ECE560: Computer Systems Performance Evaluation

Lecture #9 –
Stochastic Processes
(Part I)

Instructor: Dr. Liudong Xing
Spring 2024

Administration Issues

- Homework #2 solution posted
- Homework #3
  - Due: **February 21, Wednesday**
- Project proposal (refer to Guidelines)
  - Due: **February 23, Friday**
- Lecture #8 (Statistical Inference)
  - **Self-study lecture**

Review of Lectures #5-7

In last three lectures, we reviewed

- Probability
  - Sample spaces & events
  - Axioms of probability
  - Field, σ-field, and probability measure
  - Conditional probability, law of total probability, Bayes' Formula
  - Independence (pair-wise, mutually)
- Random Variable (*r.v.*)
  - Basic concepts
  - Discrete & continuous r.vs
    - Concepts, parameters, examples
  - Jointly distributed r.v.s
    - Concepts, parameters, important functions
    - Maximum property

$$F_Y(y) = F_{X1}(y)F_{X2}(Y)...F_{Xn}(y) \quad \forall y \in R$$

- Minimum property

$$F_Y(y) = 1-(1-F_{X1}(y))(1-F_{X2}(Y))...(1-F_{Xn}(y)) \quad \forall y \in R$$

Topics

- **Basic concepts of stochastic processes**
- Important stochastic processes
  - Counting processes
  - Poisson processes
  - Birth-and-death processes
  - Markov processes *Lecture #10*

Chapter 4 in Allen's book

Stochastic Processes (1)

- Definition:

A family of r.v.'s $\{X(t) \mid t \in T\}$ that is indexed by a parameter t (such as time) is known as a "stochastic process" (or chance/random process)

- Index set T: the set of all possible values of t
  - Each element of T is referred to as a parameter
- State space: the set of all possible values assumed by r.v.'s X(t)
  - Each of these values is called a state of the SP

Stochastic Processes (2)

- Classification

		Index Set T: Discrete	Index Set T: Continuous
State Space	Discrete	**1. Discrete parameter Discrete state**	**2. Continuous parameter Discrete state**
State Space	Continuous	**3**	**4**

- Since parameter is usually referred to as "time", and discrete-state process is often referred to as a "chain",
  1. Discrete **time**, stochastic **chain**
  2. Continuous **time,** stochastic **chain**
  3. Discrete **time**, continuous state process
  4. Continuous **time**, continuous state process

Example Stochastic Processes

- Number of commands received by a time-sharing system during time interval (0 , t1)
  - {X(t), 0 < t < t1}
  - Type?
- Number of students attending the *n*th lecture
  - { X_n, n = 1, 2, 3, 4, . . .}
  - Type?
- Average time to run a batch job at a computing center on the *n*th day of the week.
  - { X_n, n = 1, 2, 3, 4, . . ., 7}
  - Type?
- The waiting time of an inquiry message that arrives at time t, until processing is begun
  - { X(t), t >= 0}
  - Type?

Stochastic Processes (3)

- Intuition behind the concept of a SP

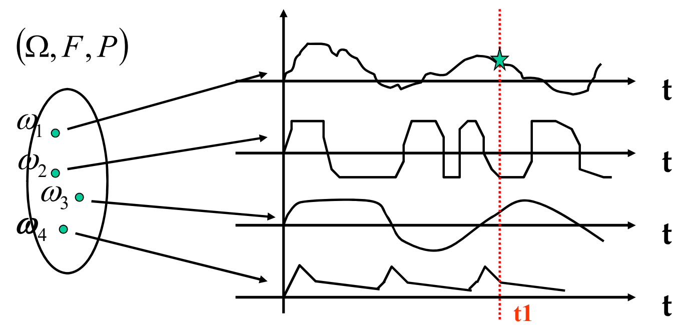

	ω fixed	ω varied
t fixed	**A real #**	$X_{t_1}(\omega) = a \ \ r.v.$: $X_{t_1}(\omega) : \Omega \to R$
t varied	$X_t(\omega_1)$= **a "*sample path*" or "*realization*" of process X @ ω_1 = a deterministic func. of t**	**A family of r.v.'s = SP**

Agenda

✓ Basic concepts of stochastic processes

- Important stochastic processes
  - **Counting processes**
  - Poisson processes
  - Birth-and-death processes
  - Markov processes *Lecture #10*

Chapter 4 in Allen's book

Important Stochastic Processes (1)

- **Counting processes**
  - A counting process is referred to as a SP representing the number of events of some kind that have occurred (after time 0 but before time t) by the time t:

    $\{N(t), t \geq 0\}$: Continuous parameter, discrete state
  - Events can be:
    - Arrival of an inquiry at the central processing system of a computer
    - A phone call to an airline reservation center
    - The occurrence of a hw or sw failure in a computer system
    - Job arrivals to a file server
    - etc.

Counting Processes (Cont'd)

- Formal definition: A SP $\{N(t), t \geq 0\}$ constitutes a **counting process** provided that
  - $N(0)=0$
  - $N(t)$ assumes only non-negative integer values
  - $s < t$ implies that $N(s) \leq N(t)$
  - $N(t) - N(s)$ is the number of events that have occurred in the interval $(s, t]$ (after but not later than t)

Agenda

✓ Basic concepts of stochastic processes

- Important stochastic processes
  - √ Counting processes
  - **Poisson processes**
  - Birth-and-death processes
  - Markov processes *Lecture #10*

Chapter 4 in Allen's book

o(h) notation

- Definition 4.1.2: a function *f* is o(*h*) (read "f is little –oh of h" and written "f=o(h)") if

$$\lim_{h \to 0} \frac{f(h)}{h} = 0 \quad \text{or} \quad \lim_{h \to \infty} \frac{f(h)}{h} = \infty$$

- Any quantity having an order of magnitude bigger than *h*

Remarks on *o(h)*

- Used in algorithm complexity evaluation in computer science
- The o(h) is actually ω(h)
- The o(h) in computer science is:
  - any quantity having an order of magnitude **smaller** than h

$$\lim_{h \to \infty} \frac{f(h)}{h} = 0$$

But we use definition of o(h) in the textbook!

o(h) notation (Cont'd)

– Properties of o(h)

- If f is o(h), g is o(h), then f+g is o(h)
- If f is o(h), c is a const., then cf is o(h)
- Any finite linear combination of functions, each of which is o(h), is also, o(h), i.e., $c_1, c_2, \cdots c_n$ are constants, $f_1, f_2, \cdots f_n$ are *n* functions, each of them is o(h), Then,

$$\sum_{i=1}^{n} c_i f_i \quad \text{is o(h)}$$

o(h) notation (Cont'd)

– Hands-on examples

- f(h)=h^2
- f(h)=h
- f(h)= h^n , n > 1
- f(h)=sin(h)
- f(h)=c, c is a non-zero constant

$o(h)$ notation (Cont'd)

– Application of $o(h)$

Suppose X is an Exp. r.v. with parameter λ, its c.d.f.:

$$F_X(h) = P\{X \le h\} = 1 - e^{-\lambda h}$$

What is the prob. that X is less than or equal to t+h given that it is greater than t?

Poisson Processes

- Let $\{N(t), t \ge 0\}$ be a Poisson process with rate $\lambda > 0$, then
  - It is a continuous parameter discrete state SP
  - It counts the # of occurrence of some type of events
  - It could also be interpreted as an arrival of some entity at an average rate λ
  - For each N(t)=Y, it is a Poisson r.v. describing the # of events occurring in any time interval of length *t* (e.g., (0, t], (s, s+t]), and it has *pmf* of

$$P\{Y = n\} = P\{N(t) = n\}$$

$$= P_n(t) = e^{-\lambda t}\frac{(\lambda t)^n}{n!}, n = 0, 1, 2\cdots$$

$$E[Y] = E[N(t)] = \lambda t$$

Poisson Processes (Cont'd)

- <u>Formal Definition</u>
  - A Poisson process with rate $\lambda > 0$ is a counting process $\{N(t), t \ge 0\}$ satisfying the following conditions:
    1. N(0)=0
    2. The process has independent Increments (events occurring in non-overlapping intervals are mutually independent)
    3. The process has stationary increments (The distribution of the # of events in any interval of time depends only on the length of the interval and not on the time origin)
    4. In any time interval of length h, we can have

$$\begin{cases} P\{N(h)=1\} = \Pr\{\text{exactly one event occurs}\} = \lambda h + o(h) \\ P\{N(h)\ge 2\} = \Pr\{\text{more than one event occur}\} = o(h) \\ P\{N(h)=0\} = \Pr\{\text{no event occurs}\} = 1-(\lambda h + o(h)) - o(h) \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad = 1 - \lambda h + o(h) \end{cases}$$

Poisson Processes (Cont'd)

- Remarks
  - Let $\{N(t), t \ge 0\}$ be a Poisson process with rate λ>0, and let $0 < t_1 < t_2 < \cdots$ be successive occurrence times of events, and let the inter-arrival time (i.e. the time between the occurrence of 2 consecutive events) be defined by $\tau_1 = t_1, \tau_2 = t_2 - t_1, \cdots$ then inter-arrival times $\{\tau_n\}$ are mutually indep. identically distributed exp. r.v.'s, each with mean 1/λ (***Theorem 4.2.2***)
  - Let $\{N(t), t \ge 0\}$ be a counting process s.t. the inter-arrival times of events $\{\tau_n\}$ are indep. identically distributed exp. r.v.'s, each with avg. value 1/λ, then $\{N(t), t \ge 0\}$ is a Poisson process with rate λ (***Theorem 4.2.3***)

\# of arrivals ~ Poisson $\equiv$ inter-arrival time ~ Exp.

Superposition Property

- Superposition of independent Poisson processes
  - If two Poisson streams are merged, the result is a Poisson stream with the rate equal to the sum of the input rates

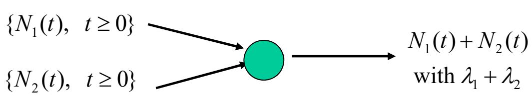

  - In general, the superposition of n indep. Poisson processes w/ respective rates $\lambda_1, \lambda_2, \cdots \lambda_n$ is also a Poisson process with rate $\lambda = \sum_{i=1}^{n} \lambda_i$

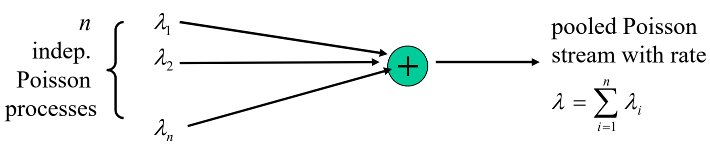

(based on the proposition that the sum of n indep. Poisson r.v.'s is still a Poisson r.v.)

Decomposition Property

- Decomposition of a Poisson process
  - If a Poisson streams is divided into 2 streams, w/ each event going to stream A w/ prob. P_A and stream B w/ prob. P_B, the resulting streams are Poisson with rates $\lambda P_A, \lambda P_B$ respectively

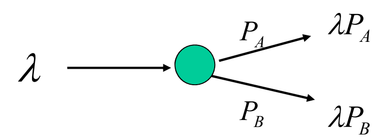

  - In general,

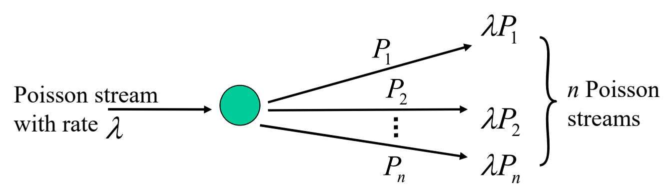

A Poisson process w/ rate λ branches out into n output paths

Hands-on Problems

A computer center has a large number of separate system components that may fail (terminals, tape drives, disks, printers, sensors, CPUs, etc.), without bringing the entire system down.

There are, on the average, 0.6 failures per day. Failures can be represented by a **Poisson process with rate $\lambda = 0.6$ (per day)**. The time between failures is observed to be exponentially distributed.

Answer the following questions:

- What is the mean time between failures?
- The number of failures in a "t"-day interval(Y_t) has the Poisson distribution with a mean of $\lambda t=0.6t$. What is the probability of exactly one failure in a 24-hour period?

Hands-on Problems (Cont'd)

- What is the probability of less than 5 failures in a week?
- Starting from a random point in time, what is the probability that no failure will occur during the next 24 hours?
- Suppose exactly 24 hours has elapsed with no failures. What is the expected time until the next failure?
- Four out of every five failures is a terminal problem, with equal probability on each failure. What is the process describing the terminal failure?
- What is the mean time between terminal failures?
- What is the probability of *k* terminal failures in a t-day interval?

Agenda

- ✓ Basic concepts of stochastic processes
- • Important stochastic processes
  - √ Counting processes
  - √ Poisson processes
  - – **Birth-and-death processes**
  - – Markov processes *Lecture #10*

Chapter 4 in Allen's book

Birth-and-Death Processes: Definition

- Let $\{X(t), t \geq 0\}$ be a continuous parameter stochastic process with discrete state space 0, 1, 2, …..

  $X(t) = n, n=0, 1, 2, \ldots$ means that
  - The system has a population of *n* elements/customers at time *t*
  - The system is in state $E_n \in \{0,1,2,\ldots\}$ at time *t*
- There exist nonnegative birth and death rates:

  $\{\lambda_n, n = 0,1,2,\ldots\}$ $\{\mu_n, n = 0,1,2,\ldots\}$

- State changes are only in increments of ± 1 and the value of E_n is never negative.

  $E_n \Rightarrow E_{n+1}$ if $n \geq 0$ or $E_n \Rightarrow E_{n-1}$ if $n \geq 1$

- If the system is in state E_n at time t, the probability of a transition to E_{n+1} during the interval (t, t+h] is $\lambda_n h + o(h)$, and to E_{n-1} is $\mu_n h + o(h)$

- The probability of more than 1 transition during an interval of length of h is *o(h)*.

Then, $\{X(t), t \geq 0\}$ is a birth-and-death process.

Birth-and-Death Processes

- State-transition diagram
- Differential-difference equations for P[$X(t)=n$]=$P_n(t)$:

$$\begin{cases} \dfrac{dP_n(t)}{dt} = -(\lambda_n + \mu_n)P_n(t) + \lambda_{n-1}P_{n-1}(t) + \mu_{n+1}P_{n+1}(t), \quad n \ge 1 \\ \dfrac{dP_0(t)}{dt} = -\lambda_0 P_0(t) + \mu_1 P_1(t) \\ P_i(0) = 1, \; P_j(0) = 0, \; j \ne i \end{cases}$$

Initial condition (system is initially in state Ei)

- Solving the above infinite set of differential-difference equations w/ the initial conditions, can give you the solution for $P_n(t)$ for all n, all t.

Remarks:

- **But analytically the time-dependent solutions are very difficult to obtain, except for some very special cases – pure-birth process**
- **People uses "steady-steady solutions"**

Special Case: Pure-Birth Processes

- Any Birth-and-Death process for which all death rates μ_n are 0 is called a ~
  - Example: for the pure- birth process with

$$\begin{cases} \lambda_n = \lambda > 0, \; \mu_n = 0 \\ \textbf{initial conditions: } p_0(0) = 1, \; p_j(0) = 0, \; j > 0 \end{cases}$$

We have the following set of equations:

$$\begin{cases} \dfrac{dP_n(t)}{dt} = -\lambda P_n(t) + \lambda P_{n-1}(t), \quad n \ge 1 \\ \dfrac{dP_0(t)}{dt} = -\lambda P_0(t) \end{cases}$$

The solution of the set of equations satisfying the given initial condition is:

$$P_n(t) = e^{-\lambda t}\frac{(\lambda t)^n}{n!}, \quad n \ge 0, \quad t \ge 0$$

A Poisson process is a pure-birth process with a constant birth rate

Steady-State Solutions

- In general, finding the time-dep. solutions is very difficult. However,
  - If $\lim_{t\to\infty} p_n(t) = p_n$ for each n, then the system is in statistical equilibrium/stationary/in the steady-state!
  - A steady-state (Equilibrium) solution exists:
    - Can be obtained by taking limits as $t \to \infty$ on both side of time-dep. equations and using the fact that

$$\begin{cases} \lim_{t\to\infty} p_n(t) = p_n \\ \lim_{t\to\infty} \frac{dp_n(t)}{dt} = 0 \qquad \forall \mathrm{n} = 0,\ 1,\ 2\cdots \end{cases}$$

    - That is, we obtain a set of difference equations:

$$\begin{cases} 0 = -(\lambda_n + \mu_n)\boldsymbol{P}_n + \lambda_{n-1}\boldsymbol{P}_{n-1} + \mu_{n+1}\boldsymbol{P}_{n+1}, \quad \mathbf{n \geq 1} \\ 0 = -\lambda_0 \boldsymbol{P}_0 + \mu_1 \boldsymbol{P}_1 \end{cases}$$

$$\Updownarrow$$

$$\begin{cases} \lambda_{n-1}P_{n-1} + \mu_{n+1}P_{n+1} = (\lambda_n + \mu_n)P_n \quad (1) \\ \mu_1 P_1 = \lambda_0 P_0 \quad (2) \end{cases}$$

-- Called "Balance Equations"

Steady-State Solutions (continued)

- **Balance equations**: when a system reaches equilibrium, for each state En, we have

  Rate of entering En = Rate of leaving En

- Also, $\sum_{n=0}^{\infty} p_n = 1 \quad (3)$
- Now it should be easy for us to derive the steady state solutions by using (1) to (3),

$$p_1 = \frac{\lambda_0}{\mu_1} p_0\ ,\ p_2 = \frac{\lambda_1}{\mu_2} p_1 = \frac{\lambda_1}{\mu_2}\frac{\lambda_0}{\mu_1} p_0 \cdots$$

$$p_n = \frac{\lambda_{\mathrm{n-1}}\lambda_{\mathrm{n-2}}\cdots\lambda_1\lambda_0}{\mu_{\mathrm{n}}\mu_{\mathrm{n-1}}\cdots\mu_2\mu_1} p_0 \qquad n \geq 1$$

- Substitute them into (3),

$$\therefore \quad p_0\left[1 + \frac{\lambda_0}{\mu_1} + \frac{\lambda_1}{\mu_2}\frac{\lambda_0}{\mu_1} + \cdots + \frac{\lambda_{\mathrm{n-1}}\lambda_{\mathrm{n-2}}\cdots\lambda_1\lambda_0}{\mu_{\mathrm{n}}\mu_{\mathrm{n-1}}\cdots\mu_2\mu_1}\right] = 1$$

$\Rightarrow$ $p_0 =$

$p_n =$

Birth-and-Death Processes: An Example

- Consider a queuing system with one server and no waiting line. And assume
  - Poisson arrivals with rate λ
  - Exponential service with rate μ

  The state transition diagram is:

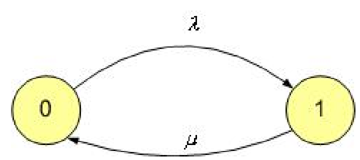

Find P_0, P_1?

Next Topics

- Stochastic Processes (Cont'd)
  - Markov processes

Things to Do

- Read Allen's Ch. 4
- Project proposal
  - Due **Friday, February 23**