Service System Arrival Processes

From Service System Data to Arrival Process Models

IEOR 4615, Service Engineering, Professor Whitt

Lecture 16, March 24 and 26, 2015

OUTLINE

The standard model for an arrival process in a service system is a nonhomogeneous Poisson process (NHPP).

- Stochastic Point Process Models (Definitions)
  - Three Equivalent Representations
  - The Arrival Rate
- Plots of Call Center Arrival Counts (Looking at Data)
  - Different Time Scales
  - Identifying Predictable and Unpredictable Variability
- Poisson Review (Theory)
- More on Poisson Process Models

Stochastic Point Processes: 3 Equivalent Representations

- **Arrival Times**, the location of point k; A_k.
- **Interarrival Times**, the intervals between successive points:

  $X_k = A_k - A_{k-1}, k \geq 1.$
- **Counting Process**, the number of points in the interval $[0, t]$

  $A(t) = \max\{k \geq 0 : A_k \leq t\}, t \geq 0.$
- **Basic Inverse Relation**:
  - The two stochastic processes $\{A_k : k \geq 0\}$ and $\{A(t) : t \geq 0\}$ are **inverse processes**. (You see them both from a plot of either one.)

Sample Path of a Counting Process

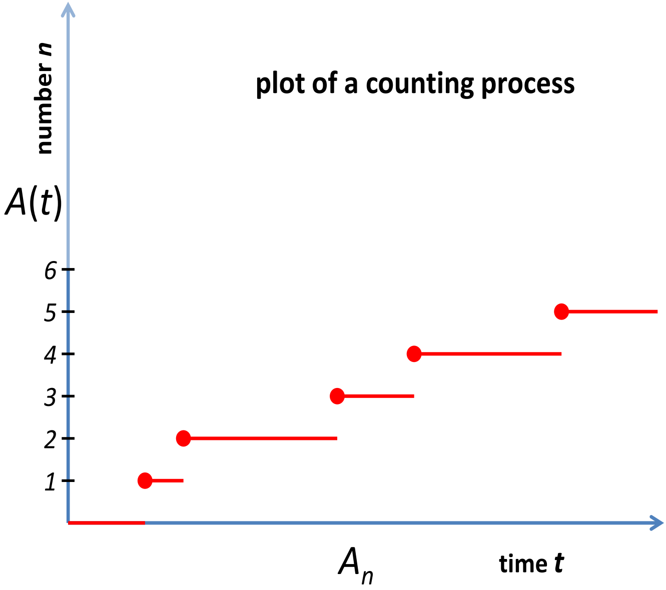

The Arrival Rate Function

Let $\{A(t) : t \geq 0\}$ be a counting process. An important partial characterization is the **cumulative arrival rate function**

$$\Lambda(t) \equiv E[A(t)], \quad t \geq 0.$$

We assume that $\Lambda(t)$ is differentiable. Then

$$\Lambda(t) = \int_0^t \lambda(s)\, ds, \quad t \geq 0.$$

$\lambda(t)$ is the **arrival rate function**.

Back to OUTLINE

The standard model for an arrival process in a service system is a nonhomogeneous Poisson process (NHPP).

- Stochastic Point Process Models (Definitions)
- Plots of Call Center Arrival Counts (Looking at Data)
  - Different Time Scales
  - Identifying Predictable and Unpredictable Variability
- Poisson Review (Theory)
- More on Poisson Process Models

Looking at Call Center Arrivals: Different Time Scales

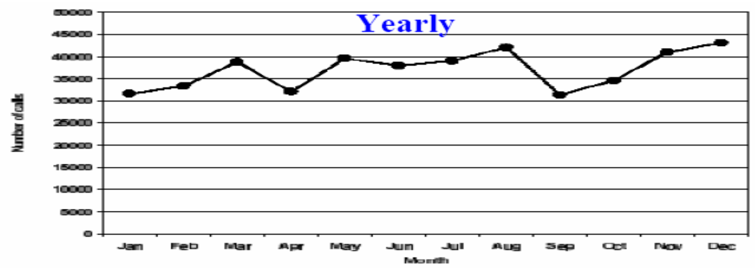

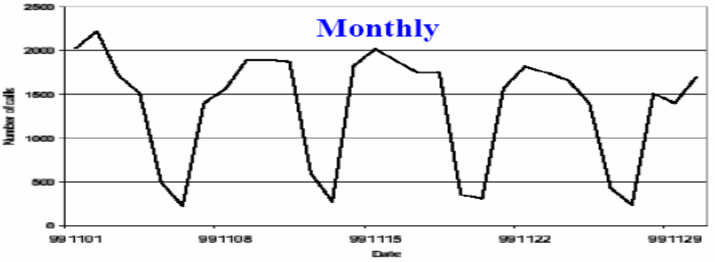

Notes on the Yearly and Monthly Plots

- The yearly plot shows monthly totals over one year.
- The monthly plot shows daily totals over one month.
- The big dips in the monthly plot are the weekends.
- All the variation has to be predictable deterministic variation.
    - For the months in the year, the counts range from 32,000 to 44,000.
    - Assuming totals are Poisson, they are approximately normal with variance equal to the mean.
    - when mean is about 40,000, the standard deviation is $\sqrt{40,000} = 200$.
    - The actual fluctuations are much greater than would be the case for Poisson counts with a fixed mean.

Looking at Call Center Arrivals: Different Time Scales

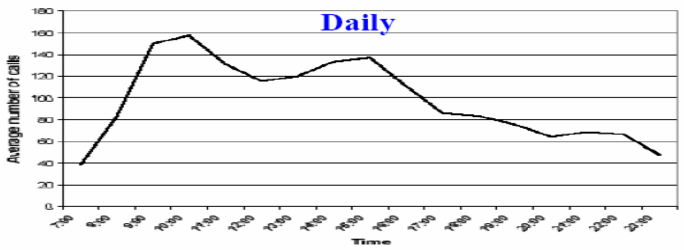

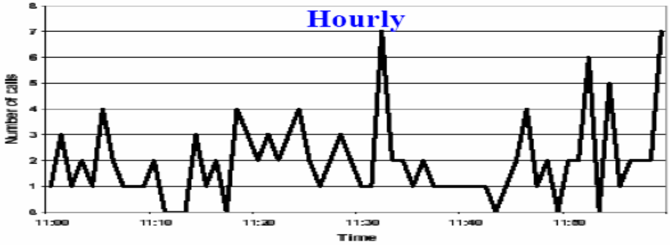

Notes on the Daily and Hourly Plots

- The daily plot shows hourly totals over one day.
- The hourly plot shows totals for minutes over the hour.
- The variation in the daily plot has to be mostly deterministic variation.
  - For the hours in the day, the counts range from 40 to 160.
  - when mean is about 100, the standard deviation is $\sqrt{100} = 10$.
  - The actual fluctuations are much greater than would be the case for Poisson counts with a fixed mean.
- For the plots within an hour, we see genuine stochastic variability.
- If the hourly total is 150, then the mean number in each minute is $150/60 = 2.5$. The plot looks roughly like i.i.d Poisson random variables with mean 2.5.

Arrival Rate Over the Day in 1959 and 1995

Q-Science

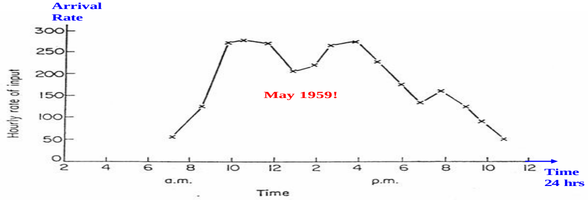

Fig. 15.1 The variation in the hourly input rates of reservations calls during a typical day (in May 1959)

(Lee A.M., Applied Q-Th)

1995 Help Desk and Customer Support Practices Report

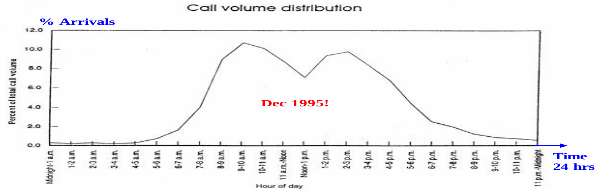

Number of respondents = 522

(Help Desk Institute)

Identifying the Predictable and Unpredictable Variability

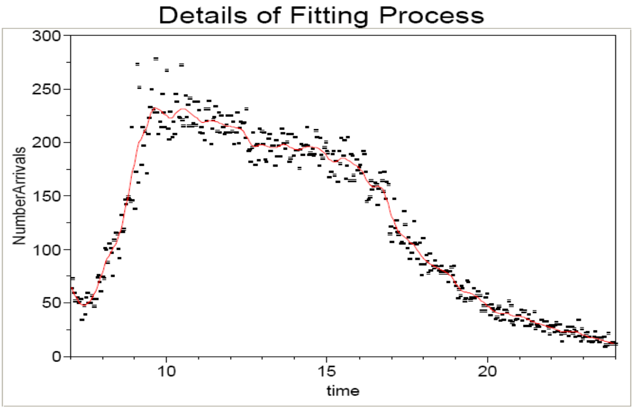

Plot for Aug 9 (Fri.)

- Divide day (7am to midnight) into time intervals of 150 seconds (=2½ minutes)
- Count number arrivals in each interval, and make scatterplot
- Fit using a nonparametric regression smoothing

Look at Multiple Days: IID NHPP's?

Number of Calls at a U.S. bank.

Mondays. March 2002-August 2002.

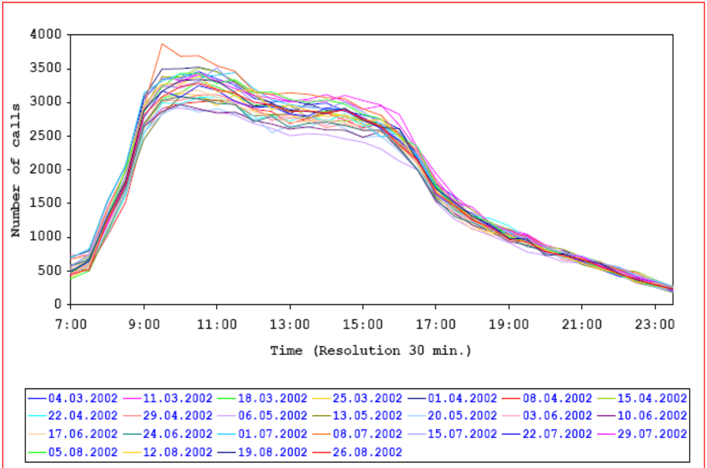

Estimate Day-To-Day Variation: 25 Mondays Overall

- Look at: **13:00-13:30**.
  - **Quick Rough Analysis**:
    - **range**: **[2500, 3200]**, **mean** ≈ 2850.
    - **5 STD DEV** ≈ 700, **STD DEV** ≈ 140.
    - **variance** $\approx 19,600 >> 2850$ **Too large!**.
  - **Actual Data Analysis**:
    - **actual sample mean = 2,842**.
    - **actual sample variance = 24,539** $>>$ **2, 842**. Thus, **Too large!**.
- Look at: **17:00-17:30**.
  - **actual sample mean = 1,705**.
  - **actual sample variance = 10,356** $>>$ **1,705**. Again, **Too large!**.

Separating Hourly Rate from the Daily Total: Normalize

Mondays

Plot shows (spline-smooth of)

$$\text{Normalized Arrivals}_{\text{given day}} = \frac{\text{Arrivals per Hour}_{\text{that day}}}{\text{Average}_{\text{that day}} \ (\text{arriv's. per hour})}$$

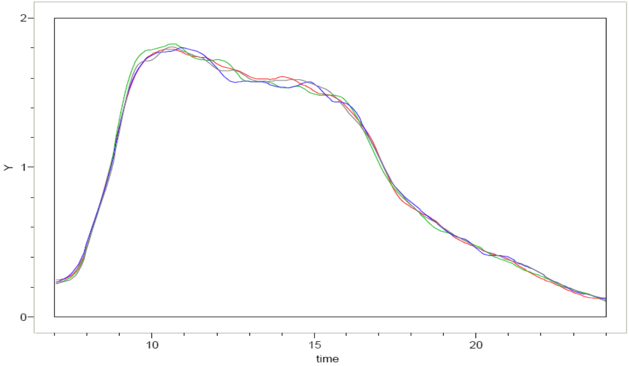

The variable **"time"** is on a 24 hour clock.

= 8/05/02

= 8/12/02

= 8/19/02

= 8/26/02

Repeat on Fridays

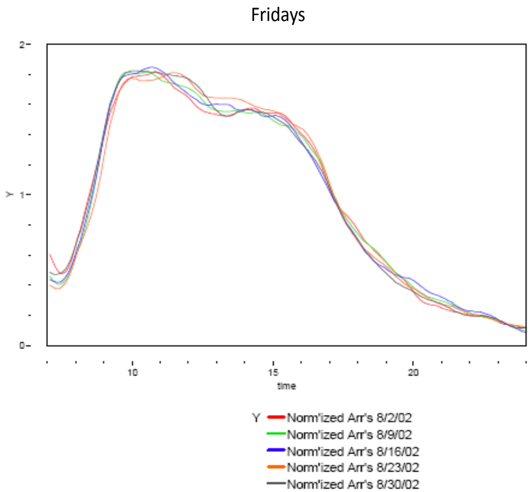

Note Dip between 7 & 7:30 am

Compare Mondays and Fridays

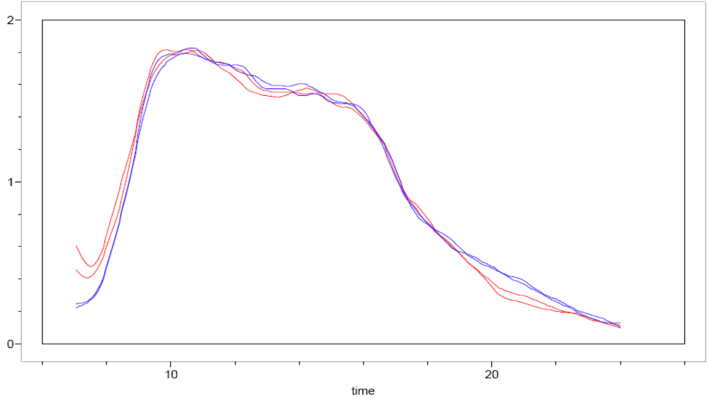

Friday

Has characteristic beginning dip, AND

Daily volume shifted to (slightly) earlier in the day

Y — Norm'ized Arr's 8/2/02

— Norm'ized Arr's 8/9/02

— Norm'ized Pred for 8/5/02

— Norm'ize Pred for 8/12/02

Back to the OUTLINE

The standard model for an arrival process in a service system is a nonhomogeneous Poisson process (NHPP).

- Stochastic Point Process Models (Definitions)
- Plots of Call Center Arrival Counts (Looking at Data)
- Poisson Review (Theory)
  - The Poisson Distribtuon
  - The Poisson Process (PP) and NHPP
- More on Poisson Process Models

Poisson Review

- Random Variable X with the Poisson Distribution [Ross §2.2.4]:

$$P(X = k) = \frac{e^{-m} m^k}{k!}, \quad k = 0, 1, 2, \ldots$$

- Mean and Variance: $E[X] = m$ and $Var(X) = m$
- Moment Generating Function (mgf) [Ross §2.6]. Since

$e^x = \sum_{k=0}^{\infty} (x^k / k!)$,

$$\psi_X(t) = E[e^{tX}] = \sum_{k=0}^{\infty} \frac{e^{-m} m^k e^{tk}}{k!} = e^{-m} \sum_{k=0}^{\infty} \frac{(me^t)^k}{k!} = e^{m(e^t - 1)}$$

- Differentiate to get moments:

$$\frac{d\psi_X(t)}{dt} = me^t e^{m(e^t-1)}, \quad \dot{\psi}_X(0) = m, \quad \ddot{\psi}_X(0) = m^2 + m$$

Poisson Distribution As Limit of Binomial Distribution

- Random Variable X with the Binomial Distribution [Ross §2.2.2]:

$$P(X=k) \quad = \quad b(k;n,p) = \frac{n!}{(n-k)!k!}p^k(1-p)^{n-k}, \quad k = 0,1,2,\ldots$$

- Let $n \to \infty$, $p \to 0$ and $np \to m$, i.e., set $p = m/n$ [Ross §2.2.4]:

$$\begin{aligned} P(X=k) &= \frac{n!}{(n-k)!k!}(m/n)^k(1-(m/n))^{n-k} \\ &= \left(\frac{n(n-1)\cdots(n-k+1)}{n^k}\right)\left(\frac{m^k}{k!}\right)\left(\frac{(1-(m/n))^n}{(1-(m/n))^k}\right) \\ &\to (1)\left(\frac{m^k}{k!}\right)\left(\frac{e^{-m}}{1}\right) = \frac{e^{-m}m^k}{k!} \quad \text{as} \quad n \to \infty. \end{aligned}$$

Normal Approximation

- If the variable X is Poisson with mean m, where m is large, then X has approximately a normal distribution, i.e., $X \approx N(m, m)$.
- Why: Use CLT with the property:
- If X_1 and X_2 are independent Poisson variables with means m_1 and m_2, then $X_1 + X_2$ is Poisson with Mean $m_1 + m_2$ [Ross, Example 2.37].
- Use mgf's:

$$\begin{aligned}
\Psi_{X_1+X_2}(t) &= E[e^{t(X_1+X_2)}] = E[e^{tX_1}e^{tX_2}] = E[e^{tX_1}]E[e^{tX_2}] \\
&= \Psi_{X_1}(t)\psi_{X_2}(t) = e^{m_1(e^t-1)}e^{m_2(e^t-1)} = e^{(m_1+m_2)(e^t-1)}.
\end{aligned}$$

A Poisson Counting Process

- Let $A(t)$ count the number of points in the interval $[0, t]$ [Ross §5.3].

  (Think of counting arrivals.)

- The stochastic process $\{A(t) : t \geq 0\}$ is a Poisson process with rate λ if
  - the process has stationary and independent increments, and
  - each increment has a Poisson distribution; e.g., the increment

    $A(s + t) - A(s) \overset{d}{=} A(t) - A(0) \overset{d}{=} A(t)$ has a Poisson distribution with

    mean λt:

$$P(A(t) = k) = \frac{e^{-\lambda t}(\lambda t)^k}{k!}, \quad k = 0, 1, 2, \ldots$$

A Nonhomogeneous Poisson Process (NHPP)

- Let $A(t)$ count the number of points in the interval $[0, t]$ [Ross §5.3].
- The stochastic process $\{A(t) : t \geq 0\}$ is an NHPP with time-varying rate $\lambda(t)$ if
    - the process has independent increments, and
    - each increment has a Poisson distribution; in particular,

$$P(A(s+t) - A(s) = k) = \frac{e^{-m(s,s+t)}(m(s, s_t))^k}{k!}, \quad k = 0, 1, 2, \ldots$$

where the mean of the increment $A(s+t) - A(s)$ is

$$E[A(s+t) - A(s)] = m(s, s+t) = \int_s^{s+t} \lambda(u)\, du.$$

For a short interval

- Let $\{A(t) : t \geq 0\}$ be an NHPP with time-varying arrival rate $\lambda(t)$,
- For a small $\epsilon > 0$,

$$\begin{aligned} P(A(t+\epsilon) - A(t) = 0) &\approx e^{-\lambda(t)\epsilon} \approx 1 - \lambda(t)\epsilon \\ P(A(t+\epsilon) - A(t) = 1) &\approx \lambda(t)\epsilon e^{-\lambda(t)\epsilon} \approx \lambda(t)\epsilon \\ P(A(t+\epsilon) - A(t) \geq 2) &\approx O(\epsilon^2) \quad \text{(very small)} \end{aligned}$$

- Hence, For a small $\epsilon > 0$, $A(t+\epsilon) - A(t)$ is approximately Bernoulli.

(Poisson approximation for binomial in reverse)

Back to the OUTLINE

The standard model for an arrival process in a service system is a nonhomogeneous Poisson process (NHPP).

- Stochastic Point Process Models (Definitions)
- Plots of Call Center Arrival Counts (Looking at Data)
- Poisson Review (Theory)
- More on Poisson Process Models
    - Review of NHPP and PP
    - Superpositions and the Palm-Khintchine theorem
    - When is a NHPP or PP a Good Model?
    - Relating NHPP's and PP's

A Nonhomogeneous Poisson Process (NHPP)

- **no batches**: Arrival occur one at a time
- **Poisson distribution**: $P(A(t) = k) = \frac{e^{-m(t)} m(t)^k}{k!}$
- **mean is the integral of the arrival rate**: $m(t) = \int_0^t \lambda(s)\, ds$, $t \geq 0$.
- **independent increments**: For all $k \geq 2$ and $0 \leq t_1 \leq t_2 \leq t_3 \leq \cdots t_{2k}$, $A(t_2) - A(t_1), A(t_4) - A(t_3), \ldots, A(t_{2k}) - A(t_{2k-1})$ are k independent random variables.
- **Poisson process** is the special case in which: $\lambda(t) = \lambda$ (constant). Then the process has **stationary increments**.

Definitions of a Poisson Process (PP)

- a **renewal process** (interarrival times i.i.d.) with an exponential distribution (having mean $1/\lambda$).
- a **Lévy process** (process with stationary and independent increments) with unit jumps.
- a **pure-birth process** (a birth-and-death CTMC with 0 death rate) having constant birth rate λ.
- an **NHPP**: with $m(t) = \Lambda(t) = \lambda t$, $t \geq 0$.
- Each definition implies, but does not assume, the Poisson distribution: $P(A(t) = k) = \frac{e^{-\lambda t}(\lambda t)^k}{k!}$, $t \geq 0$.

Properties of a NHPP, in contrast to a PP

- The interarrival times are in general **not** independent.
- The interarrival times are in general **not** stationary.
- The interarrival times are in general **not** exponentially distributed.
- The counts over intervals are still Poisson.
- The increments are still independent. (For all $k \geq 2$ and $0 \leq t_1 \leq t_2 \leq t_3 \leq \cdots t_{2k}$, $A(t_2) - A(t_1), A(t_4) - A(t_3), \ldots,$ $A(t_{2k}) - A(t_{2k-1})$ are k independent random variables, as stated two slides back.)

A Superposition of Point Processes: Good News

A superposition of n point processes can be represented as the sum of the corresponding counting processes, i.e., $A(t) \equiv A_1(t) + \cdots A_n(t)$, $t \geq 0$, where $\{A_i(t) : t \geq 0\}$ is a counting process for each i.

Theorem

(Good News) The superposition of n independent NHPP's (PP's) is itself an NHPP (PP) with an arrival rate function equal to the sum of the component arrival rate functions.

A Superposition of Point Processes: Bad News

Theorem

(Bad News) The superposition of n independent stationary (ordinary) renewal processes is itself either a stationary or ordinary renewal process if and only if all the component processes and the superposition process are PP's.

Theoretical Justification for Poisson: More Good News

Theorem

(The Law of Rare Events) The superposition of n independent i.i.d. nonstationary (stationary) point processes with intensities $\lambda(t)/n$ (λ/n) *converges in distribution to an NHPP (PP) with arrival rate function* $\lambda(t)$ (λ) *as* $n \to \infty$.

(the Palm-Khintchine theorem)

Proof of the Law of Rare Events

The essential argument: classical Poisson approximation of the binomial distribution, as on slide 8. If the probability of heads in one toss of a biased coin is p, then the number of heads in n independent coin tosses, S_n, has the binomial distribution, i.e.,

$$P(S_n = k) = b(k; n, p) = \frac{n!}{k!(n-k)!} p^k (1-p)^{(n-k)}.$$

Theorem

(*Poisson as the limit of the binomial distribution*) *If* $n \to \infty$ *and* $p \to 0$ *so that* $np \to m > 0$, *then*

$$P(S_n = k) = b(k; n, p) \to \frac{e^{-m} m^k}{k!} \quad as \quad n \to \infty.$$

Examples Where PP and NHPP Are Suspect

- Likely to be smoother, **less variable or less bursty**:
  - **scheduled arrivals**, as at doctor's office.
  - **enforced separation**, as in landings at airports.
  - **interarrival times** are **sums** of independent steps.
- Likely to be **more variable or more bursty**
  - **overflow processes**, in finite capacity systems, because they tend to occur in clumps, when the main system is overloaded.
  - **batch arrivals**, as in arrivals to amusement parks, arrivals at hostital ED because of accident, where customers may use resources as individuals.
  - **interarrival times are mixtures of independent steps**.

Examples Where PP and NHPP Should Be Good

- **arrivals at call center**,
- **arrivals at supermarket or bank**,
- **walk-in arrivals at hospital ED**.

All of these examples involve **single customers acting independently of others**. Walk-in patients at hospital ED are likely to be **not scheduled**.

How to Simulate a PP or an NHPP

- **Simulate a PP** by generating i.i.d. exponential interarrival times X_n with the desired mean $1/\lambda$. That can be done with i.i.d. uniforms U_n on $[0, 1]$: let $X_n = -\log_e(1 - U_n)/\lambda$:

$$\begin{aligned} P(-\log_e(1 - U_n)/\lambda \le x) &= P(1 - U_n \ge e^{-\lambda x}) \\ &= P(U_n \le 1 - e^{-\lambda x}) = 1 - e^{-\lambda x}. \end{aligned}$$

- **Simulate a NHPP** with rate function $\lambda(t)$ by simulating a PP with rate $\bar{\lambda}$ (see above) such that $\lambda(t) \le \bar{\lambda}$, Let a point at time t in the PP be a point in the NHPP with probability $p(t) \equiv \lambda(t)/\bar{\lambda}$. (Use independent thinning property of NHPP.)

How to Construct an NHPP from a PP

Let $\{N(t) : t \geq 0\}$ be a rate-1 PP and let $\Lambda(t) \equiv \int_0^t \lambda(s)\, ds$ be the cumulative arrival rate function of the NHPP. Then let

$$A(t) = N(\Lambda(t)), \quad t \geq 0.$$

Theorem

The stochastic process $\{A(t) : t \geq 0\}$ defined above is a NHPP with the specified cumulative arrival rate function $\Lambda(t)$.

How to Construct a PP from a NHPP

Let $\{A(t) : t \geq 0\}$ be an NHPP with cumulative arrival rate function $\Lambda(t) \equiv \int_0^t \lambda(s)\, ds$. Then let $\Lambda^{-1}(t)$ be the **inverse function** of $\Lambda(t)$, defined by

$$\begin{aligned} \Lambda^{-1}(t) &= \inf\{s > 0 : \Lambda(s) = t\}, \quad t \geq 0, \quad \text{or} \\ \Lambda^{-1}(t) &= y \quad \text{where} \quad t = \Lambda(y) = \int_0^y \lambda(s)\, ds. \end{aligned}$$

Theorem

The stochastic process $\{N(t) : t \geq 0\}$ where $N(t) \equiv A(\Lambda^{-1}(t))$ for the NHPP $\{A(t) : t \geq 0\}$ with cumulative arrival rate function $\Lambda(t)$ is a rate-1 PP.

Basic References

- R. W. Hall. *Queueing Models for Services and Manufacturing*, Prentice-Hall, Englewood Cliffs, NJ, 1991. [The course textbook. This lecture relates to Chapters 3 and 6.]
- S. M. Ross. *Introduction to Probability Models*, 10th edition, Elsevier, Amsterdam, 2010. [Used in IEOR 3106 and 4106; see Chs. 5, 6, 8.]
- K. Sigman. *Stationary Marked Point Processes, An Intuitive Approach*, Chapman and Hall, New York, 1995. [More advanced, focusing on stationary point processes, in continuous and discrete time. Includes generalization to marks, thus can include service times together with the arrival process, and thus the full work brought to the system.]

Data & Statistical References

- M. Armony et al. Patient Flow in Hospitals: a Data-based Queueing-science Perspective. See course web page: http://www.columbia.edu/∼ww2040/4615S15/Surveys.html
- L. Brown et al. Statistical Analysis of a Telephone Call Center: A Queueing Science Perspective. *Journal of the American Statistical Association* (JASA) 100 (2005) 36-50. (see same web page)
- P. A. W. Lewis. Some Results on Tests for Poisson Processes. *Biometrika* 52 (1965) 67-77.

WW Data and Statistical References

- S-H. Kim and WW. Choosing Arrival Process Models for Service Systems: Tests of a Nonhomogeneous Poisson Process. *Naval Research Logistics*, vol. 61, No. 1, 2014, pp. 66-90.
- S-H. Kim and WW. Are Call Center and Hospital Arrivals Well Modeled by Nonhomogeneous Poisson Processes? *Manufacturing and Service Operations Management*, vol. 16, No. 3, 2014, pp. 464-480.
- S-H. Kim and WW. Poisson and Non-Poisson Properties in Appointment-Generated Arrival Processes: the Case of an Endocrinology Clinic. *Operations Research Letters*, vol 43, 2015, pp. 247-253.

Stationary Point Processes

- There are **two forms** of stationary point processes, defined in continuous time or discrete time; see Sigman (1995).
- A point process is said to be **stationary** in **continuous time** if the counting process $\{A(t) : t \geq 0\}$ has **stationary increments**. (That primarily means the arrival rate function is a constant function.) For example, an NHPP is a Poisson process that is also a stationary point process.
- However, for a general stationary point process (Not a PP), the increments typically are not independent and do not have a Poisson distribution.

What Makes a Good Model?

- What is the **purpose**?
  - Use as a **component of a queueing model**.
  - Use to help determine **staffing levels**, e.g., call center agents or hospital beds, nurses and doctors.
- Is the model **realistic**?
  - We need to **analyze arrival data**.
    - Look at **plots** and do **statistical analysis**.
- Can we use it to perform **useful analysis**?
  - Our answer: **POISSON PROCESSES**.
    - **Statistical tests** (for homogeneous and nonhomogeneous).
    - Estimate **arrival rate function**: **fitting** and **forecasting**.