Computer Networks
Modeling arrivals and service with Poisson

Saad Mneimneh
Computer Science
Hunter College of CUNY
New York

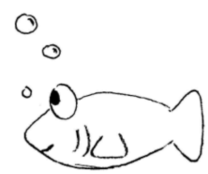

Now remember this:
A poisson has a short memory span of few months, but a Poisson process is memoryless.

1 Introduction

In computer networks, packet arrivals and service are modeled as a stochastic process in which events occur at times $t_1, t_2, \ldots$ For instance, in the figure below, $t_1, t_2, \ldots$ can be interpreted as the packet arrival times, or the service completion times. Accordingly, T_i, defined as $t_i - t_{i-1}$ for $i > 1$, can be interpreted as the inter-arrival times of packets (intervals between subsequent arrivals), or the delays experienced by the served packets (assuming that the server is always busy). Similarly, $A(t)$ denotes the number of packets that arrive in $[0, t]$, or the number of packets served in $[0, t]$. In the rest of this document, we will refer to packets rather than service, but it should be clear that the discussion applies to both.

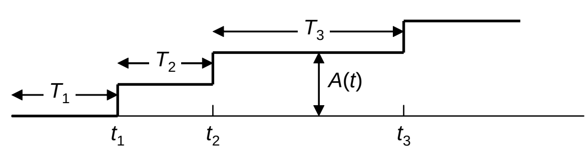

Figure 1: Arrival/service process

2 Poisson

The packet inter-arrival times are typically modeled as a Poisson process, i.e. T_i are independent and identically distributed (IID, so we drop i from the term T_i in the following expression), and obey an an exponential distribution:

$$F_T(t) = P(T \leq t) = 1 - e^{-\lambda t}$$

where λ is a parameter. We will give λ a name shortly.

The probability density function for T is therefore (the derivative of $1 - e^{-\lambda t}$ with respect to t):

$$f_T(t) = \lambda e^{-\lambda t}$$

Therefore,

$$P(t_1 \leq T \leq t_2) = \int_{t_1}^{t_2} \lambda e^{-\lambda t} = -e^{-\lambda t}|_{t_1}^{t_2} = -e^{-\lambda t_2} + e^{-\lambda t_1}$$

The expected value of T can be obtained as:

$$E[T] = \int_0^t t\lambda e^{-\lambda t} dt = \frac{1}{\lambda}$$

Therefore, the parameter λ is called the arrival rate, or simply rate [1]. Similarly,

$$E[T^2] = \int_0^t t^2 \lambda e^{-\lambda t} dt = \frac{2}{\lambda^2}$$

Therefore, the variance is:

$$\sigma^2(T) = E[(T - E[T])^2] = E[T^2 + E[T]^2 - 2E[T]T]$$

By the linearity of expectation,

$$\sigma^2(T) = E[T^2] + E[T]^2 - 2E[T]E[T] = E[T^2] - E[T]^2 = \frac{1}{\lambda^2}$$

3 Poisson process is memoryless

Now we prove a unique property of the exponential process, known as the memoryless property. Consider the waiting time until some arrival occurs. The memoryless property states that given that no arrival has occurred by time τ, the distribution of the remaining waiting time is the same as it was originally. Mathematically,

$$P(T > \tau + t | T > \tau) = P(T > t)$$

The proof is simple as a direct consequence of the exponential distribution:

[1] A mathematical result, known as the law of large numbers, says that if X_i are IID, then $\lim_{n\to\infty} \frac{\sum_{i=1}^n X_i}{n} = E[X]$. Therefore, the inverse of the rate can be expressed as $\lim_{n\to\infty} \frac{t_n}{A(t_n)} = \lim_{n\to\infty} \frac{\sum_{i=1}^n T_i}{n} = E[T] = \frac{1}{\lambda}$.

$$P(T > \tau + t | T > \tau) = \frac{P(T > \tau + t \text{ and } T > \tau)}{P(T > \tau)} = \frac{P(T > \tau + t)}{P(T > \tau)} = \frac{e^{-\lambda(\tau+t)}}{e^{-\lambda\tau}} = e^{-\lambda t}$$

As a direct application of the memoryless property, consider Z_1 to be the waiting time until the first arrival occurs after some time t. Assume also that $A(t) = n$, and $t_n = \tau \leq t$. Then

$$P(Z_1 > z | A(t) = n, t_n = \tau) = P(T_{n+1} > z + t - \tau | T_{n+1} > t - \tau, t_n = \tau)$$

Note that T_{n+1} is independent of all earlier inter-arrival times, and thus of t_n. Therefore,

$$P(Z_1 > z | A(t) = n, t_n = \tau) = P(T_{n+1} > z + t - \tau | T_{n+1} > t - \tau)$$

Using the memoryless property:

$$P(Z_1 > z | A(t) = n, t_n = \tau) = P(T_{n+1} > z) = e^{-\lambda z}$$

Since the result does not depend on τ, we see that conditional on $A(t)$, Z_1 is independent of t_n.

$$P(Z_1 > z | A(t) = n) = e^{-\lambda z}$$

Since the result also does not depend on n, Z_1 is independent of $A(t)$.

$$P(Z_1 > z) = e^{-\lambda z}$$

The same argument applies if we conditioned not only on t_n, but also on $t_1, t_2, \ldots t_{n-1}$. Therefore,

$$P(Z_1 > z | \{A(\tau), 0 \leq \tau \leq t\}) = e^{-\lambda z}$$

If we now define Z_i for $i > 1$ to be the inter-arrival time from the $i - 1^{st}$ arrival to the i^{th} arrival after time t, we see that conditional on $A(t) = n$, $Z_i = T_{i+n}$ and, therefore, $Z_1, Z_2, \ldots$ are IID exponentially distributed conditional on $A(t) = n$. Since the distribution is independent of n, $Z_1, Z_2, \ldots$ are IID exponentially distributed and independent of $A(t)$. It should be also clear that $Z_1, Z_2, \ldots$ are independent of $\{A(\tau), 0 \leq \tau \leq t\}$.

The memoryless property shows that the portion of a Poisson process starting at some time $t > 0$, is a probabilistic replica of the process starting at $t = 0$. Therefore, $\tilde{A}(t, t+\delta) = A(t+\delta) - A(t)$ has the same distribution as $A(\delta)$. This property is called the *stationary increment property*. We have shown also that $\tilde{A}(t, t+\delta)$ is independent of all arrivals in $[0, t]$. Therefore, for $t_1 < t_2 < \ldots < t_k$, $A(t_1), \tilde{A}(t_1, t_2), \ldots, \tilde{A}(t_{k-1}, t_k)$ are independent. This property is called the *independent increment property*.

4 A paradox

In this section we consider a hypothetic example of a Poisson process applied to bus arrivals/departures. Suppose a bus arrives at (then departs) a station according to a Poisson process with average inter-arrival time of 20 minutes (i.e. $\lambda = 0.05$ bus/minute). When a customer arrives to the station at time t, what is the average waiting time until the next bus? According to the previous section, the waiting time is Z_1, which is exponentially distributed (with parameter λ) and independent of everything prior to t. Therefore, $E[Z_1] = 1/\lambda = 20$ minutes. Now, when a customer arrives to the station at time t, what is the average time since the last departure? Let us call this time Y_1. We don't know how many arrivals/departures we had in $[0, t]$, so let us condition on $A(t)$.

$$P(Y_1 > y | A(t) = n) = P(T_{n+1} > y) = e^{-\lambda y}$$

Since this is independent of n, Y_1 is independent of $A(t)$ and is exponentially distributed (with parameter λ). Therefore, $E[Y_1] = 1/\lambda = 20$ minutes.

By definition, the inter-arrival time is the time from the last departure to the next arrival. Therefore, the average inter-arrival time is $E[Y_1 + Z_1] = E[Y_1] + E[Z_1] = 20 + 20 = 40$ minutes, by the linearity of expectation. But it should be 20 minutes only! What happened? This is an example of what is known as *random incidence*. The customer is more likely to fall in a large interval. For an intuition, consider the following example of randomly throwing a ball into bins.

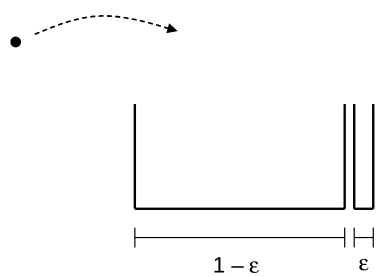

Figure 2: Random incidence

The average bin size is:

$$\frac{(1-\epsilon)+\epsilon}{2} = \frac{1}{2}$$

The ball falling at random will observe a bin of size $(1-\epsilon)$ with probability $(1-\epsilon)$, and a bin of size ϵ with probability ϵ. Therefore, the average bin size observed by the ball is:

$$(1-\epsilon)(1-\epsilon) + \epsilon \cdot \epsilon = (1-\epsilon)^2 + \epsilon^2 \approx 1 - 2\epsilon > \frac{1}{2}$$

For a mathematical interpretation of the paradox, let $Z(t)$ be the time until the first arrival after t.

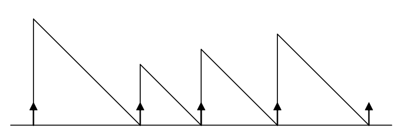

Figure 3: Time until next arrival

Similarly, let $Y(t)$ be the time since the last departure before t.

Now, $Y(t) + Z(t)$ is shown below:

The time average of $Y(t) + Z(t)$ can be computed as:

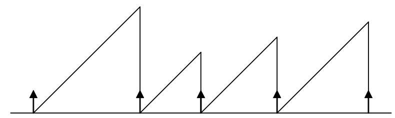

Figure 4: Time since last departure

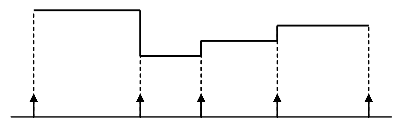

Figure 5: Y(t)+Z(t)

$$\lim_{t\to\infty}\frac{1}{t}\int_0^t (Y(t)+Z(t))dt = \frac{\text{sum of areas of squares}}{\text{length}} = \lim_{n\to\infty}\frac{\sum_{i=1}^n T_i^2}{\sum_{i=1}^n T_i} = \lim_{n\to\infty}\frac{\sum_{i=1}^n T_i^2/n}{\sum_{i=1}^n T_i/n}$$

But $\lim_{n\to\infty}\sum_{i=1}^n T_i^2/n = E[T^2]$ and $\lim_{n\to\infty}\sum_{i=1}^n T_i/n = E[T]$. Therefore,

$$\lim_{t\to\infty}\frac{1}{t}\int_0^t (Y(t)+Z(t))dt = \frac{2/\lambda^2}{1/\lambda} = \frac{2}{\lambda}$$

5 Poisson as a counting process

Another characterization of a Poisson process can be obtained by starting to observe the relation between t and T:

$$t_n = \sum_{i=1}^n T_i = T_1 + T_2 \dots T_n$$

From this relation, we can show that the probability density function for the arrival times is:

$$f_{t_n}(t) = \frac{\lambda^n t^{n-1} e^{-\lambda t}}{(n-1)!}$$

To describe a Poisson process as a counting process (that counts the number of arrivals), we are interested in obtaining the probability mass function for $A(t)$, i.e. $P(A(t) = n)$ for some integer n. Note that

$$\begin{aligned}
P(t_{n+1} \in (t, t+\delta]) &= P(A(t) = n \text{ and } Z_1 \le \delta) \\
&= P(A(t) = n) \cdot P(Z_1 \le \delta | A(t) = n) \\
&= P(A(t) = n) \cdot P(Z_1 \le \delta) \text{ (memoryless property)} \\
&= P(A(t) = n) \cdot (1 - e^{-\lambda\delta})
\end{aligned}$$

Therefore,

$$P(A(t) = n) = \frac{P(t_{n+1} \in (t, t+\delta])}{1 - e^{-\lambda\delta}}$$

Now,

$$P(t_{n+1} \in (t, t+\delta]) = \int_t^{t+\delta} \frac{\lambda^{n+1}\tau^n e^{-\lambda\tau}}{n!} d\tau \approx \frac{\lambda^{n+1}t^n e^{-\lambda t}}{n!}\delta$$

as it can be seen from the figure below:

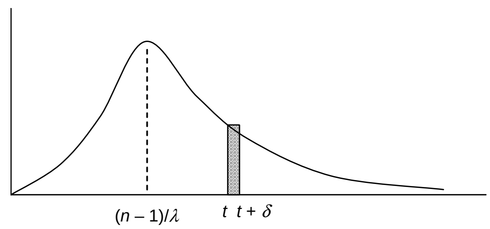

Figure 6: Area under $f_{t_n}(t)$ around $[t, t+\delta]$

Therefore,

$$P(A(t) = n) = \frac{\lambda^{n+1}t^n e^{-\lambda t}}{n!(1 - e^{-\lambda\delta})/\delta}$$

Note that $e^{-\lambda\delta} = 1 - \lambda\delta + (\lambda\delta)^2/2! - (\lambda\delta)^3/3! + \ldots$. Therefore, $(1 - e^{-\lambda\delta})/\delta = \lambda + o(\delta)/\delta$, where $\lim_{\delta\to 0} o(\delta)/\delta = 0$. We can now take the limit as $\delta \to 0$ (which justifies the integral approximation done above) to obtain:

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

It is not hard to show that any counting process that satisfies the above probability mass function, in addition to the stationary and independent increment properties, is a Poisson process. For instance $P(T_1 > t) = P(A(t) = 0) = e^{-\lambda t}$. Also, $P(T_n > t|t_{n-1} = \tau) = \lim_{\delta\to 0} P(\tilde{A}(t, t+\tau) = 0|\tilde{A}(\tau - \delta, \tau) = 1, \tilde{A}(0, \tau - \delta) = n - 2) = P(A(t) = 0)$ (stationary and independent increment).

From this probability mass function, we can obtain the following:

$$E[A(t)] = \lambda t$$
$$E[A^2(t)] = \lambda t + (\lambda t)^2$$
$$\sigma^2(A(t)) = \lambda t$$

6 PASTA

Let $N(t)$ be the number of customers in the system at time t. Define

$$p_n(t) = P(N(t) = n)$$

to be the probability of having n customers in the system at time t. If the system has a steady state distribution, then we can also define

$$p_n = \lim_{t\to\infty} p_n(t)$$

This of course means that $\bar{N} = \sum_{i=0}^{\infty} np_n = \lim_{t\to\infty} \frac{1}{t}\int_0^t N(\tau)d\tau$. Let us now consider the same probability as seen by an arriving customer:

$$a_n(t) = P(N(t) = n|\text{a customer arrives just after time } t)$$

and

$$a_n = \lim_{t\to\infty} a_n(t)$$

Is $a_n = p_n$, $n = 0, 1, \ldots$? We provide two examples to show that this is not necessary.

6.1 Example 1

Assume customer inter-arrival times are unifromly distributed between 2 and 4 seconds. Assume also that customer service times are all equal to 1. Then an arriving customer always finds an empty system. Therefore, $a_0 = 1$. On the other hand, the average number of customers in the system is $N = 1/3$ (apply Little's theorem on $\lambda = 1/3$ and $T = 1$). Therefore, $p_0 \neq 1$.

6.2 Example 2

Assume customer inter-arrival times are exponentially distributed (Poisson process). Assume also that customer service times and future arrival times are correlated, e.g. the service of the n^{th} customer is half the inter-arrival time between customers n and $n+1$. Clearly, an arriving customer always finds an empty system. Therefore, $a_0 = 1$. However, the average number of customers in the system is easily seen to be $1/2$. Therefore, $p_0 \neq 1$.

It turns out that this cannot happen with a Poisson process where service times and future arrivals are independent.

$$\begin{aligned} a_n(t) &= \lim_{\delta\to 0} P(N(t) = n|\tilde{A}(t, t+\delta) \\ &= \lim_{\delta\to 0} P(A(t) - D(t) = n|\tilde{A}(t, t+\delta) = 1) \end{aligned}$$

where $D(t)$ is the number of customers that depart in $[0, t]$. Since $\tilde{A}(t, t+\delta)$ and $A(t)$ are independent, and service times are independent of future arrivals, $A(t) - D(t)$ and $\tilde{A}(t, t+\delta)$ are independent. Therefore, $\lim_{\delta\to 0} P(A(t) - D(t) = n|\tilde{A}(t, t+\delta) = 1) = P(A(t) - D(t) = n) = P(N(t) = n) = p_n(t)$. Taking the limit as $t \to \infty$, we have $a_n = p_n$. This property of a Poisson process is called PASTA (Poisson Arrivals See Time Averages).

7 Merging and splitting Poisson processes

Consider the number of arrivals of a Poisson process in the interval $[t, t+\delta]$.

$$P(\tilde{A}(t, t+\delta) = 0) = e^{-\lambda\delta} = 1 - \lambda\delta + o(\delta)$$
$$P(\tilde{A}(t, t+\delta) = 1) = \lambda\delta e^{-\lambda\delta} = \lambda\delta + o(\delta)$$
$$P(\tilde{A}(t, t+\delta) \geq 2) = 1 - (1 + \lambda\delta)e^{-\lambda\delta} = o(\delta)$$

where $\lim_{\delta\to 0} o(\delta)/\delta = 0$.

It is not hard to show that any process that satisfies the above probabilities, in addition to the stationary and independent increment properties, is a Poisson process.

For instance $P(A(t+\delta)=0) = P(A(t)=0, \tilde{A}(t,t+\delta)=0) = P(A(t)=0)P(\tilde{A}(t,t+\delta)=0) = P(A(t)=0)(1-\lambda\delta+o(\delta))$. Therefore,

$$\lim_{\delta\to\infty} \frac{P(A(t+\delta)=0) - P(A(t)=0)}{\delta} = -\lambda P(A(t)=0)$$

$$\frac{d}{dt}P(A(t)=0)) = -\lambda P(A(t)=0)$$

So $P(A(t)=0) = e^{-\lambda t}$. In general, one can show that $P(A(t)=n) = (\lambda t)^n e^{-\lambda t}/n!$. This alternative definition of a Poisson process allows us to think about merging and splitting Poisson processes.

7.1 Merging

We can show that the sum of two Poisson processes with rates λ and μ is a Poisson process with rate $\lambda+\mu$.

$$P(\tilde{A}(t,t+\delta)=0) = (1-\lambda\delta+o(\delta))(1-\mu\delta+o(\delta)) = 1-(\lambda+\mu)\delta+o(\delta)$$

$$P(\tilde{A}(t,t+\delta)=1) = (\lambda\delta+o(\delta))(1-\mu\delta+o(\delta))+(\mu\delta+o(\delta))(1-\lambda\delta+o(\delta)) = (\lambda+\mu)\delta+o(\delta)$$

$$P(\tilde{A}(t,t+\delta)\geq 2) = o(\delta)$$

Since the two processes satisfy the stationary and independent increment properties, the resulting process does too. Therefore, the resulting process is a Poisson process with rate $\lambda+\mu$.

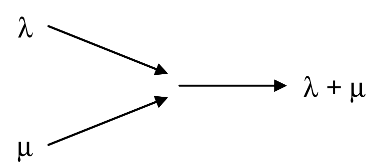

Figure 7: Merging two Poisson processes

7.2 Splitting

A Poisson process with rate λ can be split into two independent Poisson processes as follows: each arrival is independently directed to process 1 with probability p and to process 2 with probability $1-p$.

$$P(\tilde{A}_1(t,t+\delta)=0) = (1-\lambda\delta)+(1-p)\lambda\delta+o(\delta) = 1-p\lambda\delta+o(\delta)$$

$$P(\tilde{A}_1(t,t+\delta)=1) = p(\lambda\delta+o(\delta))+o(\delta) = p\lambda\delta+o(\delta)$$

$$P(\tilde{A}_1(t,t+\delta)\geq 2) = o(\delta)$$

Similar calculation can be done for process 2 by exchanging p and $1-p$. Since the original process satisfies the stationary and independent increments properties, so do process 1 and process 2. Therefore, process 1 and process 2 are both Poisson processes with rates $p\lambda$ and $(1-p)\lambda$ respectively.

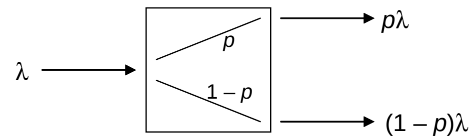

Figure 8: Splitting a Poisson process

It remains to show that process 1 and process 2 are independent. Note that conditional on the original process, process 1 and process 2 are not independent. In fact, one determines the other.

$$P(A_1(t) = m, A_2(t) = n | A(t) = m + n) = \binom{m+n}{m} p^m (1-p)^n$$

where $\binom{m+n}{m} = \frac{(m+n)!}{m!n!}$. This is simply the binomial distribution, since, given $m + n$ arrivals to the original process, each independently goes to process 1 with probability p.

$$P(A_1(t) = m, A_2(t) = n | A(t) = m + n) = \frac{P(A_1(t) = m, A_2(t) = n)}{P(A(t) = m + n)}$$

Combining the two equalities, we get:

$$P(A_1(t) = m, A_2(t) = n) = \frac{(p\lambda t)^m e^{-\lambda p t}}{m!} \frac{[(1-p)\lambda t]^n e^{-\lambda(1-p)t}}{n!}$$
$$P(A_1(t) = m, A_2(t) = n) = P(A_1(t) = m) P(A_2(t) = n)$$

which proves that $A_1(t)$ and $A_2(t)$ are independent. To show that process 1 and process 2 are independent we must show that for any $t_1 < t_2 < \ldots < t_k$, $\{\tilde{A}_1(t_{i-1}, t_i), 1 \leq i \leq k\}$ and $\{\tilde{A}_2(t_{j-1}, t_j), 1 \leq j \leq k\}$ are independent. The argument above shows this independence for $i = j$. For $i \neq j$, the independence follows from the independent increment property of $A(t)$.

References

Dimitri Bertsekas and Robert Gallager, Data Networks
Rober Gallager, Discrete Stochastic Processes