CPSC 531: System Modeling and Simulation

Carey Williamson
Department of Computer Science
University of Calgary
Fall 2017

Stochastic Process:

Collection of random variables indexed over time

- **Example:**
  - $N(t)$: number of jobs in the system at time t
  - The number $N(t)$ at any time t is a random variable
  - Can find the probability distribution functions for $N(t)$ at each possible value of t

- Notation: $\{N(t): t \geq 0\}$

- Counting Process:
  A stochastic process that represents the total number of events occurring in the time interval $[0, t]$

- Poisson Process:
  The counting process $\{N(t), t \geq 0\}$ is a Poisson process with rate λ, if:
  - $N(0) = 0$
  - The process has independent increments
  - The number of events in any interval of length t follows a Poisson distribution with mean λt. That is, for all $s, t \geq 0$

$$\mathbb{P}(N(t+s) - N(s) = n) = \frac{(\lambda t)^n}{n!} e^{-\lambda t}$$

Property: equal mean and variance: $E[N(t)] = V[N(t)] = \lambda t$

Poisson Arrival Process

- A common modeling assumption in simulation and/or analysis is that of Poisson arrivals (aka Poisson arrival process)

- Poisson Arrivals Model:
  - Arrivals occur randomly (i.e., at "random" times)
  - No two arrivals occur at exactly the same time
  - Inter-arrival times are exponentially distributed and independent
  - The counting process (number of events in any interval of length t) follows a Poisson distribution with mean λt. That is, for all $s, t \geq 0$

$$\mathbb{P}(N(s+t) - N(s) = n) = \frac{(\lambda t)^n}{n!} e^{-\lambda t}$$

Property: equal mean and variance: $E[N(t)] = V[N(t)] = \lambda t$

Interarrival Times

- Consider the interarrival times of a Poisson arrival process with rate λ, denoted by $A_1, A_2, \ldots$, where A_i is the elapsed time between arrival i and arrival $i+1$

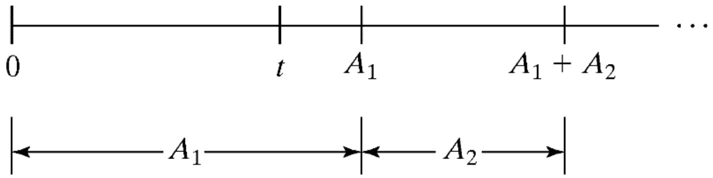

- Interarrival times, $A_1, A_2, \ldots$ are independent identically distributed exponential random variables with mean $1/\lambda$

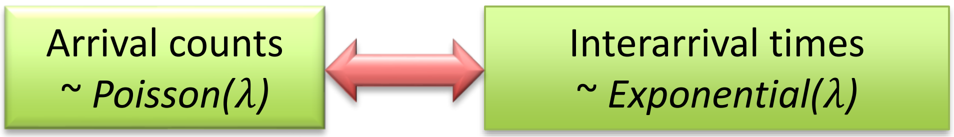

Pooling Property

- If you combine multiple Poisson processes together (pooling), then the resulting process is also Poisson
- Aggregate rate is the sum of the individual rates being pooled
- Pooling:
  - $N_1(t)$: Poisson process with rate λ_1
  - $N_2(t)$: Poisson process with rate λ_2
  - $N(t) = N_1(t) + N_2(t)$: Poisson process with rate $\lambda_1 + \lambda_2$

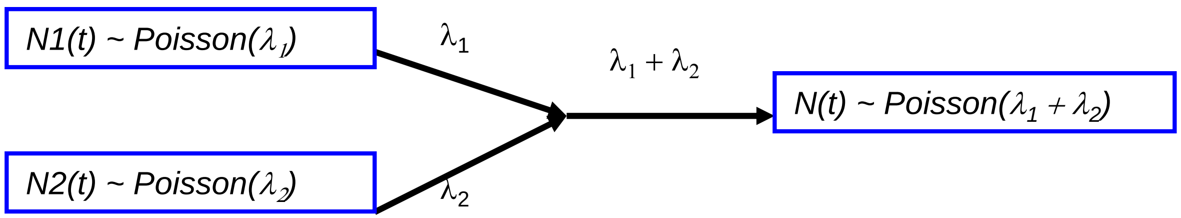

Splitting Property

- If you split a Poisson process “randomly”, then the resulting individual processes are also Poisson
- Individual rates sum to that of the original process

- Splitting:
  - $N(t)$: Poisson process with rate λ
  - Each event is classified as Type 1 (probability p) or Type 2 (probability $1-p$)
  - $N_1(t)$: The number of Type 1 events is a Poisson process with rate $p\lambda$
  - $N_2(t)$: The number of Type 2 events is a Poisson process with rate $(1-p)\lambda$
  - $N(t) = N_1(\mathrm{t}) + N_2(t)$

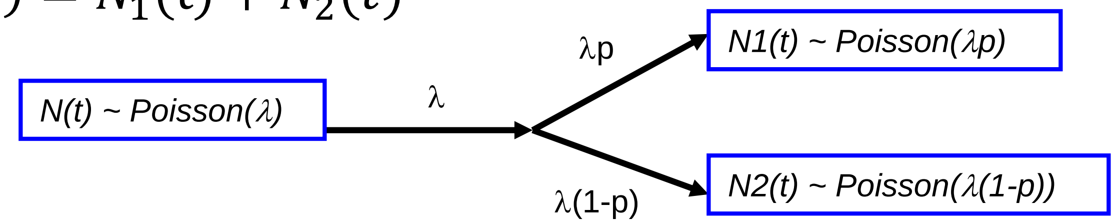

More on Poisson Distribution

- $\{N(t), t \geq 0\}$: a Poisson process with arrival rate λ

- Probability of no arrivals in a small time interval h:

$$\mathbb{P}(N(h) = 0) = e^{-\lambda h} \approx 1 - \lambda h$$

- Probability of one arrivals in a small time interval h:

$$\mathbb{P}(N(h) = 1) = \lambda h \cdot e^{-\lambda h} \approx \lambda h$$

- Probability of two or more arrivals in a small time interval h:

$$\mathbb{P}(N(h) \geq 2) = 1 - \big(\mathbb{P}(N(h) = 0) + \mathbb{P}(N(t) = 1)\big) \approx 0$$

Aside: Poisson Point Process

- The discussion so far has focused on the temporal aspects of a Poisson process (i.e., in time domain)
- Similar properties apply to the spatial domain (i.e., location) in one or more dimensions

- Poisson Point Process:
  - Items are dispersed randomly (i.e., at "random" locations)
  - No two items occur at exactly the same place
  - Inter-item distances are exponentially distributed and independent
  - The counting process (number of events in any region of area A) follows a Poisson distribution with mean λA. That is, for all $s, A \geq 0$

$$\mathbb{P}(N(A) = n) = \frac{(\lambda A)^n}{n!} e^{-\lambda A}$$

Property: equal mean and variance: $E[N(A)] = V[N(A)]$

Bonus Material: Queueing Theory Basics

- Queueing theory is a well-established area of performance modeling that studies the behaviour of queues
- Classic textbook: Queueing Systems: Vol 1, by L. Kleinrock

- The foundation of queueing theory is built using the types of probability models that we have just been studying
- The goal in this short presentation is to show you the basics of the M/M/1 queuing model, for which $N = \rho/(1-\rho)$

- This is only a preview; we will revisit this material in much more depth in late November and/or early December

Notation

- λ: The average arrival rate (in customers per time unit)
  - The mean inter-arrival time is 1/λ
- μ: The average service rate (in customers per time unit)
  - The mean service time requirement is 1/μ
- ρ: The average load offered to the system
  - $\rho = \lambda/\mu < 1.0$

- Kendall notation for queueing systems:
  - Arrival process: either M (for Markovian) or G (for General)
  - Service time process: either D (for Deterministic), M, or G
  - N: The number of servers
- Example: M/M/1 is a single-server queue with a Poisson arrival process and exponential service times for customers

M/M/1 System Model

Markov chain model of classic M/M/1 queue
Birth-death process representing system occupancy
Fixed arrival rate λ
Fixed service rate μ

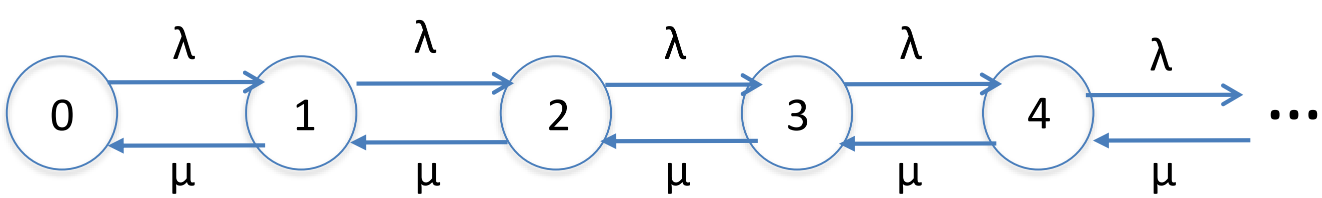

Mean system occupancy: $N = \rho / (1 - \rho)$

$p_n = p_0 (\lambda/\mu)^n$

Ergodicity requirement: $\rho = \lambda/\mu < 1$

$U = 1 - p_0 = \rho$