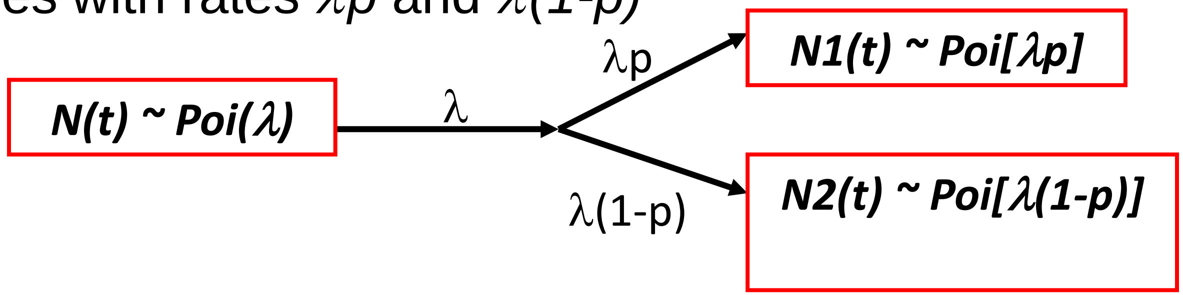

- Pooling:
  - Suppose two Poisson processes are pooled together
  - *N1(t)* + *N2(t)* = *N(t)*, where *N(t)* is a Poisson processes with rates $\lambda_1 + \lambda_2$

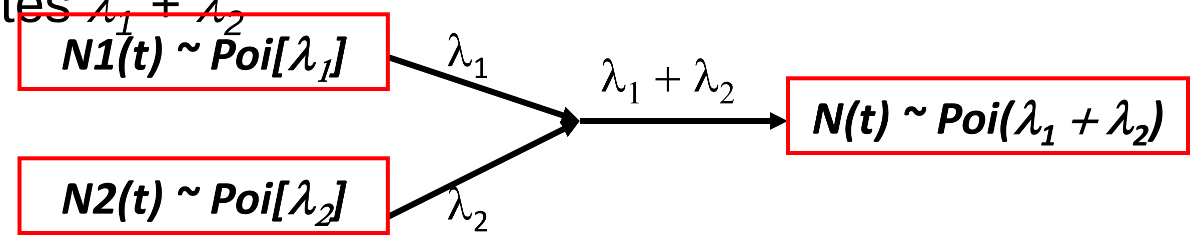

Merging of Poisson Processes

- $\{N_1(t), t \geq 0\}$ and $\{N_2(t), t \geq 0\}$ are two independent Poisson processes with respective rates λ_1 and λ_2,
- $\{N_i (t)\}$ corresponds to type i arrivals.
- The *merged process* $N(t) = N_1(t) + N_2(t), t \geq 0$. Then $\{N(t), t \geq 0\}$ is a *Poisson process* with rate $\lambda = \lambda_1 + \lambda_2$.

- Z_k is the inter-arrival time between the $(k - 1)^{th}$ and k^{th} arrival in the merged process
- $I_k = i$ if the k^{th} arrival in the merged process is a type i arrival,
- For any $k = 1, 2, \ldots,$

  $P\{I_k = i \mid Z_k = t\} = \lambda_i/(\lambda_1+\lambda_2)$, i= 1, 2, *independently of t* .

Splitting of Poisson Processes

- $\{N(t), t \geq 0\}$ is a Poisson process with rate λ.
- Each arrival of the process is classified as being a type 1 arrival or type 2 arrival with respective probabilities p_1 and p_2, independently of all other arrivals.
- $N_i(t)$ is the number of type i arrivals up to time t .
- $\{N_1(t)\}$ and $\{N_2(t)\}$ are *two independent Poisson processes* having respective rates λp_1 and λp_2.

Extra Notes: Derivations, etc.

Mean of the Exponential Distribution

$$E[X_n] = \int_0^\infty x f_{X_n}(x)\,dx = \int_0^\infty \lambda x e^{-\lambda x}\,dx$$

$$\text{Let } u = x,\ dv = \lambda e^{-\lambda x}dx \Rightarrow du = dx,\ v = -e^{-\lambda x}$$

$$E[X_n] = uv\Big|_0^\infty - \int_0^\infty v\,du = -xe^{-\lambda x}\Big|_0^\infty + \int_0^\infty e^{-\lambda x}dx$$

$$E[X_n] = \left[\frac{\infty}{\infty} - 0\right] - \frac{1}{\lambda}e^{-\lambda x}\Big|_0^\infty$$

L'Hopital Theorem

$$\text{Lim}_{x\to\infty}\frac{-x}{e^{\lambda x}} = \text{Lim}_{x\to\infty}\frac{-1}{\lambda e^{\lambda x}} = 0$$

$$E[X_n] = \left[0 + \frac{1}{\lambda}\right] = \frac{1}{\lambda}$$

Variance of the Exponential Distribution

$$E\left[X_n^2\right] = \int_0^\infty x^2 f_{X_n}(x)\,dx = \int_0^\infty \lambda x^2 e^{-\lambda x}\,dx$$

Let $u = x^2,\ dv = \lambda e^{-\lambda x}dx \Rightarrow du = 2x\,dx,\ v = -e^{-\lambda x}$

$$E\left[X_n^2\right] = uv\Big|_0^\infty - \int_0^\infty v\,du = -x^2 e^{-\lambda x}\Big|_0^\infty + 2\int_0^\infty x e^{-\lambda x}\,dx$$

$$E\left[X_n^2\right] = \left[\left(\frac{\infty}{\infty}\right) - 0\right] + \frac{2}{\lambda}\int_0^\infty x\lambda e^{-\lambda x}\,dx$$

L'Hopital Theorem

$$\text{Lim}_{x\to\infty} \frac{-x^2}{e^{\lambda x}} = 0$$

$$E\left[X_n^2\right] = \frac{2}{\lambda} E\left[X_n\right] = \frac{2}{\lambda^2}$$

$$Var\left[X_n\right] = E\left[X_n^2\right] - E\left[X_n\right]^2 = \frac{2}{\lambda^2} - \frac{1}{\lambda^2} = \frac{1}{\lambda^2}$$