5. Stochastic processes (1)

Contents

Stochastic processes (1)

- Consider some quantity in a teletraffic (or any) system
- It typically **evolves** in time **randomly**
  - Example 1: the number of occupied channels in a telephone link at time t or at the arrival time of the n^{th} customer
  - Example 2: the number of packets in the buffer of a statistical multiplexer at time t or at the arrival time of the n^{th} customer
- This kind of evolution is described by a **stochastic process**
  - At any individual time t (or n) the system can be described by a random variable
  - Thus, the stochastic process is a collection of random variables

Stochastic processes (2)

- **Definition**: A (real-valued) **stochastic process** $X = (X_t \mid t \in I)$ is a collection of random variables X_t
  - taking values in some (real-valued) set S, $X_t(\omega) \in S$, and
  - indexed by a real-valued (time) parameter $t \in I$.
- Stochastic processes are also called **random processes** (or just **processes)**
- The index set $I \subset \Re$ is called the **parameter space** of the process
- The value set $S \subset \Re$ is called the **state space** of the process
- **Note**: Sometimes notation X_t is used to refer to the whole stochastic process (instead of a single random variable related to the time t)

Stochastic processes (3)

- Each (individual) random variable X_t is a mapping from the sample space Ω into the real values $\Re$:

$$X_t : \Omega \to \Re, \quad \omega \mapsto X_t(\omega)$$

- Thus, a stochastic process X can be seen as a mapping from the sample space Ω into the set of real-valued functions $\Re^I$ (with $t \in I$ as an argument):

$$X : \Omega \to \Re^I, \quad \omega \mapsto X(\omega)$$

- Each sample point $\omega \in \Omega$ is associated with a real-valued function $X(\omega)$. Function $X(\omega)$ is called a **realization** (or a **path** or a **trajectory**) of the process.

Summary

- Given the sample point $\omega \in \Omega$
  - $X(\omega) = (X_t(\omega) \mid t \in I)$ is a real-valued function (of $t \in I$)
- Given the time index $t \in I$,
  - $X_t = (X_t(\omega) \mid \omega \in \Omega)$ is a random variable (as $\omega \in \Omega$)
- Given the sample point $\omega \in \Omega$ and the time index $t \in I$,
  - $X_t(\omega)$ is a real value

Example

- Consider traffic process $X = (X_t \mid t \in [0,T])$ in a link between two telephone exchanges during some time interval $[0,T]$
  - X_t denotes the number of occupied channels at time t
- Sample point $\omega \in \Omega$ tells us
  - what is the number X_0 of occupied channels at time 0,
  - what are the remaining holding times of the calls going on at time 0,
  - at what times new calls arrive, and
  - what are the holding times of these new calls.
- From this information, it is possible to construct the realization $X(\omega)$ of the traffic process X
  - Note that all the randomness in the process is included in the sample point ω
  - Given the sample point, the realization of the process is just a (deterministic) function of time

Traffic process

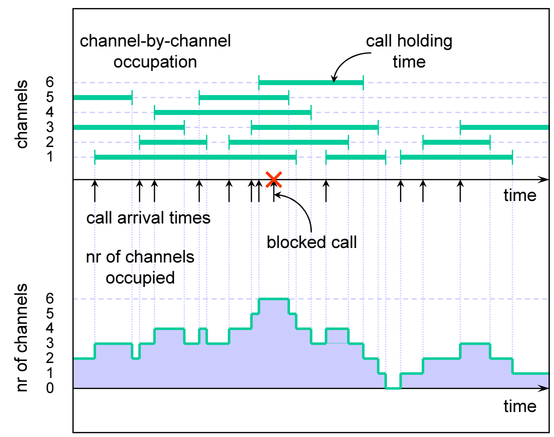

Categories of stochastic processes

- Reminder:
  - Parameter space: set I of indices $t \in I$
  - State space: set S of values $X_t(\omega) \in S$
- Categories:
  - Based on the parameter space:
    - **Discrete-time processes**: parameter space discrete
    - **Continuous-time processes**: parameter space continuous
  - Based on the state space:
    - **Discrete-state processes**: state space discrete
    - **Continuous-state processes**: state space continuous
- In this course we will concentrate on the discrete-state processes (with either a discrete or a continuous parameter space (time))
  - Typical processes describe the number of customers in a queueing system (the state space being thus $S = \{0,1,2,...\}$)

Examples

- Discrete-time, discrete-state processes
  - Example 1: the number of occupied channels in a telephone link at the arrival time of the n^{th} customer, $n = 1,2,...$
  - Example 2: the number of packets in the buffer of a router output link at the arrival time of the n^{th} customer, $n = 1,2,...$
- Continuous-time, discrete-state processes
  - Example 3: the number of occupied channels in a telephone link at time $t > 0$
  - Example 4: the number of packets in the buffer of router output link at time $t > 0$

Notation

- For a **discrete-time process**,
  - the parameter space is typically the set of positive integers, $I = \{1,2,\ldots\}$
  - Index t is then (often) replaced by n: X_n, $X_n(\omega)$
- For a **continuous-time process**,
  - the parameter space is typically either a finite interval, $I = [0, T]$, or all non-negative real values, $I = [0, \infty)$
  - In this case, index t is (often) written not as a subscript but in parentheses: $X(t)$, $X(t;\omega)$

Distribution

- The **stochastic characterization** of a stochastic process X is made by giving **all** possible **finite-dimensional distributions**

$$P\{X_{t_1} \le x_1, \ldots, X_{t_n} \le x_n\}$$

  where $t_1,\ldots, t_n \in I$, $x_1,\ldots, x_n \in S$ and $n = 1,2,...$
- In general, this is not an easy task because of **dependencies** between the random variables X_t (with different values of time t)
- For discrete-state processes it is sufficient to consider probabilities of the form

$$P\{X_{t_1} = x_1, \ldots, X_{t_n} = x_n\}$$

  - cf. discrete distributions

Dependence

- The most simple (but not so interesting) example of a stochastic process is such that all the random variables X_t are **independent** of each other. In this case

$$P\{X_{t_1} \le x_1,\ldots,X_{t_n} \le x_n\} = P\{X_{t_1} \le x_1\}\cdots P\{X_{t_n} \le x_n\}$$

- The most simple non-trivial example is a discrete state **Markov process**. In this case

$$P\{X_{t_1} = x_1,\ldots,X_{t_n} = x_n\} =$$
$$P\{X_{t_1} = x_1\}\cdot P\{X_{t_2} = x_2 \mid X_{t_1} = x_1\}\cdots P\{X_{t_n} = x_n \mid X_{t_{n-1}} = x_{n-1}\}$$

- This is related to the so called **Markov property**:
  - Given the current state (of the process), the future (of the process) does not depend on the past (of the process), i.e. *how* the process has arrived to the current state

Stationarity

- **Definition**: Stochastic process X is **stationary** if all finite-dimensional distributions are invariant to time shifts, that is:

$$P\{X_{t_1+\Delta} \le x_1,\ldots,X_{t_n+\Delta} \le x_n\} = P\{X_{t_1} \le x_1,\ldots,X_{t_n} \le x_n\}$$

  for all Δ, n, $t_1,\ldots,t_n$ and $x_1,\ldots,x_n$
- It follows (by choosing $n = 1$) that all (individual) random variables X_t of a stationary process are identically distributed, i.e. for all $t \in I$

$$P\{X_t \le x\} = F(x)$$

  This is called the **stationary distribution** of the process.

Stochastic processes in teletraffic theory

- In this course (and, more generally, in teletraffic theory) various stochastic processes are needed to describe
  - the arrivals of customers to the system (**arrival process**)
  - the state of the system (**state process**)
- Note that the latter is also often called as **traffic process**

Arrival process

- An arrival process can be described as
  - a **point process** $(\tau_n \mid n = 1,2,\ldots)$ where τ_n tells the arrival time of the n^{th} customer (discrete-time, continuous-state)
    - non-decreasing: $\tau_{n+1} \ge \tau_n$ kaikilla n (thus non-stationary!)
    - typically it is assumed that the interarrival times $\tau_n - \tau_{n-1}$ are independent and identically distributed (IID) $\Rightarrow$ renewal process
    - then it is sufficient to specify the interarrival time distribution
    - exponential IID interarrival times $\Rightarrow$ Poisson process
  - a **counter process** $(A(t) \mid t \ge 0)$ where $A(t)$ tells the number of arrivals up to time t (continuous-time, discrete-state)
    - non-decreasing: $A(t+\Delta) \ge A(t)$ for all $t,\Delta \ge 0$ (thus non-stationary!)
    - independent and identically distributed (IID) increments $A(t+\Delta) - A(t)$ with Poisson distribution $\Rightarrow$ Poisson process

State process

- In simple cases
  - the state of the system is described just by an integer
    - e.g. the number $X(t)$ of calls or packets at time t
  - This yields a state process that is continuous-time and discrete-state
- In more complicated cases,
  - the state process is e.g. a vector of integers (cf. loss and queueing network models)
- Typically we are interested in
  - whether the state process has a stationary distribution
  - if so, what it is?
- Although the state of the system did not follow the stationary distribution at time 0, in many cases state distribution approaches the stationary distribution as t tends to ∞

Contents

- Basic concepts
- Poisson process

Bernoulli process

- **Definition**: **Bernoulli process** with success probability p is an infinite series $(X_n \mid n = 1,2,...)$ of independent and identical random experiments of Bernoulli type with joint success probability p
- Bernoulli process is clearly discrete-time and discrete-state
  - Parameter space: $I = \{1,2,\ldots\}$
  - State space: $S = \{0,1\}$
- Finite dimensional distributions (note: X_n's are IID):

$$P\{X_1 = x_1,..., X_n = x_n\} = P\{X_1 = x_1\} \cdots P\{X_n = x_n\}$$

$$= \prod_{i=1}^{n} p^{x_i} (1-p)^{1-x_i} = p^{\sum_i x_i} (1-p)^{n-\sum_i x_i}$$

- Bernoulli process is stationary (with $\mathrm{Bernoulli}(p)$ as the stationary distribution)

Definition of a Poisson process

- Poisson process is the continuous-time counterpart of a Bernoulli process
  - It is a point process $(\tau_n \mid n = 1,2,...)$ where τ_n tells tells the occurrence time of the n^{th} event, (e.g. arrival of a client)
  - "failure" in Bernoulli process is now an arrival of a client
- **Definition 1**: A point process $(\tau_n \mid n = 1,2,...)$ is a **Poisson process** with **intensity** λ if the probability that there is an event during a short time interval $(t, t+h]$ is $\lambda h + o(h)$ independently of the other time intervals
  - $o(h)$ refers to any function such that $o(h)/h \to 0$ as $h \to 0$
  - new events happen with a constant intensity λ: $(\lambda h + o(h))/h \to \lambda$
  - probability that there are no arrivals in $(t, t+h]$ is $1 - \lambda h + o(h)$
- Defined as a point process, Poisson process is discrete-time and continuous-state
  - Parameter space: $I = \{1,2,\ldots\}$
  - State space: $S = (0, \infty)$

Poisson process, another definition

- Consider the interarrival time $\tau_n - \tau_{n\text{-}1}$ between two events ($\tau_0 = 0$)
  - Since the intensity that something happens remains constant λ, the ending of the interarrival time within a short period of time $(t, t{+}h]$, after it has lasted already the time t, does not depend on t (or on other previous arrivals)
  - Thus, the interarrival times are independent and, additionally, they have the memoryless property. This property can be only the one of exponential distribution (of continuous-time distributions)
- **Definition 2**: A point process $(\tau_n \mid n = 1,2,...)$ is a **Poisson process** with **intensity** λ if the interarrival times $\tau_n - \tau_{n-1}$ are independent and identically distributed (IID) with joint distribution $\mathrm{Exp}(\lambda)$

Poisson process, yet another definition (1)

- Consider finally the number of events $A(t)$ during time interval $[0,t]$
  - In a Bernoulli process, the number of successes in a fixed interval would follow a binomial distribution. As the "time slice" tends to 0, this approaches a Poisson distribution.
  - Note that $A(0)=0$
- **Definition 3**: A counter process $(A(t) \mid t \geq 0)$ is a **Poisson process** with **intensity** λ if its increments in disjoint intervals are independent and follow a Poisson distribution as follows:

$$A(t+\Delta) - A(t) \sim \mathrm{Poisson}(\lambda\Delta)$$

- Defined as a counter process, Poisson process is continuous-time and discrete-state
  - Parameter space: $I = [0, \infty)$
  - State space: $S = \{0,1,2,\ldots\}$

Poisson process, yet another definition (2)

- One dimensional distribution: $A(t) \sim \mathrm{Poisson}(\lambda t)$
  - $E[A(t)] = \lambda t$, $D^2[A(t)] = \lambda t$
- Finite dimensional distributions (due to independence of disjoint intervals):

$$\begin{aligned} &P\{A(t_1) = x_1,\ldots, A(t_n) = x_n\} = \\ &\quad P\{A(t_1) = x_1\}P\{A(t_2) - A(t_1) = x_2 - x_1\}\cdots \\ &\quad P\{A(t_n) - A(t_{n-1}) = x_n - x_{n-1}\} \end{aligned}$$

- Poisson process, defined as a counter process is not stationary, but it has stationary increments
  - thus, it doesn't have a stationary distribution, but independent and identically distributed increments

Three ways to characterize the Poisson process

- It is possible to show that all three definitions for a Poisson process are, indeed, equivalent

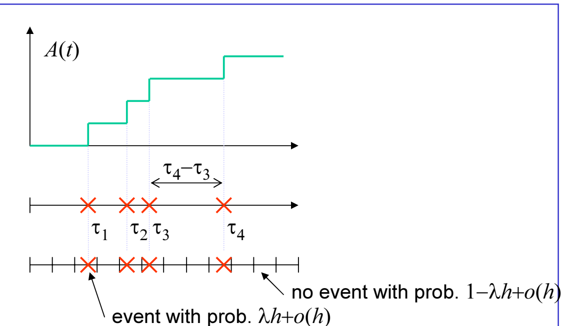

Properties (1)

- **Property 1** (**Sum**): Let $A_1(t)$ and $A_2(t)$ be two independent Poisson processes with intensities λ_1 and λ_2. Then the sum (superposition) process $A_1(t) + A_2(t)$ is a Poisson process with intensity $\lambda_1 + \lambda_2$.
- Proof: Consider a short time interval $(t, t+h]$
  - Probability that there are no events in the superposition is

$$(1-\lambda_1 h + o(h))(1-\lambda_2 h + o(h)) = 1-(\lambda_1+\lambda_2)h + o(h)$$

  - On the other hand, the probability that there is exactly one event is

$$(\lambda_1 h + o(h))(1-\lambda_2 h + o(h)) + (1-\lambda_1 h + o(h))(\lambda_2 h + o(h))$$
$$= (\lambda_1+\lambda_2)h + o(h)$$

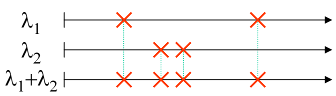

Properties (2)

- **Property 2** (**Random sampling**): Let τ_n be a Poisson process with intensity λ. Denote by σ_n the point process resulting from a random and independent sampling (with probability p) of the points of τ_n. Then σ_n is a Poisson process with intensity $p\lambda$.
- Proof: Consider a short time interval $(t, t+h]$
  - Probability that there are no events after the random sampling is

$$(1-\lambda h + o(h)) + (1-p)(\lambda h + o(h)) = 1 - p\lambda h + o(h)$$

  - On the other hand, the probability that there is exactly one event is

$$p(\lambda h + o(h)) = p\lambda h + o(h)$$

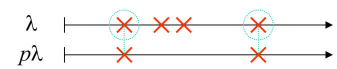

Properties (3)

- **Property 3** (**Random sorting**): Let τ_n be a Poisson process with intensity λ. Denote by $\sigma_n^{(1)}$ the point process resulting from a random and independent sampling (with probability p) of the points of τ_n. Denote by $\sigma_n^{(2)}$ the point process resulting from the remaining points. Then $\sigma_n^{(1)}$ and $\sigma_n^{(2)}$ are independent Poisson processes with intensities λp and $\lambda(1-p)$.
- Proof: Due to property 2, it is enough to prove that the resulting two processes are independent. Proof will be ignored on this course.

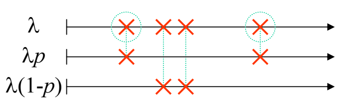

Properties (4)

- **Property 4** (**PASTA**): Consider any simple (and stable) teletraffic model with Poisson arrivals. Let $X(t)$ denote the state of system at time t (continuous-time process) and Y_n denote the state of the system seen by the nth arriving customer (discrete-time process). Then the stationary distribution of $X(t)$ is the same as the stationary distribution of Y_n.
- Thus, we can say that
  - arriving customers see the system in the stationary state
  - PASTA= "Poisson Arrivals See Time Avarages"
- PASTA property is only valid for Poisson arrivals
  - and it is not valid for other arrival processes
  - consider e.g. your own PC. Whenever you start a new session, the system is idle. In the continuous time, however, the system is not only idle but also busy (when you use it).

Example (1)

- Connection requests arrive at a server according to a Poisson process with intensity $\lambda = 5$ requests in a minute.
- What is the probability that exactly 2 new requests arrive during the next 30 seconds?
  - The number of new arrivals during this time interval follows Poisson distribution with the parameter $\lambda\Delta = (5/60)\cdot 30 = 2.5$, i.e.

$$A(t+30) - A(t) \sim \mathrm{Poisson}(2.5)$$

  - Thus

$$P\{A(t+30) - A(t) = 2\} = \frac{2.5^2}{2!} e^{-2.5} = 0.257$$

Example (2)

- Consider the system described on the previous slide.
- A new connection request has just arrived at the server. What is the probability that it takes more than 30 seconds before the next request arrives?
  - Consider the process as a point process. The interarrival time follows exponential distribution with parameter .

$$P\{\tau_{i+1} - \tau_i \geq 30\} = 1 - P\{\tau_{i+1} - \tau_i \leq 30\}$$
$$= e^{-5/60\cdot 30} = e^{-2.5} = 0.082$$

  - Consider the process as a counter process, cf. slide 29. Now we can restate the question above as ”What is the probability that there are no arrivals during 30 seconds?”

$$P\{A(t+30) - A(t) = 0\} = \frac{2.5^0}{0!} e^{-2.5} = e^{-2.5} = 0.082$$