Tutorial 1

Gidon Rosalki

2025-10-30

Notice: If you find any mistakes, please open an issue at `https://github.com/robomarvin1501/notes_networking`

1 Introduction

There will be 4 exercises, with submission in pairs. There will be a 5 point bonus if the exercise is typed, and there will be 2 interviews per pair through the semester, which will have a binary pass / fail marking system. A pass is required in order for the homework to be accepted.

The final grade is comprised of 25% from the homework, 75% from the final exam, with a two point bonus for attendance in lectures. 2 lectures can be missed to still get the bonus. In order to pass the course one needs at least 55% in the final exam, a homework average of at least 55%, and a final grade of at least 60%.

We are going to be studying the basic building blocks of modern computer networks. Computers communicate with clearly defined languages, called protocols. We will mostly be discussing established protocols, internet architecture. Towards the end we will begin discussing things likes the performance analysis of network protocols, mobile communication, and security, which has become a bigger issue since the internet was not originally founded on a basis of security.

2 Probability

Today we will be discussing probability. The concept of probability is the measure of the chance that some event will occur, for example, "a coin toss will land on heads". Most things in life (and in CS) are not deterministic, so we need to model when events could happen.
In networking, it is mostly a game of chance, a transmission over WiFi may get lost in background noise, a link can fail, so transmitted messages may not arrive. We will mostly use probability for modelling link / packet failure rates, and analysing efficiency of protocols / networks.

2.1 Definitions

Ω is the **sample space**, with elements ω_i. It describes the states in our system. It may be either *finite*, or *infinite*. An example is in a coin toss it may be $\Omega = \{H, T\}$. An **event** is a subset $A \subseteq \Omega$, so symbolises a set of possible outcomes. Events are disjoint if they have no intersection, and the complement of the event is every other event in Ω: $\overline{A} = \Omega \setminus A$.
A **probability function** is a function $\mathbb{P} : X \to [0, 1]$ that satisfies the following:

- $\forall A \subseteq \Omega \ \mathbb{P}[A] \in [0, 1]$
- $\mathbb{P}[\Omega] = 1$

For a set of disjoint events $\{A_1, \ldots, A_n\}$ it holds that:

$$\mathbb{P}\left[\bigcup_{i=1}^{n} A_i\right] = \sum_{i=1}^{n} \mathbb{P}[A_i]$$

We have **conditional probability** to denote the chance that B will occur, if we *already know* that A occurred:

$$A, B \subseteq \Omega,\ \mathbb{P}(B) > 0 \ : \ \mathbb{P}(B|A) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}$$
$$\mathbb{P}(A \cap B) = \mathbb{P}(A|B)\,\mathbb{P}(B)$$

This has some useful properties:

- A is independent of B if
$$\mathbb{P}(A|B) = \mathbb{P}(A) \implies \mathbb{P}(A \cap B) = \mathbb{P}(A) \cdot \mathbb{P}(B)$$
- $\mathbb{P}(A \cup B) \leq \mathbb{P}(A) + \mathbb{P}(B)$

- Bayes theorem:

$$\mathbb{P}(A|B) = \frac{\mathbb{P}(B|A)\,\mathbb{P}(A)}{\mathbb{P}(B)}$$

We also have complete probability: For a disjoint set

$$\Omega = \{B_i\}_{i=1}^{n}$$
$$\mathbb{P}(A) = \sum_{i=1}^{n} \mathbb{P}(A|B_i) \cdot \mathbb{P}(B_i)$$

2.1.1 Example

There are only two types of packets in some network: "Good" packets which survive in the network at least T seconds with probability e^{-T}, and "Bad" packets, which survive in the network at least T seconds with probability e^{-1000T}. The probability of creating a good packet is p.

Example 1. *Assuming a single (either good or bad) packet was created at time $T = 0$, what is the probability that it still exists at time $T = t$?*

Solution. Let us define the following events:

$$A = \text{ a packet survived for } t \text{ seconds}$$
$$B = \text{ a good packet was created}$$

Thus:

$$\mathbb{P}(A) = \mathbb{P}(A|B) \cdot \mathbb{P}(B) + \mathbb{P}\left(A|\overline{B}\right) \cdot \mathbb{P}\left(\overline{B}\right)$$
$$= e^{-t} \cdot p + e^{-1000t} \cdot (1 - p)$$

□

3 Random variables

A random variable X is a function $X : \Omega \to \mathbb{R}$. It may be said that X describes some numerical property that our sample space could have. It may either be discrete (like a coin toss), or continuous (such as time). For example, in a coin toss $\Omega = \{H, T\}$, we may define a random variable (called an indicator in this case) I_H:

$$I_H = \begin{cases} 1, & \text{if } \omega = H \\ 0, & \text{if } \omega = T \end{cases}$$

For a sample space Ω, and a random variable X, we have

$$\mathbb{P}(X = x) = \mathbb{P}(\{\omega \subseteq \Omega : X(\omega) = x\})$$

4 Expected values

The **expected value** is the "averaged" value of a series of experiments:

$$\mathbb{E}[X] = \sum_i x_i \cdot \mathbb{P}(X = x_i)$$

It has some useful properties:

$$\mathbb{E}\left[\sum_i a_i X_i\right] = \sum_i a_i \mathbb{E}[X_i]\, Var[X] = \mathbb{E}\left[\left[X - \mathbb{E}[X]^2\right]\right] = \mathbb{E}\left[X^2\right] - \mathbb{E}[X]^2$$

There are some useful random variables, which one should know, since they come up a lot:

	Bernoulli	Binomial	Geometric	Poisson
Intuition	Can either fail or succeed	n independent Bernoulli experiments, with X returning the sum	We do independent Bernoulli, until we succeed, X is the number of tries	Counting the number of events that occurred in some period of time
Probability mass function	$\mathbb{P}(X=1) = p \wedge \mathbb{P}(X=0) = 1-p$	$\mathbb{P}(X=k) = \binom{n}{k} p^k (1-p)^{n-k}$	$\mathbb{P}(X=k) = (1-p)^{k-1} p$	$\mathbb{P}(X=k) = \frac{\lambda^k}{k!} e^{-\lambda}$
Notation	$X \sim Ber(p)$	$X \sim Bin(n,p)$	$X \sim Geo(p)$	$X \sim Pois(\lambda)$
Exp	$\mathbb{E}[X] = p$	$\mathbb{E}[X] = np$	$\mathbb{E}[X] = \frac{1}{p}$	$\mathbb{E}[X] = \lambda$
Var	$Var[X] = p(1-p)$	$Var[X] = np(1-p)$	$Var[X] = \frac{1-p}{p^2}$	$Var[X] = \lambda$

Table 1:

4.1 Examples

We want to send a message from node S to node D which are connected by a chain of n links. The probability of any link to fail is p (independently). At time $T = 0$, S sends a message to D.

Example 2. *What is the probability that node D got the message?*

Solution. This is the probability that no link would fail. p is the probability of a given link failing, so $1-p$ is the probability of the link succeeding. To send a message, all n links successfully send the message, the probability of which is independent of each other, so the probability is the union of them all succeeding, ie $(1-p)^n$ □

Example 3. *Assume that if the message does not reach D, then S will send it again until it succeeds. What is the expected number of packets we need to send until node D gets the packet?*

Solution. We define X as "the number of packets to send from S until a packet reaches D successfully" – we note that $X \sim Geo((1-p)^n)$ and so:

$$\mathbb{E}[X] = \frac{1}{(1-p)^n}$$

□

Example 4. *We now add the following mechanism: Every node keeps on sending the message to its next hop until it reaches it, except the first node that sends only once. What is the probability that node D got the message?*

Solution. All the other hops will eventually succeed, so just the probability that the first succeeds $(1p)$. □

Example 5. *What is the expected number of packets that a single node will send until success?*

Solution. Let us define $X \sim Geo(1-p)$, and so $\mathbb{E}[X] = \dfrac{1}{1-p}$ □

4.2 Continuous random variables

In the case that the results of our experiment are continuous (for example, measuring time), we need continuous random variables. This is similar to the area under a curve, and the probability of any singleton is 0. We are only interested in probabilities such as $\mathbb{P}(a \le X \le b)$. It behaves like discrete random variables, however:

$$\mathbb{P}(a \le X \le b) = \int_a^b f(x)\,dx$$

$$\int_{-\infty}^{\infty} f(x)\,dx = 1$$

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} x f(x)\,dx$$

$$Var[X] = \int_{-\infty}^{\infty} (x - \mathbb{E}[X])^2\,dx$$

	Uniform	Exponential	Gaussian (Normal)
Notation	$X \sim Uni(a,b)$	$X \sim Exp(\lambda)$	$X \sim N(\mu,\sigma)$
Probability distribution function	$f(x) = \begin{cases} \dfrac{1}{b-a}, & \text{if } x \in [a,b] \\ 0, & \text{else} \end{cases}$	$f(x) = \begin{cases} \lambda e^{-\lambda x}, & \text{if } x > 0 \\ 0, & \text{else} \end{cases}$	$f(x) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$
Exp	$\mathbb{E}[X] = \frac{a+b}{2}$	$\mathbb{E}[X] = \frac{1}{\lambda}$	$\mathbb{E}[X] = \mu$
Var	$Var[X] = \frac{(b-a)^2}{12}$	$Var[X] = \frac{1}{\lambda^2}$	$Var[X] = \sigma$

Table 2:

5 Memorylessness

Some random variables have the following property:

$$\mathbb{P}(x > S + t | x > t) = \mathbb{P}(x > S)$$

Examples of this are geometric, and exponential random variables. Intuitively, it means that our placement in time does not matter. For example, the time it takes for the first customer to arrive, vs the time it takes for the next customer after the 9th customer.

5.1 Stochastic processes

The random (Stochastic) process: We would like to model a series of events in a non deterministic way, so observing the same process twice might (and should) give us different results (even if the initial starting point is the same). The process models the evolution of a system through time, where we assume that the system is represented by some random variable. In this course, we only care about discrete processes. The total number of successes after n steps in a stochastic process of i.i.d. Bernoulli trials has a binomial distribution.
In an actual network the time a packet leaves an intermediate node (switch/router) is not pre-determined. Hence, we will usually model the packets going in and out of nodes as a stochastic process. We could also model power/traffic demand and available resources to account for them.

Formally: Suppose we have random variables $X_i > 0$, that represent the time it takes for the ith packet to arrive. Let us define:

$$\begin{aligned} S_1 &= X_1 \\ S_2 &= S_1 + X_2 \\ &\vdots \\ S_i &= S_{i-1} + X_i \end{aligned}$$

We call S_i the **arrival epoch** (the specific time at which an event occurs), $N(t) = n$ for $S_n \leq t < S_{n+1}$.

A **counting process** is a random process that counts the number of arrival epochs until time t. A counting process is defined by $\{N_t | t \geq 0\}$, where we have the following:

- $N_0 = 0$
- $N_t \geq 0$
- $s \leq t \implies N_s \leq N_t$ (monotonicity), and if $s < t$, then $N_t - N_s$ is the number of events that occurred in $(s, t]$

Consider as an example, the number of customers that arrive in a store.

A **Poisson process** $\{N_t | t \geq 0\}$, with a rate of λ is a counting process, where in addition:

- Independent increments: Number of events in any two disjoint intervals is independent
- Stationary increments: The number of events in any interval of length t is a Poisson random variable with parameter λt (depends on the interval's length and not on its timing). ie. $\mathbb{E}[N_t] = \lambda t$
- No bunching: The probability of > 1 arrivals in a tiny interval is negligible

Given this, is the number of people boarding the bus a Poisson process?
No. The passengers board in groups (at bus stops), and during rush hour more people board.

The inter-arrival times between events in a Poisson counting process are I.I.D and have an exponential distribution with parameter λ.

Example 6 (Poisson process)**.** *In each interval of length T, the number of transmitted packets is a Poisson random variable with parameter gT. What is the probability to have only a single packet transmitted?*

Solution. Recall $X_o \sim Poi(gt)$, and $\mathbb{P}_{t=T}(X = i) = \frac{(gT)^i}{i!} e^{-gT}$. Let us define a random variable X that represents the number of packets transmitted at some interval. Then, $X \sim Poi(gT)$ and we get:

$$\mathbb{P}_{t=T}(X = 1) = gT \cdot e^{-gT}$$

□