

Chapter 2

POISSON PROCESSES

2.1 Introduction

A Poisson process is a simple and widely used stochastic process for modeling the times at which arrivals enter a system. It is in many ways the continuous-time version of the Bernoulli process. Section 1.4.1 characterized the Bernoulli process by a sequence of IID binary random variables (rv's), Y_1, Y_2, \dots where $Y_i = 1$ indicates an arrival at increment i and $Y_i = 0$ otherwise. We observed (without any careful proof) that the process could also be characterized by a sequence of geometrically distributed interarrival times.

For the Poisson process, arrivals may occur at arbitrary positive times, and the probability of an arrival at any particular instant is 0. This means that there is no very clean way of describing a Poisson process in terms of the probability of an arrival at any given instant. It is more convenient to define a Poisson process in terms of the sequence of interarrival times, X_1, X_2, \dots , which are defined to be IID. Before doing this, we describe arrival processes in a little more generality.

2.1.1 Arrival processes

An *arrival process* is a sequence of increasing rv's, $0 < S_1 < S_2 < \dots$, where $S_i < S_{i+1}$ means that $S_{i+1} - S_i$ is a positive rv, *i.e.*, a rv X such that $F_X(0) = 0$. The rv's S_1, S_2, \dots are called arrival epochs¹ and represent the successive times at which some random repeating phenomenon occurs. Note that the process starts at time 0 and that multiple arrivals can't occur simultaneously (the phenomenon of bulk arrivals can be handled by the simple extension of associating a positive integer rv to each arrival). We will sometimes permit simultaneous arrivals or arrivals at time 0 as events of zero probability, but these can be ignored. In order to fully specify the process by the sequence S_1, S_2, \dots of rv's, it is necessary to specify the joint distribution of the subsequences S_1, \dots, S_n for all $n > 1$.

Although we refer to these processes as arrival processes, they could equally well model departures from a system, or any other sequence of incidents. Although it is quite common,

¹Epoch is a synonym for time. We use it here to avoid overusing the word time in overlapping ways.

especially in the simulation field, to refer to incidents or arrivals as events, we shall avoid that here. The n th arrival epoch S_n is a rv and $\{S_n \leq t\}$, for example, is an event. This would make it confusing to refer to the n th arrival itself as an event.

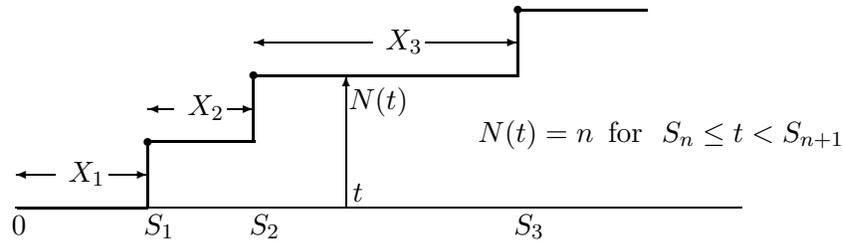


Figure 2.1: A sample function of an arrival process and its arrival epochs $\{S_1, S_2, \dots\}$, its interarrival times $\{X_1, X_2, \dots\}$, and its counting process $\{N(t); t > 0\}$

As illustrated in Figure 2.1, any arrival process $\{S_n; n \geq 1\}$ can also be specified by either of two alternative stochastic processes. The first alternative is the sequence of interarrival times, $\{X_i; i \geq 1\}$. The X_i here are positive rv's defined in terms of the arrival epochs by $X_1 = S_1$ and $X_i = S_i - S_{i-1}$ for $i > 1$. Similarly, each arrival epoch S_n is specified by $\{X_i; i \geq 1\}$ as

$$S_n = \sum_{i=1}^n X_i. \quad (2.1)$$

Thus the joint distribution of X_1, \dots, X_n for all $n > 1$ is sufficient (in principle) to specify the arrival process. Since the interarrival times are IID in most cases of interest, it is usually simpler to specify the joint distribution of the X_i than of the S_i .

The second alternative is the counting process $\{N(t); t > 0\}$, where for each $t > 0$, the rv $N(t)$ is the aggregate number of arrivals² up to and including time t .

The counting process $\{N(t); t > 0\}$, illustrated in Figure 2.1, is an uncountably infinite family of rv's where the *counting rv* $N(t)$, for each $t > 0$, is the aggregate number of arrivals in the interval $(0, t]$, *i.e.*, the number of arrivals in the set $\{\tau : 0 < \tau \leq t\}$. By convention, a parenthesis is used to indicate that the corresponding end point of an interval is not contained in the interval and a bracket is used if it is. Thus, as another example, $(a, b) = \{\tau : a < \tau < b\}$. The rv $N(0)$ is defined to be 0 with probability 1, which means that we are considering only arrivals at strictly positive times.

The counting process $\{N(t); t > 0\}$ for any arrival process has the property that $N(\tau) \geq N(t)$ for all $\tau \geq t > 0$ (*i.e.*, $N(\tau) - N(t)$ is a nonnegative random variable for $\tau \geq t > 0$).

For any given integer $n \geq 1$ and time $t > 0$, the n th arrival epoch, S_n , and the counting random variable, $N(t)$, are related by

$$\{S_n \leq t\} = \{N(t) \geq n\}. \quad (2.2)$$

²Thus, for the Bernoulli process with an increment size of 1, $N(n)$ is the rv denoted as S_n in Section 1.4.1

To see this, note that $\{S_n \leq t\}$ is the event that the n th arrival occurs at some epoch $\tau \leq t$. This event implies that $N(\tau) = n$, and thus that $\{N(t) \geq n\}$. Similarly, $\{N(t) = m\}$ for some $m \geq n$ implies $\{S_m \leq t\}$, and thus that $\{S_n \leq t\}$. This equation is essentially obvious from Figure 2.1, but is one of those peculiar obvious things that is often difficult to see. An alternate form, which is occasionally more transparent, comes from taking the complement of both sides of (2.2), getting

$$\{S_n > t\} = \{N(t) < n\}. \quad (2.3)$$

For example, the event $\{S_1 > t\}$ means that the first arrival occurs after t , which means $\{N(t) < 1\}$ (i.e., $\{N(t) = 0\}$). These relations will be used constantly in going back and forth between arrival epochs and counting rv's. In principle, (2.2) or (2.3) can be used to specify joint CDF of arrival epochs in terms of joint CDF's of counting variables and vice versa, so either characterization can be used to specify an arrival process.

In summary, then, an arrival process can be specified by the joint distributions of the arrival epochs, or of the interarrival times, or of the counting rv's. In principle, specifying any one of these specifies the others also.³

2.2 Definition and properties of a Poisson process

The nicest, and in many ways most important, arrival processes are those for which the interarrival times are IID. These arrival processes are called *renewal processes* and form the subject of Chapter 5.

Poisson processes are the nicest, and in many ways most important, class of renewal processes. They are characterized by having exponentially distributed interarrival times. We will soon see why these exponentially distributed interarrival times simplify these processes and make it possible to look at them in so many simple ways. We start by being explicit about the definitions of renewal and Poisson processes.

Definition 2.2.1. *A renewal process is an arrival process for which the sequence of interarrival times is a sequence of positive IID rv's.*

Definition 2.2.2. *A Poisson process is a renewal process in which the interarrival times have an exponential CDF; i.e., for some real $\lambda > 0$, each X_i has the density⁴ $f_X(x) = \lambda \exp(-\lambda x)$ for $x \geq 0$.*

The parameter λ is called the rate of the process. We shall see later that for any interval of size t , λt is the expected number of arrivals in that interval, giving some interpretation to denoting λ as a rate.

³By definition, a stochastic process is a collection of rv's, so one might ask whether an arrival process (as a stochastic process) is 'really' the arrival-epoch process $\{S_n; n \geq 1$ or the interarrival process $X_i; i > 0\}$ or the counting process $\{N(t); t > 0\}$. The arrival-epoch process comes to grips with the actual arrivals, the interarrival process is often the simplest, and the counting process 'looks' most like a stochastic process in time since $N(t)$ is a rv for each $t \geq 0$. It seems preferable, since the descriptions are so clearly equivalent, to view arrival processes in terms of whichever description is most convenient.

⁴With this density, $\Pr\{X_i > 0\} = 1$, so that we can regard X_i as a positive random variable. Since events of probability zero can be ignored, the density $\lambda \exp(-\lambda x)$ for $x \geq 0$ and zero for $x < 0$ is effectively the same as the density $\lambda \exp(-\lambda x)$ for $x > 0$ and zero for $x \leq 0$.

2.2.1 Memoryless property

What makes the Poisson process unique among renewal processes is the memoryless property of the exponential distribution.

Definition 2.2.3. Memoryless random variables: *A rv X possesses the memoryless property if X is a positive rv (i.e., $\Pr\{X > 0\} = 1$) for which*

$$\Pr\{X > t + x\} = \Pr\{X > x\} \Pr\{X > t\} \quad \text{for all } x, t \geq 0. \quad (2.4)$$

Note that (2.4) is a statement about the complementary CDF of X . There is no intimation that the event $\{X > t + x\}$ in the equation has any particular relation to the events $\{X > t\}$ or $\{X > x\}$.

For an exponential rv X of rate $\lambda > 0$, $\Pr\{X > x\} = e^{-\lambda x}$ for $x \geq 0$. This satisfies (2.4) for all $x \geq 0$, $t \geq 0$, so X is memoryless. Conversely, an arbitrary rv X is memoryless only if it is exponential. To see this, let $h(x) = \ln[\Pr\{X > x\}]$ and observe that since $\Pr\{X > x\}$ is non-increasing in x , $h(x)$ is also. In addition, (2.4) says that $h(t + x) = h(x) + h(t)$ for all $x, t \geq 0$. These two statements (see Exercise 2.6) imply that $h(x)$ must be linear in x , and $\Pr\{X > x\}$ must be exponential in x .

Since a memoryless rv X must be exponential, there must be some $\lambda > 0$ such that $\Pr\{X > t\} = e^{-\lambda t} > 0$ for all $t \geq 0$. Since $\Pr\{X > t\} > 0$, we can rewrite (2.4) as

$$\Pr\{X > t + x \mid X > t\} = \Pr\{X > x\}. \quad (2.5)$$

If X is interpreted as the waiting time until some given arrival, then (2.5) states that, given that the arrival has not occurred by time t , the distribution of the remaining waiting time (given by x on the left side of (2.5)) is the same as the original waiting time distribution (given on the right side of (2.5)), i.e., the remaining waiting time has no ‘memory’ of previous waiting.

Example 2.2.1. Suppose X is the waiting time, starting at time 0, for a bus to arrive, and suppose X is memoryless. After waiting from 0 to t , the distribution of the remaining waiting time from t is the same as the original distribution starting from 0. The still waiting customer is, in a sense, no better off at time t than at time 0. On the other hand, if the bus is known to arrive regularly every 16 minutes, then it will certainly arrive within a minute for a person who has already waited 15 minutes. Thus regular arrivals are not memoryless. The opposite situation is also possible. If the bus frequently breaks down, then a 15 minute wait can indicate that the remaining wait is probably very long, so again X is not memoryless. We study these non-memoryless situations when we study renewal processes in Chapter 5.

Although memoryless distributions must be exponential, it can be seen that if the definition of memoryless is restricted to integer times, then the geometric distribution becomes memoryless, and it can be seen as before that this is the only memoryless integer-time distribution. In this respect, the Bernoulli process (which has geometric interarrival times)

is like a discrete-time version of the Poisson process (which has exponential interarrival times).

We now use the memoryless property of exponential rv's to find the distribution of the first arrival in a Poisson process after an arbitrary given time $t > 0$. We not only find this distribution, but also show that this first arrival after t is independent of all arrivals up to and including t . More precisely, we prove the following theorem.

Theorem 2.2.1. *For a Poisson process of rate λ , and any given $t > 0$, the length of the interval from t until the first arrival after t is a nonnegative rv Z with the CDF $1 - \exp[-\lambda z]$ for $z \geq 0$. This rv is independent of both $N(t)$ and of the $N(t)$ arrival epochs before time t . It is also independent of the set of rv's $\{N(\tau); \tau \leq t\}$.*

The basic idea behind this theorem is that Z , conditional on the epoch τ of the last arrival before t , is simply the remaining time until the next arrival. Since the interarrival time starting at τ is exponential and thus memoryless, Z is independent of $\tau \leq t$, and of all earlier arrivals. The following proof carries this idea out in detail.

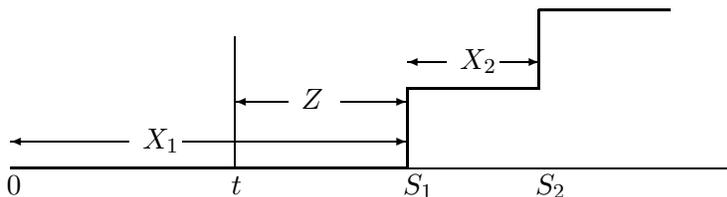


Figure 2.2: For arbitrary fixed $t > 0$, consider the event $\{N(t) = 0\}$. Conditional on this event, Z is the distance from t to S_1 ; i.e., $Z = X_1 - t$.

Proof: Let Z be the interval from t until the first arrival after t . We first condition on $N(t) = 0$ (see Figure 2.2). Given $N(t) = 0$, we see that $X_1 > t$ and $Z = X_1 - t$. Thus,

$$\begin{aligned} \Pr\{Z > z \mid N(t)=0\} &= \Pr\{X_1 > z + t \mid N(t)=0\} \\ &= \Pr\{X_1 > z + t \mid X_1 > t\} \end{aligned} \quad (2.6)$$

$$= \Pr\{X_1 > z\} = e^{-\lambda z}. \quad (2.7)$$

In (2.6), we used the fact that $\{N(t) = 0\} = \{X_1 > t\}$, which is clear from Figure 2.1 (and also from (2.3)). In (2.7) we used the memoryless condition in (2.5) and the fact that X_1 is exponential.

Next consider the conditions that $N(t) = n$ (for arbitrary $n > 1$) and $S_n = \tau$ (for arbitrary $\tau \leq t$). The argument here is basically the same as that with $N(t) = 0$, with a few extra details (see Figure 2.3).

Conditional on $N(t) = n$ and $S_n = \tau$, the first arrival after t is the first arrival after the arrival at S_n , i.e., $Z = z$ corresponds to $X_{n+1} = z + (t - \tau)$.

$$\Pr\{Z > z \mid N(t)=n, S_n=\tau\} = \Pr\{X_{n+1} > z+t-\tau \mid N(t)=n, S_n=\tau\} \quad (2.8)$$

$$= \Pr\{X_{n+1} > z+t-\tau \mid X_{n+1} > t-\tau, S_n=\tau\} \quad (2.9)$$

$$= \Pr\{X_{n+1} > z+t-\tau \mid X_{n+1} > t-\tau\} \quad (2.10)$$

$$= \Pr\{X_{n+1} > z\} = e^{-\lambda z}, \quad (2.11)$$

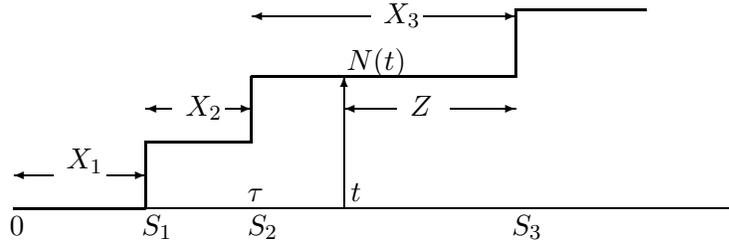


Figure 2.3: Given $N(t) = 2$, and $S_2 = \tau$, X_3 is equal to $Z + (t - \tau)$. Also, the event $\{N(t)=2, S_2=\tau\}$ is the same as the event $\{S_2=\tau, X_3>t-\tau\}$.

where (2.9) follows because, given $S_n = \tau \leq t$, we have $\{N(t) = n\} = \{X_{n+1} > t - \tau\}$ (see Figure 2.3). Eq. (2.10) follows because X_{n+1} is independent of S_n . Eq. (2.11) follows from the memoryless condition in (2.5) and the fact that X_{n+1} is exponential.

The same argument applies if, in (2.8), we condition not only on S_n but also on S_1, \dots, S_{n-1} . Since this is equivalent to conditioning on $N(\tau)$ for all τ in $(0, t]$, we have

$$\Pr\{Z > z \mid \{N(\tau), 0 < \tau \leq t\}\} = \exp(-\lambda z). \quad (2.12)$$

□

Next consider the sequence of interarrival times starting from a given time t . The first of these is taken to be the rv Z (now renamed Z_1) of the time from t to the first arrival epoch after t . For $m \geq 2$, let Z_m be the time from the $(m-1)$ st arrival after t to the m th arrival epoch after t . Thus, given $N(t) = n$ and $S_n = \tau$, we see that $Z_m = X_{m+n}$ for $m \geq 2$. It follows that, conditional on $N(t) = n$ and $S_n = \tau$, Z_1, Z_2, \dots , are IID exponentially distributed rv's. (see Exercise 2.8). Since this is independent of $N(t)$ and S_n , it follows that that Z_1, Z_2, \dots are unconditionally IID and also independent of $N(t)$ and S_n . It should also be clear that Z_1, Z_2, \dots are independent of $\{N(\tau); 0 < \tau \leq t\}$.

The above argument shows that the portion of a Poisson process starting at an arbitrary time $t > 0$ is a probabilistic replica of the process starting at 0; that is, the time until the first arrival after t is an exponentially distributed rv with parameter λ , and all subsequent interarrival times are independent of this first arrival epoch and of each other, and all have the same exponential distribution.

Definition 2.2.4. A counting process $\{N(t); t > 0\}$ has the stationary increment property if $N(t') - N(t)$ has the same CDF as $N(t' - t)$ for every $t' > t > 0$.

Let us define $\tilde{N}(t, t') = N(t') - N(t)$ as the number of arrivals in the interval $(t, t']$ for any given $t' \geq t$. We have just shown that for a Poisson process, the rv $\tilde{N}(t, t')$ has the same distribution as $N(t' - t)$, which means that a Poisson process has the stationary increment property. Thus, the distribution of the number of arrivals in an interval depends on the size of the interval but not on its starting point.

Definition 2.2.5. A counting process $\{N(t); t > 0\}$ has the independent increment property if, for every integer $k > 0$, and every k -tuple of times $0 < t_1 < t_2 < \dots < t_k$, the k -tuple of rv's $N(t_1), \tilde{N}(t_1, t_2), \dots, \tilde{N}(t_{k-1}, t_k)$ are statistically independent.

For the Poisson process, Theorem 2.2.1 says that for any t , the time Z_1 until the next arrival after t is independent of $N(\tau)$ for all $\tau \leq t$. Letting $t_1 < t_2 < \dots < t_{k-1} < t$, this means that Z_1 is independent of $N(t_1), \tilde{N}(t_1, t_2), \dots, \tilde{N}(t_{k-1}, t)$. We have also seen that the subsequent interarrival times after Z_1 , and thus $\tilde{N}(t, t')$ are independent of $N(t_1), \tilde{N}(t_1, t_2), \dots, \tilde{N}(t_{k-1}, t)$. Renaming t as t_k and t' as t_{k+1} , we see that $\tilde{N}(t_k, t_{k+1})$ is independent of $N(t_1), \tilde{N}(t_1, t_2), \dots, \tilde{N}(t_{k-1}, t_k)$. Since this is true for all k , the Poisson process has the independent increment property. In summary, we have proved the following:

Theorem 2.2.2. *Poisson processes have both the stationary increment and independent increment properties.*

Note that if we look only at integer times, then the Bernoulli process also has the stationary and independent increment properties.

2.2.2 Probability density of S_n and joint density of S_1, \dots, S_n

Recall from (2.1) that, for a Poisson process, S_n is the sum of n IID rv's, each with the density function $f_X(x) = \lambda \exp(-\lambda x)$, $x \geq 0$. Also recall that the density of the sum of two independent rv's can be found by convolving their densities, and thus the density of S_2 can be found by convolving $f_X(x)$ with itself, S_3 by convolving the density of S_2 with $f_X(x)$, and so forth. The result, for $t \geq 0$, is called the *Erlang density*,⁵

$$f_{S_n}(t) = \frac{\lambda^n t^{n-1} \exp(-\lambda t)}{(n-1)!}. \quad (2.13)$$

We can understand this density (and other related matters) better by viewing the above mechanical derivation in a slightly different way involving the joint density of S_1, \dots, S_n . For $n = 2$, the joint density of X_1 and S_2 (or equivalently, S_1 and S_2) is given by

$$f_{X_1 S_2}(x_1, s_2) = f_{X_1}(x_1) f_{S_2|X_1}(s_2|x_1) = \lambda e^{-\lambda x_1} \lambda e^{-\lambda(s_2-x_1)} \quad \text{for } 0 \leq x_1 \leq s_2,$$

where, since $S_2 = X_1 + X_2$, we have used the fact that the density of S_2 conditional on X_1 is just the exponential interarrival density evaluated at $S_2 - X_1$. Thus,

$$f_{X_1 S_2}(x_1, s_2) = \lambda^2 \exp(-\lambda s_2) \quad \text{for } 0 \leq x_1 \leq s_2. \quad (2.14)$$

This says that the joint density does not contain x_1 , except for the constraint $0 \leq x_1 \leq s_2$. Thus, for fixed s_2 , the joint density, and thus the conditional density of X_1 given $S_2 = s_2$ is uniform over $0 \leq x_1 \leq s_2$. The integration over x_1 in the convolution equation is then simply multiplication by the interval size s_2 , yielding the marginal PDF $f_{S_2}(s_2) = \lambda^2 s_2 \exp(-\lambda s_2)$, in agreement with (2.13) for $n = 2$.

The following theorem shows that this same curious behavior exhibits itself for the sum of an arbitrary number n of IID exponential rv's.

⁵Another (somewhat rarely used) name for the Erlang density is the *gamma density*.

Theorem 2.2.3. *Let X_1, X_2, \dots be IID rv's with the density $f_X(x) = \lambda e^{-\lambda x}$ for $x \geq 0$. Let $S_n = X_1 + \dots + X_n$ for each $n \geq 1$. Then for each $n \geq 2$*

$$f_{S_1 \dots S_n}(s_1, \dots, s_n) = \lambda^n \exp(-\lambda s_n) \quad \text{for } 0 \leq s_1 \leq s_2 \leq \dots \leq s_n. \quad (2.15)$$

Proof: Replacing X_1 with S_1 in (2.14), we see the theorem holds for $n = 2$. This serves as the basis for the following inductive proof. Assume (2.15) holds for given n . Then

$$\begin{aligned} f_{S_1 \dots S_{n+1}}(s_1, \dots, s_{n+1}) &= f_{S_1 \dots S_n}(s_1, \dots, s_n) f_{S_{n+1}|S_1 \dots S_n}(s_{n+1}|s_1, \dots, s_n) \\ &= \lambda^n \exp(-\lambda s_n) f_{S_{n+1}|S_1 \dots S_n}(s_{n+1}|s_1, \dots, s_n), \end{aligned} \quad (2.16)$$

where we used (2.15) for the given n . Now $S_{n+1} = S_n + X_{n+1}$. Since X_{n+1} is independent of S_1, \dots, S_n ,

$$f_{S_{n+1}|S_1 \dots S_n}(s_{n+1}|s_1, \dots, s_n) = \lambda \exp(-\lambda(s_{n+1} - s_n)).$$

Substituting this into (2.16) yields (2.15) □

The interpretation here is the same as with S_2 . The joint density does not contain any arrival time other than s_n , except for the ordering constraint $0 \leq s_1 \leq s_2 \leq \dots \leq s_n$, and thus this joint density is constant over all choices of arrival times satisfying the ordering constraint for a fixed s_n . Mechanically integrating this over s_1 , then s_2 , etc. we get the Erlang formula (2.13). The Erlang density then is the joint density in (2.15) times the volume $s_n^{n-1}/(n-1)!$ of the region of s_1, \dots, s_{n-1} satisfying $0 < s_1 < \dots < s_{n-1} < s_n$. This will be discussed further later.

Note that (2.15), for all n , specifies the joint distribution for all arrival epochs, and thus fully specifies a Poisson process.

2.2.3 The PMF for $N(t)$

The Poisson counting process, $\{N(t); t > 0\}$ consists of a nonnegative integer rv $N(t)$ for each $t > 0$. In this section, we show that the PMF for this rv is the well-known Poisson PMF, as stated in the following theorem. We give two proofs for the theorem, each providing its own type of understanding and each showing the close relationship between $\{N(t) = n\}$ and $\{S_n = t\}$.

Theorem 2.2.4. *For a Poisson process of rate λ , and for any $t > 0$, the PMF for $N(t)$, i.e., the number of arrivals in $(0, t]$, is given by the Poisson PMF,*

$$p_{N(t)}(n) = \frac{(\lambda t)^n \exp(-\lambda t)}{n!}. \quad (2.17)$$

Proof 1: This proof, for given n and t , is based on two ways of calculating $\Pr\{t < S_{n+1} \leq t + \delta\}$ for some vanishingly small δ . The first way is based on the already known density of S_{n+1} and gives

$$\Pr\{t < S_{n+1} \leq t + \delta\} = \int_t^{t+\delta} f_{S_{n+1}}(\tau) d\tau = f_{S_{n+1}}(t) (\delta + o(\delta)).$$

The term $o(\delta)$ is used to describe a function of δ that goes to 0 faster than δ as $\delta \rightarrow 0$. More precisely, a function $g(\delta)$ is said to be of order $o(\delta)$ if $\lim_{\delta \rightarrow 0} \frac{g(\delta)}{\delta} = 0$. Thus $\Pr\{t < S_{n+1} \leq t + \delta\} = f_{S_{n+1}}(t)(\delta + o(\delta))$ is simply a consequence of the fact that S_{n+1} has a continuous probability density in the interval $[t, t + \delta]$.

The second way to calculate $\Pr\{t < S_{n+1} \leq t + \delta\}$ is to first observe that the probability of more than 1 arrival in $(t, t + \delta]$ is $o(\delta)$. Ignoring this possibility, $\{t < S_{n+1} \leq t + \delta\}$ occurs if exactly n arrivals are in the interval $(0, t]$ and one arrival occurs in $(t, t + \delta]$. Because of the independent increment property, this is an event of probability $p_{N(t)}(n)(\lambda\delta + o(\delta))$. Thus

$$p_{N(t)}(n)(\lambda\delta + o(\delta)) + o(\delta) = f_{S_{n+1}}(t)(\delta + o(\delta)).$$

Dividing by δ and taking the limit $\delta \rightarrow 0$, we get

$$\lambda p_{N(t)}(n) = f_{S_{n+1}}(t).$$

Using the density for f_{S_n} in (2.13), we get (2.17). \square

Proof 2: The approach here is to use the fundamental relation that $\{N(t) \geq n\} = \{S_n \leq t\}$. Taking the probabilities of these events,

$$\sum_{i=n}^{\infty} p_{N(t)}(i) = \int_0^t f_{S_n}(\tau) d\tau \quad \text{for all } n \geq 1 \text{ and } t > 0.$$

The term on the right above is the CDF of S_n and the term on the left is the complementary CDF of $N(t)$. The complementary CDF and the PMF of $N(t)$ uniquely specify each other, so the theorem is equivalent to showing that

$$\sum_{i=n}^{\infty} \frac{(\lambda t)^i \exp(-\lambda t)}{i!} = \int_0^t f_{S_n}(\tau) d\tau. \quad (2.18)$$

If we take the derivative with respect to t of each side of (2.18), we find that almost magically each term except the first on the left cancels out, leaving us with

$$\frac{\lambda^n t^{n-1} \exp(-\lambda t)}{(n-1)!} = f_{S_n}(t).$$

Thus the derivative with respect to t of each side of (2.18) is equal to the derivative of the other for all $n \geq 1$ and $t > 0$. The two sides of (2.18) are also equal in the limit $t \rightarrow \infty$, so it follows that (2.18) is satisfied everywhere, completing the proof. \square

2.2.4 Alternate definitions of Poisson processes

Definition 2 of a Poisson process: A Poisson counting process $\{N(t); t > 0\}$ is a counting process that satisfies (2.17) (i.e., has the Poisson PMF) and has the independent and stationary increment properties.

We have seen that the properties in Definition 2 are satisfied starting with Definition 1 (using IID exponential interarrival times), so Definition 1 implies Definition 2. Exercise

2.4 shows that IID exponential interarrival times are implied by Definition 2, so the two definitions are equivalent.

It may be somewhat surprising at first to realize that a counting process that has the Poisson PMF at each t is not necessarily a Poisson process, and that the independent and stationary increment properties are also necessary. One way to see this is to recall that the Poisson PMF for all t in a counting process is equivalent to the Erlang density for the successive arrival epochs. Specifying the probability density for S_1, S_2, \dots as Erlang specifies the *marginal* densities of S_1, S_2, \dots but need not specify the *joint* densities of these rv's. Figure 2.4 illustrates this in terms of the joint density of S_1, S_2 , given as

$$f_{S_1, S_2}(s_1, s_2) = \lambda^2 \exp(-\lambda s_2) \quad \text{for } 0 \leq s_1 \leq s_2$$

and 0 elsewhere. The figure illustrates how the joint density can be changed without changing the marginals.

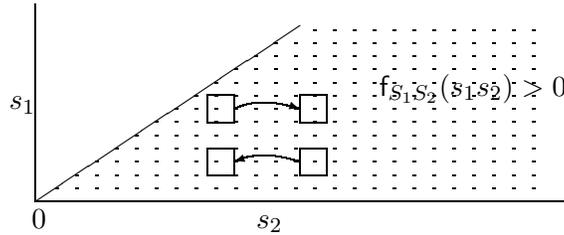


Figure 2.4: The joint density of S_1, S_2 is non-zero in the region shown. It can be changed, while holding the marginals constant, by reducing the joint density by ϵ in the upper left and lower right squares above and increasing it by ϵ in the upper right and lower left squares.

There is a similar effect with the Bernoulli process in that a discrete counting process for which the number of arrivals from 0 to t , for each integer t , is a binomial rv, but the process is not Bernoulli. This is explored in Exercise 2.5.

The next definition of a Poisson process is based on its incremental properties. Consider the number of arrivals in some very small interval $(t, t + \delta]$. Since $\tilde{N}(t, t + \delta)$ has the same distribution as $N(\delta)$, we can use (2.17) to get

$$\begin{aligned} \Pr\left\{\tilde{N}(t, t + \delta) = 0\right\} &= e^{-\lambda\delta} \approx 1 - \lambda\delta + o(\delta) \\ \Pr\left\{\tilde{N}(t, t + \delta) = 1\right\} &= \lambda\delta e^{-\lambda\delta} \approx \lambda\delta + o(\delta) \\ \Pr\left\{\tilde{N}(t, t + \delta) \geq 2\right\} &\approx o(\delta). \end{aligned} \tag{2.19}$$

Definition 3 of a Poisson process: *A Poisson counting process is a counting process that satisfies (2.19) and has the stationary and independent increment properties.*

We have seen that Definition 1 implies Definition 3. The essence of the argument the other way is that for any interarrival interval X , $F_X(x + \delta) - F_X(x)$ is the probability of an arrival

in an appropriate infinitesimal interval of width δ , which by (2.19) is $\lambda\delta + o(\delta)$. Turning this into a differential equation (see Exercise 2.7), we get the desired exponential interarrival intervals. Definition 3 has an intuitive appeal, since it is based on the idea of independent arrivals during arbitrary disjoint intervals. It has the disadvantage that one must do a considerable amount of work to be sure that these conditions are mutually consistent, and probably the easiest way is to start with Definition 1 and derive these properties. Showing that there is a unique process that satisfies the conditions of Definition 3 is even harder, but is not necessary at this point, since all we need is the use of these properties. Section 2.2.5 will illustrate better how to use this definition (or more precisely, how to use (2.19)).

What (2.19) accomplishes in Definition 3, beyond the assumption of independent and stationary increments, is the prevention of bulk arrivals. For example, consider a counting process in which arrivals always occur in pairs, and the intervals between successive pairs are IID and exponentially distributed with parameter λ (see Figure 2.5). For this process, $\Pr\{\tilde{N}(t, t+\delta)=1\} = 0$, and $\Pr\{\tilde{N}(t, t+\delta)=2\} = \lambda\delta + o(\delta)$, thus violating (2.19). This process has stationary and independent increments, however, since the process formed by viewing a pair of arrivals as a single incident is a Poisson process.

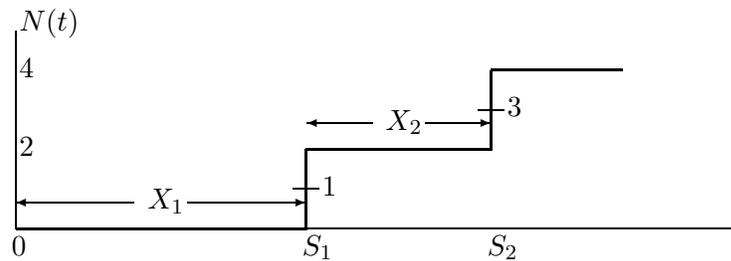


Figure 2.5: A counting process modeling bulk arrivals. X_1 is the time until the first pair of arrivals and X_2 is the interval between the first and second pair of arrivals.

2.2.5 The Poisson process as a limit of shrinking Bernoulli processes

The intuition of Definition 3 can be achieved in a less abstract way by starting with the Bernoulli process, which has the properties of Definition 3 in a discrete-time sense. We then go to an appropriate limit of a sequence of these processes, and find that this sequence of Bernoulli processes converges in some sense to the Poisson process.

Recall that a Bernoulli process is an IID sequence, Y_1, Y_2, \dots of binary random variables for which $\Pr\{Y_i=1\} = p$ and $\Pr\{Y_i=0\} = 1-p$. We can visualize $Y_i = 1$ as an *arrival* at time i and $Y_i = 0$ as no arrival, but we can also ‘shrink’ the time scale of the process so that for some integer $j > 0$, Y_i is an arrival or no arrival at time $i2^{-j}$. We consider a sequence indexed by j of such shrinking Bernoulli processes, and in order to keep the arrival rate constant, we let $p = \lambda 2^{-j}$ for the j th process. Thus for each unit increase in j , the Bernoulli process shrinks by replacing each slot with two slots, each with half the previous arrival probability. The expected number of arrivals per unit time is then λ , matching the Poisson process that

we are approximating.

If we look at this j th process relative to Definition 3 of a Poisson process, we see that for these regularly spaced increments of size $\delta = 2^{-j}$, the probability of one arrival in an increment is $\lambda\delta$ and that of no arrival is $1 - \lambda\delta$, and thus (2.19) is satisfied, and in fact the $o(\delta)$ terms are exactly zero. For arbitrary sized increments, it is clear that disjoint increments have essentially independent arrivals. The increments are not quite stationary, since, for example, an increment of size 2^{-j-1} might contain a time that is a multiple of 2^{-j} or might not, depending on its placement. However, for any fixed increment of size δ , the number of multiples of 2^{-j} (i.e., the number of possible arrival points) is either $\lfloor \delta 2^j \rfloor$ or $1 + \lfloor \delta 2^j \rfloor$. Thus in the limit $j \rightarrow \infty$, the increments are both stationary and independent.

For each j , the j th Bernoulli process has an associated Bernoulli counting process $N_j(t) = \sum_{i=1}^{\lfloor t2^j \rfloor} Y_i$. This is the number of arrivals up to time t and is a discrete rv with the binomial PMF. That is, $\mathbf{p}_{N_j(t)}(n) = \binom{\lfloor t2^j \rfloor}{n} p^n (1-p)^{\lfloor t2^j \rfloor - n}$ where $p = \lambda 2^{-j}$. We now show that this PMF approaches the Poisson PMF as j increases.⁶

Theorem 2.2.5 (Poisson's theorem). *Consider the sequence of shrinking Bernoulli processes with arrival probability $\lambda 2^{-j}$ and time-slot size 2^{-j} . Then for every fixed time $t > 0$ and fixed number of arrivals n , the counting PMF $\mathbf{p}_{N_j(t)}(n)$ approaches the Poisson PMF (of the same λ) with increasing j , i.e.,*

$$\lim_{j \rightarrow \infty} \mathbf{p}_{N_j(t)}(n) = \mathbf{p}_{N(t)}(n). \quad (2.20)$$

Proof: We first rewrite the binomial PMF, for $\lfloor t2^j \rfloor$ variables with $p = \lambda 2^{-j}$ as

$$\begin{aligned} \lim_{j \rightarrow \infty} \mathbf{p}_{N_j(t)}(n) &= \lim_{j \rightarrow \infty} \binom{\lfloor t2^j \rfloor}{n} \left(\frac{\lambda 2^{-j}}{1 - \lambda 2^{-j}} \right)^n \exp[\lfloor t2^j \rfloor \ln(1 - \lambda 2^{-j})] \\ &= \lim_{j \rightarrow \infty} \binom{\lfloor t2^j \rfloor}{n} \left(\frac{\lambda 2^{-j}}{1 - \lambda 2^{-j}} \right)^n \exp(-\lambda t) \end{aligned} \quad (2.21)$$

$$= \lim_{j \rightarrow \infty} \frac{\lfloor t2^j \rfloor \cdot \lfloor t2^j \rfloor - 1 \cdots \lfloor t2^j \rfloor - n + 1}{n!} \left(\frac{\lambda 2^{-j}}{1 - \lambda 2^{-j}} \right)^n \exp(-\lambda t) \quad (2.22)$$

$$= \frac{(\lambda t)^n \exp(-\lambda t)}{n!}. \quad (2.23)$$

We used $\ln(1 - \lambda 2^{-j}) = -\lambda 2^{-j} + o(2^{-j})$ in (2.21) and expanded the combinatorial term in (2.22). In (2.23), we recognized that $\lim_{j \rightarrow \infty} \lfloor t2^j \rfloor - i = \lambda t$ for $0 \leq i \leq n - 1$. \square

Since the binomial PMF (scaled as above) has the Poisson PMF as a limit for each n , the CDF of $N_j(t)$ also converges to the Poisson CDF for each t . In other words, for each $t > 0$, the counting random variables $N_j(t)$ of the Bernoulli processes converge in distribution to $N(t)$ of the Poisson process.

⁶This limiting result for the binomial distribution is very different from the asymptotic results in Chapter 1 for the binomial. Here the parameter p of the binomial is shrinking with increasing j , whereas there, p is constant while the number of variables is increasing.

This does not say that the Bernoulli *counting processes* converge to the Poisson counting process in any meaningful sense, since the joint distributions are also of concern. The following corollary treats this.

Corollary 2.2.1. *For any finite integer $k > 0$, let $0 < t_1 < t_2 < \dots < t_k$ be any set of time instants. Then the joint CDF of $N_j(t_1), N_j(t_2), \dots, N_j(t_k)$ approaches the joint CDF of $N(t_1), N(t_2), \dots, N(t_k)$ as $j \rightarrow \infty$.*

Proof: We can rewrite the joint PMF for each Bernoulli process as

$$\begin{aligned} \mathbf{p}_{N_j(t_1), \dots, N_j(t_k)}(n_1, \dots, n_k) &= \mathbf{p}_{N_j(t_1), \tilde{N}_j(t_1, t_2), \dots, \tilde{N}_j(t_{k-1}, t_k)}(n_1, n_2 - n_1, \dots, n_k - n_{k-1}) \\ &= \mathbf{p}_{N_j(t_1)}(n_1) \prod_{\ell=2}^k \mathbf{p}_{\tilde{N}_j(t_{\ell-1}, t_{\ell})}(n_{\ell} - n_{\ell-1}) \end{aligned} \quad (2.24)$$

where we have used the independent increment property for the Bernoulli process. For the Poisson process, we similarly have

$$\mathbf{p}_{N(t_1), \dots, N(t_k)}(n_1, \dots, n_k) = \mathbf{p}_{N(t_1)}(n_1) \prod_{\ell=2}^k \mathbf{p}_{\tilde{N}(t_{\ell-1}, t_{\ell})}(n_{\ell} - n_{\ell-1}) \quad (2.25)$$

Taking the limit of (2.24) as $j \rightarrow \infty$, we recognize from Theorem 2.2.5 that each term of (2.24) goes to the corresponding term in (2.25). For the \tilde{N} rv's, this requires a trivial generalization in Theorem 2.2.5 to deal with the arbitrary starting time. \square

We conclude from this that the sequence of Bernoulli processes above converges to the Poisson process in the sense of the corollary. Recall from Section 1.7.5 that there are a number of ways in which a sequence of rv's can converge. As one might imagine, there are many more ways in which a sequence of stochastic processes can converge, and the corollary simply establishes one of these. Note however that showing only that $p_{N_j(t)}(n)$ approaches $p_{N(t)}(n)$ for each t is too weak to be very helpful, because it shows nothing about the time-evolution of the process. On the other hand, we don't even know how to define a joint distribution over an infinite number of epochs, let alone deal with limiting characteristics. Considering arbitrary finite sets of epochs forms a good middle ground.

Both the Poisson process and the Bernoulli process are so easy to analyze that the convergence of shrinking Bernoulli processes to Poisson is rarely the easiest way to establish properties about either. On the other hand, this convergence is a powerful aid to the intuition in understanding each process. In other words, the relation between Bernoulli and Poisson is very useful in suggesting new ways of looking at problems, but is usually not the best way to analyze those problems.

2.3 Combining and splitting Poisson processes

Suppose that $\{N_1(t); t > 0\}$ and $\{N_2(t); t > 0\}$ are independent Poisson counting processes⁷ of rates λ_1 and λ_2 respectively. We want to look at the sum process where $N(t) = N_1(t) +$

⁷Two processes $\{N_1(t); t > 0\}$ and $\{N_2(t); t > 0\}$ are said to be independent if for all positive integers k and all sets of times $0 < t_1 < t_2 < \dots < t_k$, the random variables $N_1(t_1), \dots, N_1(t_k)$ are independent

$N_2(t)$ for all $t \geq 0$. In other words, $\{N(t); t > 0\}$ is the process consisting of all arrivals to both process 1 and process 2. We shall show that $\{N(t); t > 0\}$ is a Poisson counting process of rate $\lambda = \lambda_1 + \lambda_2$. We show this in three different ways, first using Definition 3 of a Poisson process (since that is most natural for this problem), then using Definition 2, and finally Definition 1. We then draw some conclusions about the way in which each approach is helpful. Since $\{N_1(t); t > 0\}$ and $\{N_2(t); t > 0\}$ are independent and both possess the stationary and independent increment properties, it follows from the definitions that $\{N(t); t > 0\}$ also possesses the stationary and independent increment properties. Using the approximations in (2.19) for the individual processes, we see that

$$\begin{aligned} \Pr\{\tilde{N}(t, t + \delta) = 0\} &= \Pr\{\tilde{N}_1(t, t + \delta) = 0\} \Pr\{\tilde{N}_2(t, t + \delta) = 0\} \\ &= (1 - \lambda_1 \delta)(1 - \lambda_2 \delta) \approx 1 - \lambda \delta. \end{aligned}$$

where $\lambda_1 \lambda_2 \delta^2$ has been dropped. In the same way, $\Pr\{\tilde{N}(t, t + \delta) = 1\}$ is approximated by $\lambda \delta$ and $\Pr\{\tilde{N}(t, t + \delta) \geq 2\}$ is approximated by 0, both with errors proportional to δ^2 . It follows that $\{N(t); t > 0\}$ is a Poisson process.

In the second approach, we have $N(t) = N_1(t) + N_2(t)$. Since $N(t)$, for any given t , is the sum of two independent Poisson rv's, it is also a Poisson rv with mean $\lambda t = \lambda_1 t + \lambda_2 t$. If the reader is not aware that the sum of two independent Poisson rv's is Poisson, it can be derived by discrete convolution of the two PMF's (see Exercise 1.22). More elegantly, one can observe that we have already implicitly shown this fact. That is, if we break an interval I into disjoint subintervals, I_1 and I_2 , then the number of arrivals in I (which is Poisson) is the sum of the number of arrivals in I_1 and in I_2 (which are independent Poisson). Finally, since $N(t)$ is Poisson for each t , and since the stationary and independent increment properties are satisfied, $\{N(t); t > 0\}$ is a Poisson process.

In the third approach, X_1 , the first interarrival interval for the sum process, is the minimum of X_{11} , the first interarrival interval for the first process, and X_{21} , the first interarrival interval for the second process. Thus $X_1 > t$ if and only if both X_{11} and X_{21} exceed t , so

$$\Pr\{X_1 > t\} = \Pr\{X_{11} > t\} \Pr\{X_{21} > t\} = \exp(-\lambda_1 t - \lambda_2 t) = \exp(-\lambda t).$$

Using the memoryless property, each subsequent interarrival interval can be analyzed in the same way.

The first approach above was the most intuitive for this problem, but it required constant care about the order of magnitude of the terms being neglected. The second approach was the simplest analytically (after recognizing that sums of independent Poisson rv's are Poisson), and required no approximations. The third approach was very simple in retrospect, but not very natural for this problem. If we add many independent Poisson processes together, it is clear, by adding them one at a time, that the sum process is again Poisson. What is more interesting is that when many independent counting processes (not necessarily Poisson) are added together, the sum process often tends to be approximately Poisson if

of $N_2(t_1), \dots, N_2(t_k)$. Here it is enough to extend the independent increment property to independence between increments over the two processes; equivalently, one can require the interarrival intervals for one process to be independent of the interarrivals for the other process.

the individual processes have small rates compared to the sum. To obtain some crude intuition about why this might be expected, note that the interarrival intervals for each process (assuming no bulk arrivals) will tend to be large relative to the mean interarrival interval for the sum process. Thus arrivals that are close together in time will typically come from different processes. The number of arrivals in an interval large relative to the combined mean interarrival interval, but small relative to the individual interarrival intervals, will be the sum of the number of arrivals from the different processes; each of these is 0 with large probability and 1 with small probability, so the sum will be approximately Poisson.

2.3.1 Subdividing a Poisson process

Next we look at how to break $\{N(t); t > 0\}$, a Poisson counting process of rate λ , into two processes, $\{N_1(t); t > 0\}$ and $\{N_2(t); t > 0\}$. Suppose that each arrival in $\{N(t); t > 0\}$ is sent to the first process with probability p and to the second process with probability $1 - p$ (see Figure 2.6). Each arrival is switched independently of each other arrival and independently of the arrival epochs. It may be helpful to visualize this as the combination of two independent processes. The first is the Poisson process of rate λ and the second is a Bernoulli process $\{X_n; n \geq 1\}$ where $p_{X_n}(1) = p$ and $p_{X_n}(2) = 1 - p$. The n th arrival of the Poisson process is, with probability p , labeled as a type 1 arrival *i.e.*, labeled as $X_n = 1$. With probability $1 - p$, it is labeled as type 2, *i.e.*, labeled as $X_n = 2$.

We shall show that the resulting processes are each Poisson, with rates $\lambda_1 = \lambda p$ and $\lambda_2 = \lambda(1 - p)$ respectively, and that furthermore the two processes are independent. Note that, conditional on the original process, the two new processes are not independent; in fact one completely determines the other. Thus this independence might be a little surprising.

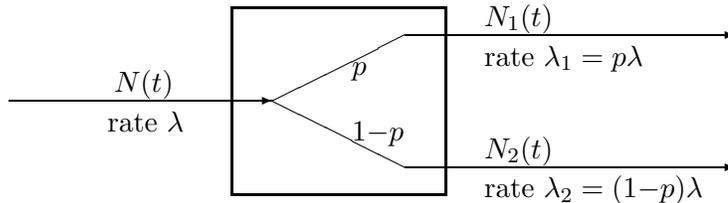


Figure 2.6: Each arrival is independently sent to process 1 with probability p and to process 2 otherwise.

First consider a small increment $(t, t + \delta]$. The original process has an arrival in this incremental interval with probability $\lambda\delta$ (ignoring δ^2 terms as usual), and thus process 1 has an arrival with probability $\lambda\delta p$ and process 2 with probability $\lambda\delta(1 - p)$. Because of the independent increment property of the original process and the independence of the division of each arrival between the two processes, the new processes each have the independent increment property, and from above have the stationary increment property. Thus each process is Poisson. Note now that we cannot verify that the two processes are independent from this small increment model. We would have to show that the number of arrivals for process 1 and 2 are independent over $(t, t + \delta]$. Unfortunately, leaving out the terms of order δ^2 , there is at most one arrival to the original process and no possibility of

an arrival to each new process in $(t, t + \delta]$. If it is impossible for both processes to have an arrival in the same interval, they cannot be independent. It is possible, of course, for each process to have an arrival in the same interval, but this is a term of order δ^2 . Thus, without paying attention to the terms of order δ^2 , it is impossible to demonstrate that the processes are independent.

To demonstrate that process 1 and 2 are independent, we first calculate the joint PMF for $N_1(t)$, $N_2(t)$ for arbitrary t . Conditioning on a given number of arrivals $N(t)$ for the original process, we have

$$\Pr\{N_1(t)=m, N_2(t)=k \mid N(t)=m+k\} = \frac{(m+k)!}{m!k!} p^m (1-p)^k. \quad (2.26)$$

Equation (2.26) is simply the binomial distribution, since, given $m+k$ arrivals to the original process, each independently goes to process 1 with probability p . Since the event $\{N_1(t) = m, N_2(t) = k\}$ is a subset of the conditioning event above,

$$\Pr\{N_1(t)=m, N_2(t)=k \mid N(t)=m+k\} = \frac{\Pr\{N_1(t)=m, N_2(t)=k\}}{\Pr\{N(t)=m+k\}}.$$

Combining this with (2.26), we have

$$\Pr\{N_1(t)=m, N_2(t)=k\} = \frac{(m+k)!}{m!k!} p^m (1-p)^k \frac{(\lambda t)^{m+k} e^{-\lambda t}}{(m+k)!}. \quad (2.27)$$

Rearranging terms, we get

$$\Pr\{N_1(t)=m, N_2(t)=k\} = \frac{(p\lambda t)^m e^{-\lambda p t}}{m!} \frac{[(1-p)\lambda t]^k e^{-\lambda(1-p)t}}{k!}. \quad (2.28)$$

This shows that $N_1(t)$ and $N_2(t)$ are independent. To show that the processes are independent, we must show that for any $k > 1$ and any set of times $0 \leq t_1 \leq t_2 \leq \dots \leq t_k$, the sets $\{N_1(t_i); 1 \leq i \leq k\}$ and $\{N_2(t_j); 1 \leq j \leq k\}$ are independent of each other. It is equivalent to show that the sets $\{\tilde{N}_1(t_{i-1}, t_i); 1 \leq i \leq k\}$ and $\{\tilde{N}_2(t_{j-1}, t_j); 1 \leq j \leq k\}$ (where t_0 is 0) are independent. The argument above shows this independence for $i = j$, and for $i \neq j$, the independence follows from the independent increment property of $\{N(t); t > 0\}$.

2.3.2 Examples using independent Poisson processes

We have observed that if the arrivals of a Poisson process are split into two new arrival processes, with each arrival of the original process independently entering the first of the new processes with some fixed probability p , then each new process is Poisson and independent of the other. The most useful consequence of this is that any two independent Poisson processes can be viewed as being generated from a single process in this way. Thus, if one process has rate λ_1 and the other has rate λ_2 , they can be viewed as coming from a process of rate $\lambda_1 + \lambda_2$. Each arrival to the combined process is then labeled as a first-process arrival with probability $p = \lambda_1/(\lambda_1 + \lambda_2)$ and as a second-process arrival with probability $1 - p$.

Example 2.3.1. The above point of view is very useful for finding probabilities such as $\Pr\{S_{1k} < S_{2j}\}$ where S_{1k} is the epoch of the k th arrival to the first process and S_{2j} is the epoch of the j th arrival to the second process. The problem can be rephrased in terms of a combined process to ask: out of the first $k + j - 1$ arrivals to the combined process, what is the probability that k or more of them are switched to the first process? (Note that if k or more of the first $k + j - 1$ go to the first process, at most $j - 1$ go to the second, so the k th arrival to the first precedes the j th arrival to the second; similarly if fewer than k of the first $k + j - 1$ go to the first process, then the j th arrival to the second process precedes the k th arrival to the first). Since each of these first $k + j - 1$ arrivals are switched independently with the same probability p , the answer is

$$\Pr\{S_{1k} < S_{2j}\} = \sum_{i=k}^{k+j-1} \frac{(k+j-1)!}{i!(k+j-1-i)!} p^i (1-p)^{k+j-1-i}. \quad (2.29)$$

Example 2.3.2 (The M/M/1 queue). Queueing theorists use a standard notation of characters separated by slashes to describe common types of queueing systems. The first character describes the arrival process to the queue. M stands for memoryless and means a Poisson arrival process; D stands for deterministic and means that the interarrival interval is fixed and non-random; G stands for general interarrival distribution. We assume that the interarrival intervals are IID (thus making the arrival process a renewal process), but many authors use GI to explicitly indicate IID interarrivals. The second character describes the service process. The same letters are used, with M indicating exponentially distributed service times. The third character gives the number of servers.⁸ It is assumed, when this notation is used, that the service times are IID, independent of the arrival epochs, and independent of which server is used.

Consider an M/M/1 queue, *i.e.*, a queueing system with a Poisson arrival system (say of rate λ) and a single server that serves arriving customers in order with a service time distribution $F(y) = 1 - \exp[-\mu y]$. The service times are independent of each other and of the interarrival intervals. During any period when the server is busy, customers then leave the system according to a Poisson process (process 2) of rate μ . We then see that if j or more customers are waiting at a given time, then (2.29) gives the probability that the k th subsequent arrival comes before the j th departure.

Example 2.3.3 (The sum of geometrically many exponential rv's). Consider waiting for a taxi on a busy street during rush hour. We will model the epochs at which taxis pass our spot on the street as a Poisson process of rate λ . We assume that each taxi is independently empty (*i.e.*, will stop to pick us up) with probability p . The empty taxis then form a Poisson process of rate $p\lambda$ and our waiting time is exponential with parameter λp . Looked at slightly differently, we will be picked by the first taxi with probability p , by the second with probability $(1-p)p$, and in general, by the m th taxi with probability $(1-p)^{m-1}p$. In other words, the number of the taxi that picks us up (numbered in order of arrival) is a geometrically distributed *rv*, say M , and our waiting time is the sum of the first M interarrival times, *i.e.*, S_M .

⁸Sometimes a fourth character is added which gives the size of the queue plus service facility. Thus, for example, an M/M/m/m queueing system would have Poisson arrivals, m independent servers, each with exponential service time. If an arrival occurs when all servers are busy, then that arrival is dropped.

As illustrated by taxis, waiting times modeled by a geometric sum of exponentials are relatively common. An important example that we shall analyze in Chapter 7 is that of the service time distribution of an M/M/1 queue in steady state (*i.e.*, after the effect of the initially empty queue has disappeared). We will show there that the number M of customers in the system immediately after an arrival is geometric with the PMF

$$\Pr\{M = m\} = (\lambda/\mu)^{m-1}(1 - \lambda/\mu) \quad \text{for } m \geq 1.$$

Now the time that the new arrival spends in the system is the sum of M service times, each of which is exponential with parameter μ . These service times are independent of each other and of M . Thus the system time (the time in queue and in service) of the new arrival is exponential with parameter $\mu(1 - \lambda/\mu) = \mu - \lambda$.

2.4 Non-homogeneous Poisson processes

The Poisson process, as we defined it, is characterized by a constant arrival rate λ . It is often useful to consider a more general type of process in which the arrival rate varies as a function of time. A *non-homogeneous Poisson process* with time varying arrival rate $\lambda(t)$ is defined⁹ as a counting process $\{N(t); t > 0\}$ which has the independent increment property and, for all $t \geq 0, \delta > 0$, also satisfies:

$$\begin{aligned} \Pr\{\tilde{N}(t, t + \delta) = 0\} &= 1 - \delta\lambda(t) + o(\delta) \\ \Pr\{\tilde{N}(t, t + \delta) = 1\} &= \delta\lambda(t) + o(\delta) \\ \Pr\{\tilde{N}(t, t + \delta) \geq 2\} &= o(\delta). \end{aligned} \tag{2.30}$$

where $\tilde{N}(t, t + \delta) = N(t + \delta) - N(t)$. The non-homogeneous Poisson process does not have the stationary increment property.

One common application occurs in optical communication where a non-homogeneous Poisson process is often used to model the stream of photons from an optical modulator; the modulation is accomplished by varying the photon intensity $\lambda(t)$. We shall see another application shortly in the next example. Sometimes a Poisson process, as we defined it earlier, is called a homogeneous Poisson process.

We can use a “shrinking Bernoulli process” again to approximate a non-homogeneous Poisson process. To see how to do this, assume that $\lambda(t)$ is bounded away from zero. We partition the time axis into increments whose lengths δ vary inversely with $\lambda(t)$, thus holding the probability of an arrival in an increment at some fixed value $p = \delta\lambda(t)$. Thus,

⁹We assume that $\lambda(t)$ is right continuous, *i.e.*, that for each t , $\lambda(t)$ is the limit of $\lambda(t + \epsilon)$ as ϵ approaches 0 from above. This allows $\lambda(t)$ to contain discontinuities, as illustrated in Figure 2.7, but follows the convention that the value of the function at the discontinuity is the limiting value from the right. This convention is required in (2.30) to talk about the distribution of arrivals just to the right of time t .

temporarily ignoring the variation of $\lambda(t)$ within an increment,

$$\begin{aligned}\Pr\left\{\tilde{N}\left(t, t + \frac{p}{\lambda(t)}\right) = 0\right\} &= 1 - p + o(p) \\ \Pr\left\{\tilde{N}\left(t, t + \frac{p}{\lambda(t)}\right) = 1\right\} &= p + o(p) \\ \Pr\left\{\tilde{N}\left(t, t + \frac{p}{\lambda(t)}\right) \geq 2\right\} &= o(p).\end{aligned}\tag{2.31}$$

This partition is defined more precisely by defining $m(t)$ as

$$m(t) = \int_0^t \lambda(\tau) d\tau.\tag{2.32}$$

Then the i th increment ends at that t for which $m(t) = i p$.

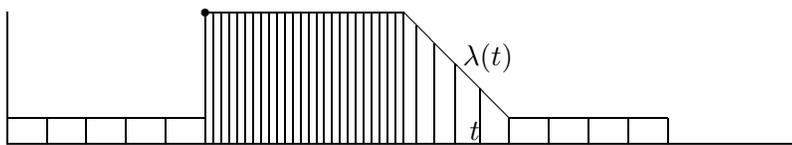


Figure 2.7: Partitioning the time axis into increments each with an expected number of arrivals equal to p . Each rectangle or trapezoid above has the same area, which ensures that the i th partition ends where $m(t) = i p$.

As before, let $\{Y_i; i \geq 1\}$ be a sequence of IID binary rv's with $\Pr\{Y_i = 1\} = p$ and $\Pr\{Y_i = 0\} = 1 - p$. Consider the counting process $\{N(t); t > 0\}$ in which Y_i , for each $i \geq 1$, denotes the number of arrivals in the interval $(t_{i-1}, t_i]$, where t_i satisfies $m(t_i) = i p$. Thus, $N(t_i) = Y_1 + Y_2 + \cdots + Y_i$. If p is decreased as 2^{-j} , each increment is successively split into a pair of increments. Thus by the same argument as in (2.23),

$$\Pr\{N(t) = n\} = \frac{[1 + o(p)][m(t)]^n \exp[-m(t)]}{n!}.\tag{2.33}$$

Similarly, for any interval $(t, \tau]$, taking $\tilde{m}(t, \tau) = \int_t^\tau \lambda(u) du$, and taking $t = t_k$, $\tau = t_i$ for some k, i , we get

$$\Pr\left\{\tilde{N}(t, \tau) = n\right\} = \frac{[1 + o(p)][\tilde{m}(t, \tau)]^n \exp[-\tilde{m}(t, \tau)]}{n!}.\tag{2.34}$$

Going to the limit $p \rightarrow 0$, the counting process $\{N(t); t > 0\}$ above approaches the non-homogeneous Poisson process under consideration, and we have the following theorem:

Theorem 2.4.1. *For a non-homogeneous Poisson process with right-continuous arrival rate $\lambda(t)$ bounded away from zero, the distribution of $\tilde{N}(t, \tau)$, the number of arrivals in $(t, \tau]$, satisfies*

$$\Pr\left\{\tilde{N}(t, \tau) = n\right\} = \frac{[\tilde{m}(t, \tau)]^n \exp[-\tilde{m}(t, \tau)]}{n!} \quad \text{where } \tilde{m}(t, \tau) = \int_t^\tau \lambda(u) du.\tag{2.35}$$

Hence, one can view a non-homogeneous Poisson process as a (homogeneous) Poisson process over a non-linear time scale. That is, let $\{N^*(s); s \geq 0\}$ be a (homogeneous) Poisson process with rate 1. The non-homogeneous Poisson process is then given by $N(t) = N^*(m(t))$ for each t .

Example 2.4.1 (The M/G/ ∞ Queue). Using the queueing notation explained in Example 2.3.2, an M/G/ ∞ queue indicates a queue with Poisson arrivals, a general service distribution, and an infinite number of servers. Since the M/G/ ∞ queue has an infinite number of servers, no arriving customers are ever queued. Each arrival immediately starts to be served by some server, and the service time Y_i of customer i is IID over i with some given CDF $G(y)$; the service time is the interval from start to completion of service and is also independent of arrival epochs. We would like to find the CDF of the number of customers being served at a given epoch τ .

Let $\{N(t); t > 0\}$ be the Poisson counting process, at rate λ , of customer arrivals. Consider the arrival times of those customers that are still in service at some fixed time τ . In some arbitrarily small interval $(t, t + \delta]$, the probability of an arrival is $\delta\lambda + o(\delta)$ and the probability of 2 or more arrivals is negligible (i.e., $o(\delta)$). The probability that a customer arrives in $(t, t + \delta]$ and is still being served at time $\tau > t$ is then $\delta\lambda[1 - G(\tau - t)] + o(\delta)$. Consider a counting process $\{N_1(t); 0 < t \leq \tau\}$ where $N_1(t)$ is the number of arrivals between 0 and t that are still in service at τ . This counting process has the independent increment property. To see this, note that the overall arrivals in $\{N(t); t > 0\}$ have the independent increment property; also the arrivals in $\{N(t); t > 0\}$ have independent service times, and thus are independently in or not in $\{N_1(t); 0 < t \leq \tau\}$. It follows that $\{N_1(t); 0 < t \leq \tau\}$ is a non-homogeneous Poisson process with rate $\lambda[1 - G(\tau - t)]$ at time $t \leq \tau$. The expected number of arrivals still in service at time τ is then

$$m(\tau) = \lambda \int_{t=0}^{\tau} [1 - G(\tau - t)] dt = \lambda \int_{t=0}^{\tau} [1 - G(t)] dt. \quad (2.36)$$

and the PMF of the number in service at time τ is given by

$$\Pr\{N_1(\tau) = n\} = \frac{m(\tau)^n \exp(-m(\tau))}{n!}. \quad (2.37)$$

Note that as $\tau \rightarrow \infty$, the integral in (2.36) approaches the mean of the service time distribution (i.e., it is the integral of the complementary CDF, $1 - G(t)$, of the service time). This means that in steady state (as $\tau \rightarrow \infty$), the distribution of the number in service at τ depends on the service time distribution only through its mean. This example can be used to model situations such as the number of phone calls taking place at a given epoch. This requires arrivals of new calls to be modeled as a Poisson process and the holding time of each call to be modeled as a random variable independent of other holding times and of call arrival times. Finally, as shown in Figure 2.8, we can regard $\{N_1(t); 0 < t \leq \tau\}$ as a splitting of the arrival process $\{N(t); t > 0\}$. By the same type of argument as in Section 2.3, the number of customers who have completed service by time τ is independent of the number still in service.

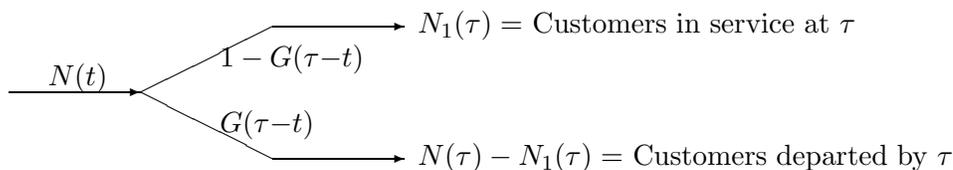


Figure 2.8: Poisson arrivals $\{N(t); t > 0\}$ can be considered to be split in a non-homogeneous way. An arrival at t is split with probability $1 - G(\tau - t)$ into a process of customers still in service at τ .

2.5 Conditional arrival densities and order statistics

A diverse range of problems involving Poisson processes are best tackled by conditioning on a given number n of arrivals in the interval $(0, t]$, i.e., on the event $N(t) = n$. Because of the incremental view of the Poisson process as independent and stationary arrivals in each incremental interval of the time axis, we would guess that the arrivals should have some sort of uniform distribution given $N(t) = n$. More precisely, the following theorem shows that the joint density of $\mathbf{S}^{(n)} = (S_1, S_2, \dots, S_n)$ given $N(t) = n$ is uniform over the region $0 < S_1 < S_2 < \dots < S_n < t$.

Theorem 2.5.1. *Let $f_{\mathbf{S}^{(n)}|N(t)=n}(\mathbf{s}^{(n)} | n)$ be the joint density of $\mathbf{S}^{(n)}$ conditional on $N(t) = n$. This density is constant over the region $0 < s_1 < \dots < s_n < t$ and has the value*

$$f_{\mathbf{S}^{(n)}|N(t)}(\mathbf{s}^{(n)} | n) = \frac{n!}{t^n}. \quad (2.38)$$

Two proofs are given, each illustrative of useful techniques.

Proof 1: Recall that the joint density of the first $n+1$ arrivals $\mathbf{S}^{(n+1)} = (S_1, \dots, S_n, S_{n+1})$ with no conditioning is given in (2.15). We first use Bayes law to calculate the joint density of $\mathbf{S}^{(n+1)}$ conditional on $N(t) = n$.

$$f_{\mathbf{S}^{(n+1)}|N(t)}(\mathbf{s}^{(n+1)} | n) p_{N(t)}(n) = p_{N(t)|\mathbf{S}^{(n+1)}}(n | \mathbf{s}^{(n+1)}) f_{\mathbf{S}^{(n+1)}}(\mathbf{s}^{(n+1)}).$$

Note that $N(t) = n$ if and only if $S_n \leq t$ and $S_{n+1} > t$. Thus $p_{N(t)|\mathbf{S}^{(n+1)}}(n | \mathbf{s}^{(n+1)})$ is 1 if $S_n \leq t$ and $S_{n+1} > t$ and is 0 otherwise. Restricting attention to the case $N(t) = n$, $S_n \leq t$ and $S_{n+1} > t$,

$$\begin{aligned} f_{\mathbf{S}^{(n+1)}|N(t)}(\mathbf{s}^{(n+1)} | n) &= \frac{f_{\mathbf{S}^{(n+1)}}(\mathbf{s}^{(n+1)})}{p_{N(t)}(n)} \\ &= \frac{\lambda^{n+1} \exp(-\lambda s_{n+1})}{(\lambda t)^n \exp(-\lambda t) / n!} \\ &= \frac{n! \lambda \exp[-\lambda(s_{n+1} - t)]}{t^n}. \end{aligned} \quad (2.39)$$

This is a useful expression, but we are interested in $\mathbf{S}^{(n)}$ rather than $\mathbf{S}^{(n+1)}$. Thus we break up the left side of (2.39) as follows:

$$f_{\mathbf{S}^{(n+1)}|N(t)}(\mathbf{s}^{(n+1)} | n) = f_{\mathbf{S}^{(n)}|N(t)}(\mathbf{s}^{(n)} | n) f_{S_{n+1}|\mathbf{S}^{(n)}N(t)}(s_{n+1} | \mathbf{s}^{(n)}, n).$$

Conditional on $N(t) = n$, S_{n+1} is the first arrival epoch after t , which by the memoryless property is conditionally independent of $\mathbf{S}^{(n)}$, again given $N(t) = n$. Thus that final term is simply $\lambda \exp(-\lambda(s_{n+1} - t))$ for $s_{n+1} > t$. Substituting this into (2.39), the result is (2.38). \square

Proof 2: This alternative proof derives (2.38) by looking at arrivals in very small increments of size δ (see Figure 2.9). For a given t and a given set of n times, $0 < s_1 < \dots < s_n < t$, we calculate the probability that there is a single arrival in each of the intervals $(s_i, s_i + \delta]$, $1 \leq i \leq n$ and no other arrivals in the interval $(0, t]$. Letting $A(\delta)$ be this event,

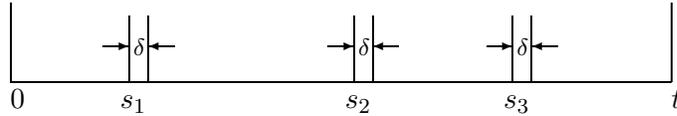


Figure 2.9: Intervals for the event $A(t)$ used to derive the joint arrival density.

$$\Pr\{A(\delta)\} = \mathbf{p}_{N(s_1)}(0) \mathbf{p}_{\tilde{N}(s_1, s_1 + \delta)}(1) \mathbf{p}_{\tilde{N}(s_1 + \delta, s_2)}(0) \mathbf{p}_{\tilde{N}(s_2, s_2 + \delta)}(1) \cdots \mathbf{p}_{\tilde{N}(s_n + \delta, t)}(0).$$

The sum of the lengths of the above intervals is t . Also $\mathbf{p}_{\tilde{N}(s_i, s_i + \delta)}(1) = \lambda \delta \exp(-\lambda \delta)$ for each i , so

$$\Pr\{A(\delta)\} = (\lambda \delta)^n \exp(-\lambda t).$$

The event $A(\delta)$ can be characterized as the event that, first, $N(t) = n$ and, second, that the n arrivals occur in $(s_i, s_i + \delta]$ for $1 \leq i \leq n$. Thus we conclude that

$$f_{\mathbf{S}^{(n)}|N(t)}(\mathbf{s}^{(n)}|n) = \lim_{\delta \rightarrow 0} \frac{\Pr\{A(\delta)\}}{\delta^n \mathbf{p}_{N(t)}(n)},$$

which simplifies to (2.38). \square

The joint density of the interarrival intervals, $\mathbf{X}^{(n)} = (X_1, \dots, X_n)$ given $N(t) = n$ can be found directly from Theorem 2.5.1 simply by making the linear transformation $X_1 = S_1$ and $X_i = S_i - S_{i-1}$ for $2 \leq i \leq n$. The density is unchanged, but the constraint region transforms into $\sum_{i=1}^n X_i < t$ with $X_i > 0$ for $1 \leq i \leq n$ (see Figure 2.10).

$$f_{\mathbf{X}^{(n)}|N(t)}(\mathbf{x}^{(n)}|n) = \frac{n!}{t^n} \quad \text{for } \mathbf{X}^{(n)} > 0, \sum_{i=1}^n X_i < t. \quad (2.40)$$

It is also instructive to compare the joint distribution of $\mathbf{S}^{(n)}$ conditional on $N(t) = n$ with the joint distribution of n IID uniformly distributed random variables, $\mathbf{U}^{(n)} = (U_1, \dots, U_n)$ on $(0, t]$. For any point $\mathbf{U}^{(n)} = \mathbf{u}^{(n)}$, this joint density is

$$f_{\mathbf{U}^{(n)}}(\mathbf{u}^{(n)}) = \frac{1}{t^n} \quad \text{for } 0 < u_i \leq t, 1 \leq i \leq n.$$

Both $f_{\mathbf{S}^{(n)}}$ and $f_{\mathbf{U}^{(n)}}$ are uniform over the volume of n -space where they are non-zero, but as illustrated in Figure 2.11 for $n = 2$, the volume for the latter is $n!$ times larger than the

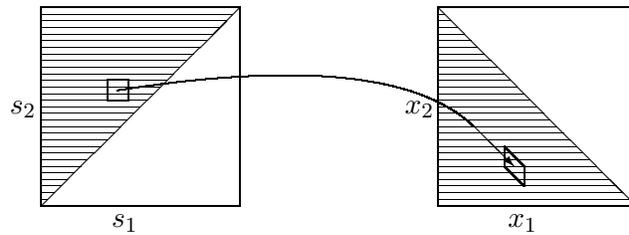


Figure 2.10: Mapping from arrival epochs to interarrival times. Note that incremental cubes in the arrival space map to parallelepipeds of the same volume in the interarrival space.

volume for the former. To explain this more fully, we can define a set of random variables S_1, \dots, S_n , not as arrival epochs in a Poisson process, but rather as the order statistics function of the IID uniform variables U_1, \dots, U_n ; that is

$$S_1 = \min(U_1, \dots, U_n); S_2 = 2^{\text{nd}} \text{ smallest } (U_1, \dots, U_n); \text{ etc.}$$

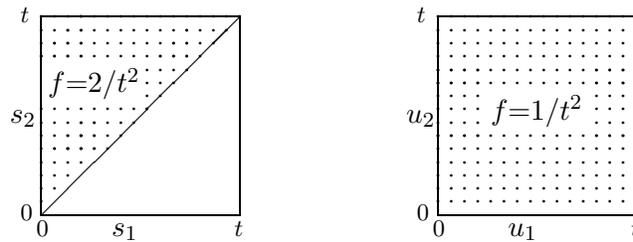


Figure 2.11: Density for the order statistics of an IID 2-dimensional uniform distribution. Note that the square over which $f_{U^{(2)}}$ is non-zero contains one triangle where $u_2 > u_1$ and another of equal size where $u_1 > u_2$. Each of these maps, by a permutation mapping, to the single triangle where $s_2 > s_1$.

The n -cube is partitioned into $n!$ regions, one where $u_1 < u_2 < \dots < u_n$. For each permutation $\pi(i)$ of the integers 1 to n , there is another region¹⁰ where $u_{\pi(1)} < u_{\pi(2)} < \dots < u_{\pi(n)}$. By symmetry, each of these regions has the same volume, which then must be $1/n!$ of the volume t^n of the n -cube.

All of these $n!$ regions map to the same region of ordered values. Thus these order statistics have the same joint probability density function as the arrival epochs S_1, \dots, S_n conditional on $N(t) = n$. Anything we know (or can discover) about order statistics is valid for Poisson arrival epochs given $N(t) = n$ and vice versa.¹¹

¹⁰As usual, we are ignoring those points where $u_i = u_j$ for some i, j , since the set of such points has 0 probability.

¹¹There is certainly also the intuitive notion, given n arrivals in $(0, t]$, and given the stationary and independent increment properties of the Poisson process, that those n arrivals can be viewed as uniformly

Next we want to find the marginal CDF's of the individual S_i , conditional on $N(t) = n$. Starting with S_1 , and viewing it as the minimum of the IID uniformly distributed variables U_1, \dots, U_n , we recognize that $S_1 > \tau$ if and only if $U_i > \tau$ for all i , $1 \leq i \leq n$. Thus,

$$\Pr\{S_1 > \tau \mid N(t)=n\} = \left[\frac{t-\tau}{t}\right]^n \quad \text{for } 0 < \tau \leq t. \quad (2.41)$$

For S_2 to S_n , the density is slightly simpler in appearance than the CDF. To find $f_{S_i|N(t)}(s_i \mid n)$, look at n uniformly distributed rv's in $(0, t]$. The probability that one of these lies in the interval $(s_i, s_i + ds_i]$ is $(n ds_i)/t$. Out of the remaining $n - 1$, the probability that $i - 1$ lie in the interval $(0, s_i]$ is given by the binomial distribution with probability of success s_i/t . Thus the desired density is

$$\begin{aligned} f_{S_i|N(t)}(s_i \mid n) dt &= \frac{s_i^{i-1}(t-s_i)^{n-i}(n-1)!}{t^{n-1}(n-i)!(i-1)!} \frac{n ds_i}{t} \\ f_{S_i|N(t)}(s_i \mid n) &= \frac{s_i^{i-1}(t-s_i)^{n-i}n!}{t^n(n-i)!(i-1)!}. \end{aligned} \quad (2.42)$$

It is easy to find the expected value of S_1 conditional on $N(t) = n$ by integrating the complementary CDF in (2.41), getting

$$\mathbb{E}[S_1 \mid N(t)=n] = \frac{t}{n+1}. \quad (2.43)$$

We come back later to find $\mathbb{E}[S_i \mid N(t) = n]$ for $2 \leq i \leq n$. First, we look at the marginal distributions of the interarrival intervals. Recall from (2.40) that

$$f_{\mathbf{X}^{(n)}|N(t)}(\mathbf{x}^{(n)} \mid n) = \frac{n!}{t^n} \quad \text{for } \mathbf{X}^{(n)} > 0, \sum_{i=1}^n X_i < t. \quad (2.44)$$

The joint density is the same for all points in the constraint region, and the constraint does not distinguish between X_1 to X_n . Thus X_1, \dots, X_n must all have the same marginal distribution, and more generally the marginal distribution of any subset of the X_i can depend only on the size of the subset. We have found the distribution of S_1 , which is the same as X_1 , and thus

$$\Pr\{X_i > \tau \mid N(t)=n\} = \left[\frac{t-\tau}{t}\right]^n \quad \text{for } 1 \leq i \leq n \text{ and } 0 < \tau \leq t. \quad (2.45)$$

$$\mathbb{E}[X_i \mid N(t)=n] = \frac{t}{n+1} \quad \text{for } 1 \leq i \leq n. \quad (2.46)$$

From this, we see immediately that for $1 \leq i \leq n$,

$$\mathbb{E}[S_i \mid N(t) = n] = \frac{it}{n+1}. \quad (2.47)$$

distributed. One way to view this is to visualize the Poisson process as the sum of a very large number k of independent processes of rate λ/k each. Then, given $N(t) = n$, with $k \gg n$, there is negligible probability of more than one arrival from any one process, and for each of the n processes with arrivals, that arrival is uniformly distributed in $(0, t]$.

One could go on and derive joint distributions of all sorts at this point, but there is one additional type of interval that must be discussed. Define $X_{n+1}^* = t - S_n$ to be the interval from the largest arrival epoch before t to t itself. Rewriting (2.44),

$$f_{\mathbf{X}^{(n)}|N(t)}(\mathbf{x}^{(n)} | n) = \frac{n!}{t^n} \quad \text{for } \mathbf{X}^{(n)} > 0, X_{n+1}^* > 0, \sum_{i=1}^n X_i + X_{n+1}^* = t.$$

The constraints above are symmetric in $X_1, \dots, X_n, X_{n+1}^*$, and, within the constraint region, the joint density of X_1, \dots, X_n (conditional on $N(t) = n$) is uniform. Note that there is no joint density over $X_1, \dots, X_n, X_{n+1}^*$ conditional on $N(t) = n$, since X_{n+1}^* is then a deterministic function of X_1, \dots, X_n . However, the density over X_1, \dots, X_n can be replaced by a density over any other n rv's out of $X_1, \dots, X_n, X_{n+1}^*$ by a linear transformation with unit determinant. Thus X_{n+1}^* has the same marginal distribution as each of the X_i . This gives us a partial check on our work, since the interval $(0, t]$ is divided into $n+1$ intervals of sizes $X_1, X_2, \dots, X_n, X_{n+1}^*$, and each of these has a mean size $t/(n+1)$. We also see that the joint CDF of any proper subset of $X_1, X_2, \dots, X_n, X_{n+1}^*$ depends only on the size of the subset and not the particular rv's in the subset.

One important consequence of this is that we can look at the arrivals in a segment $(0, t)$ of a Poisson process either forward or backward¹² in time and they 'look the same.' Looked at backwards, the interarrival intervals are $X_{n+1}^*, X_n, \dots, X_2$. These intervals are IID, and X_1 is then determined as $t - X_{n+1}^* - X_n - \dots - X_2$. We will not make any particular use of this property here, but we will later explore this property of time-reversibility for other types of processes. For Poisson processes, this reversibility is intuitively obvious from the stationary and independent properties. It is less obvious how to express this condition by equations, but that is not really necessary at this point.

2.6 Summary

We started the chapter with three equivalent definitions of a Poisson process—first as a renewal process with exponentially distributed inter-renewal intervals, second as a stationary and independent increment counting process with a Poisson distributed number of arrivals in each interval, and third essentially as a limit of shrinking Bernoulli processes. We saw that each definition provided its own insights into the properties of the process. We emphasized the importance of the memoryless property of the exponential distribution, both as a useful tool in problem solving and as an underlying reason why the Poisson process is so simple.

We next showed that the sum of independent Poisson processes is again a Poisson process. We also showed that if the arrivals in a Poisson process are independently routed to different locations with some fixed probability assignment, then the arrivals at these locations form independent Poisson processes. This ability to view independent Poisson processes either

¹²This must be interpreted carefully, since if we ask about the distribution to the next arrival from some $\tau \in (0, t)$, it is exponential (possibly occurring after t). But if we look for the distribution of the previous arrival before τ , we usually truncate the distribution at 0. Thus the proper interpretation is to look only at the interval $(0, t)$ under the condition of a given number of arrivals in that interval.

independently or as a splitting of a combined process is a powerful technique for finding almost trivial solutions to many problems.

It was next shown that a non-homogeneous Poisson process could be viewed as a (homogeneous) Poisson process on a non-linear time scale. This allows all the properties of (homogeneous) Poisson processes to be applied directly to the non-homogeneous case. The simplest and most useful result from this is (2.35), showing that the number of arrivals in any interval has a Poisson PMF. This result was used to show that the number of customers in service at any given time τ in an M/G/ ∞ queue has a Poisson PMF with a mean approaching λ times the expected service time in the limit as $\tau \rightarrow \infty$.

Finally we looked at the distribution of arrival epochs conditional on n arrivals in the interval $(0, t]$. It was found that these arrival epochs had the same joint distribution as the order statistics of n uniform IID rv's in $(0, t]$. By using symmetry and going back and forth between the uniform variables and the Poisson process arrivals, we found the distribution of the interarrival times, the arrival epochs, and various conditional distributions.

2.7 Exercises

Exercise 2.1. a) Find the Erlang density $f_{S_n}(t)$ by convolving $f_X(x) = \lambda \exp(-\lambda x)$ with itself n times.

b) Find the moment generating function of X (or find the Laplace transform of $f_X(x)$), and use this to find the moment generating function (or Laplace transform) of $S_n = X_1 + X_2 + \dots + X_n$. Invert your result to find $f_{S_n}(t)$.

c) Find the Erlang density by starting with (2.15) and then calculating the marginal density for S_n .

Exercise 2.2. a) Find the mean, variance, and moment generating function of $N(t)$, as given by (2.17).

b) Show by discrete convolution that the sum of two independent Poisson rv's is again Poisson.

c) Show by using the properties of the Poisson process that the sum of two independent Poisson rv's must be Poisson.

Exercise 2.3. The purpose of this exercise is to give an alternate derivation of the Poisson distribution for $N(t)$, the number of arrivals in a Poisson process up to time t . Let λ be the rate of the process.

a) Find the conditional probability $\Pr\{N(t) = n \mid S_n = \tau\}$ for all $\tau \leq t$.

b) Using the Erlang density for S_n , use (a) to find $\Pr\{N(t) = n\}$.

Exercise 2.4. Assume that a counting process $\{N(t); t > 0\}$ has the independent and stationary increment properties and satisfies (2.17) (for all $t > 0$). Let X_1 be the epoch of the first arrival and X_n be the interarrival time between the $n - 1$ st and the n th arrival. Use only these assumptions in doing the following parts of this exercise.

a) Show that $\Pr\{X_1 > x\} = e^{-\lambda x}$.

b) Let S_{n-1} be the epoch of the $n - 1$ st arrival. Show that $\Pr\{X_n > x \mid S_{n-1} = \tau\} = e^{-\lambda x}$.

c) For each $n > 1$, show that $\Pr\{X_n > x\} = e^{-\lambda x}$ and that X_n is independent of S_{n-1} .

d) Argue that X_n is independent of X_1, X_2, \dots, X_{n-1} .

Exercise 2.5. The point of this exercise is to show that the sequence of PMF's for a Bernoulli counting process does not specify the process. In other words, knowing that $N(t)$ satisfies the binomial distribution for all t does not mean that the process is Bernoulli. This helps us understand why the second definition of a Poisson process requires stationary and independent increments along with the Poisson distribution for $N(t)$.

a) Let Y_1, Y_2, Y_3, \dots be a sequence of binary rv's in which each rv is 0 or 1 with equal probability. Find a joint distribution for Y_1, Y_2, Y_3 that satisfies the binomial distribution, $p_{N(t)}(k) = \binom{t}{k} 2^{-t}$ for $t = 1, 2, 3$ and $0 \leq k \leq t$, but for which Y_1, Y_2, Y_3 are not independent.

One simple solution for this contains four 3-tuples with probability $1/8$ each, two 3-tuples with probability $1/4$ each, and two 3-tuples with probability 0. Note that by making the subsequent arrivals IID and equiprobable, you have an example where $N(t)$ is binomial for all t but the process is not Bernoulli. Hint: Use the binomial for $t = 3$ to find two 3-tuples that must have probability $1/8$. Combine this with the binomial for $t = 2$ to find two other 3-tuples that must have probability $1/8$. Finally look at the constraints imposed by the binomial distribution on the remaining four 3-tuples.

b) Generalize (a) to the case where Y_1, Y_2, Y_3 satisfy $\Pr\{Y_i = 1\} = p$ and $\Pr\{Y_i = 0\} = 1 - p$. Assume $p < 1/2$ and find a joint distribution on Y_1, Y_2, Y_3 that satisfies the binomial distribution, but for which the 3-tuple $(0, 1, 1)$ has zero probability.

c) More generally yet, view a joint PMF on binary t -tuples as a nonnegative vector in a 2^t dimensional vector space. Each binomial probability $p_{N(\tau)}(k) = \binom{\tau}{k} q^k (1-q)^{\tau-k}$ constitutes a linear constraint on this vector. For each τ , show that one of these constraints may be replaced by the constraint that the components of the vector sum to 1.

d) Using (c), show that at most $(t+1)t/2 + 1$ of the binomial constraints are linearly independent. Note that this means that the linear space of vectors satisfying these binomial constraints has dimension at least $2^t - (t+1)t/2 - 1$. This linear space has dimension 1 for $t = 3$ explaining the results in parts a) and b). It has a rapidly increasing dimension for $t > 3$, suggesting that the binomial constraints are relatively ineffectual for constraining the joint PMF of a joint distribution. More work is required for the case of $t > 3$ because of all the inequality constraints, but it turns out that this large dimensionality remains.

Exercise 2.6. Let $h(x)$ be a positive function of a real variable that satisfies $h(x+t) = h(x) + h(t)$ for all x, t , and let $h(1) = c$.

a) Show that for integer $k > 0$, $h(k) = kc$.

b) Show that for integer $j > 0$, $h(1/j) = c/j$.

c) Show that for all integer k, j , $h(k/j) = ck/j$.

d) The above parts show that $h(x)$ is linear in positive *rational* numbers. For very picky mathematicians, this does not guarantee that $h(x)$ is linear in positive *real* numbers. Show that if $h(x)$ is also monotonic in x , then $h(x)$ is linear in $x > 0$.

Exercise 2.7. Assume that a counting process $\{N(t); t > 0\}$ has the independent and stationary increment properties and, for all $t > 0$, satisfies

$$\begin{aligned} \Pr\{\tilde{N}(t, t+\delta) = 0\} &= 1 - \lambda\delta + o(\delta) \\ \Pr\{\tilde{N}(t, t+\delta) = 1\} &= \lambda\delta + o(\delta) \\ \Pr\{\tilde{N}(t, t+\delta) > 1\} &= o(\delta). \end{aligned}$$

a) Let $F_1^c(\tau) = \Pr\{N(\tau) = 0\}$ and show that $dF_1^c(\tau)/d\tau = -\lambda F_1^c(\tau)$.

b) Show that X_1 , the time of the first arrival, is exponential with parameter λ .

- c) Let $F_n^c(\tau) = \Pr\{\tilde{N}(t, t + \tau) = 0 \mid S_{n-1} = t\}$ and show that $dF_n^c(\tau)/d\tau = -\lambda F_n^c(\tau)$.
- d) Argue that X_n is exponential with parameter λ and independent of earlier arrival times.

Exercise 2.8. For a Poisson process, let $t > 0$ be arbitrary and let Z_1 be the duration of the interval from t until the next arrival after t . Let Z_m , for each $m > 1$, be the interarrival time from the epoch of the $m - 1^{\text{st}}$ arrival after t until the m th arrival.

- a) Given that $N(t) = n$, explain why $Z_m = X_{m+n}$ for $m > 1$ and $Z_1 = X_{n+1} - t + S_n$.
- b) Conditional on $N(t) = n$ and $S_n = \tau$, show that Z_1, Z_2, \dots are IID.
- c) Show that Z_1, Z_2, \dots are IID.

Exercise 2.9. Consider a “shrinking Bernoulli” approximation $N_\delta(m\delta) = Y_1 + \dots + Y_m$ to a Poisson process as described in Subsection 2.2.5.

- a) Show that

$$\Pr\{N_\delta(m\delta) = n\} = \binom{m}{n} (\lambda\delta)^n (1 - \lambda\delta)^{m-n}.$$

- b) Let $t = m\delta$, and let t be fixed for the remainder of the exercise. Explain why

$$\lim_{\delta \rightarrow 0} \Pr\{N_\delta(t) = n\} = \lim_{m \rightarrow \infty} \binom{m}{n} \left(\frac{\lambda t}{m}\right)^n \left(1 - \frac{\lambda t}{m}\right)^{m-n},$$

where the limit on the left is taken over values of δ that divide t .

- c) Derive the following two equalities:

$$\lim_{m \rightarrow \infty} \binom{m}{n} \frac{1}{m^n} = \frac{1}{n!}; \quad \text{and} \quad \lim_{m \rightarrow \infty} \left(1 - \frac{\lambda t}{m}\right)^{m-n} = e^{-\lambda t}.$$

- d) Conclude from this that for every t and every n , $\lim_{\delta \rightarrow 0} \Pr\{N_\delta(t) = n\} = \Pr\{N(t) = n\}$ where $\{N(t); t > 0\}$ is a Poisson process of rate λ .

Exercise 2.10. Let $\{N(t); t > 0\}$ be a Poisson process of rate λ .

- a) Find the joint probability mass function (PMF) of $N(t), N(t + s)$ for $s > 0$.
- b) Find $E[N(t) \cdot N(t + s)]$ for $s > 0$.
- c) Find $E[\tilde{N}(t_1, t_3) \cdot \tilde{N}(t_2, t_4)]$ where $\tilde{N}(t, \tau)$ is the number of arrivals in $(t, \tau]$ and $t_1 < t_2 < t_3 < t_4$.

Exercise 2.11. An elementary experiment is independently performed N times where N is a Poisson rv of mean λ . Let $\{a_1, a_2, \dots, a_K\}$ be the set of sample points of the elementary experiment and let p_k , $1 \leq k \leq K$, denote the probability of a_k .

- a) Let N_k denote the number of elementary experiments performed for which the output is a_k . Find the PMF for N_k ($1 \leq k \leq K$). (Hint: no calculation is necessary.)
- b) Find the PMF for $N_1 + N_2$.
- c) Find the conditional PMF for N_1 given that $N = n$.
- d) Find the conditional PMF for $N_1 + N_2$ given that $N = n$.
- e) Find the conditional PMF for N given that $N_1 = n_1$.

Exercise 2.12. Starting from time 0, northbound buses arrive at 77 Mass. Avenue according to a Poisson process of rate λ . Customers arrive according to an independent Poisson process of rate μ . When a bus arrives, all waiting customers instantly enter the bus and subsequent customers wait for the next bus.

- a) Find the PMF for the number of customers entering a bus (more specifically, for any given m , find the PMF for the number of customers entering the m th bus).
- b) Find the PMF for the number of customers entering the m th bus given that the inter-arrival interval between bus $m - 1$ and bus m is x .
- c) Given that a bus arrives at time 10:30 PM, find the PMF for the number of customers entering the next bus.
- d) Given that a bus arrives at 10:30 PM and no bus arrives between 10:30 and 11, find the PMF for the number of customers on the next bus.
- e) Find the PMF for the number of customers waiting at some given time, say 2:30 PM (assume that the processes started infinitely far in the past). Hint: think of what happens moving backward in time from 2:30 PM.
- f) Find the PMF for the number of customers getting on the next bus to arrive after 2:30. Hint: this is different from (a); look carefully at (e).
- g) Given that I arrive to wait for a bus at 2:30 PM, find the PMF for the number of customers getting on the next bus.

Exercise 2.13. a) Show that the arrival epochs of a Poisson process satisfy

$$f_{S^{(n)}|S_{n+1}}(\mathbf{s}^{(n)}|s_{n+1}) = n!/s_{n+1}^n.$$

Hint: This is easy if you use only the results of Section 2.2.2.

- b) Contrast this with the result of Theorem 2.5.1

Exercise 2.14. Equation (2.42) gives $f_{S_i|N(t)}(s_i | n)$, which is the density of random variable S_i conditional on $N(t) = n$ for $n \geq i$. Multiply this expression by $\Pr\{N(t) = n\}$ and sum over n to find $f_{S_i}(s_i)$; verify that your answer is indeed the Erlang density.

Exercise 2.15. Consider generalizing the bulk arrival process in Figure 2.5. Assume that the epochs at which arrivals occur form a Poisson process $\{N(t); t > 0\}$ of rate λ . At each arrival epoch, S_n , the number of arrivals, Z_n , satisfies $\Pr\{Z_n=1\} = p$, $\Pr\{Z_n=2\} = 1 - p$. The variables Z_n are IID.

a) Let $\{N_1(t); t > 0\}$ be the counting process of the epochs at which single arrivals occur. Find the PMF of $N_1(t)$ as a function of t . Similarly, let $\{N_2(t); t \geq 0\}$ be the counting process of the epochs at which double arrivals occur. Find the PMF of $N_2(t)$ as a function of t .

b) Let $\{N_B(t); t \geq 0\}$ be the counting process of the total number of arrivals. Give an expression for the PMF of $N_B(t)$ as a function of t .

Exercise 2.16. a) For a Poisson counting process of rate λ , find the joint probability density of S_1, S_2, \dots, S_{n-1} conditional on $S_n = t$.

b) Find $\Pr\{X_1 > \tau \mid S_n=t\}$.

c) Find $\Pr\{X_i > \tau \mid S_n=t\}$ for $1 \leq i \leq n$.

d) Find the density $f_{S_i|S_n}(s_i|t)$ for $1 \leq i \leq n - 1$.

e) Give an explanation for the striking similarity between the condition $N(t) = n - 1$ and the condition $S_n = t$.

Exercise 2.17. a) For a Poisson process of rate λ , find $\Pr\{N(t)=n \mid S_1=\tau\}$ for $t > \tau$ and $n \geq 1$.

b) Using this, find $f_{S_1}(\tau \mid N(t)=n)$

c) Check your answer against (2.41).

Exercise 2.18. Consider a counting process in which the rate is a rv Λ with probability density $f_\Lambda(\lambda) = \alpha e^{-\alpha\lambda}$ for $\lambda > 0$. Conditional on a given sample value λ for the rate, the counting process is a Poisson process of rate λ (i.e., nature first chooses a sample value λ and then generates a sample path of a Poisson process of that rate λ).

a) What is $\Pr\{N(t)=n \mid \Lambda=\lambda\}$, where $N(t)$ is the number of arrivals in the interval $(0, t]$ for some given $t > 0$?

b) Show that $\Pr\{N(t)=n\}$, the unconditional PMF for $N(t)$, is given by

$$\Pr\{N(t)=n\} = \frac{\alpha t^n}{(t + \alpha)^{n+1}}.$$

c) Find $f_\Lambda(\lambda \mid N(t)=n)$, the density of λ conditional on $N(t)=n$.

d) Find $E[\Lambda \mid N(t)=n]$ and interpret your result for very small t with $n = 0$ and for very large t with n large.

e) Find $E[\Lambda \mid N(t)=n, S_1, S_2, \dots, S_n]$. (Hint: consider the distribution of S_1, \dots, S_n conditional on $N(t)$ and Λ). Find $E[\Lambda \mid N(t)=n, N(\tau)=m]$ for some $\tau < t$.

Exercise 2.19. a) Use Equation (2.42) to find $E[S_i | N(t)=n]$. Hint: When you integrate $s_i f_{S_i}(s_i | N(t)=n)$, compare this integral with $f_{S_{i+1}}(s_i | N(t)=n+1)$ and use the fact that the latter expression is a probability density.

b) Find the second moment and the variance of S_i conditional on $N(t)=n$. Hint: Extend the previous hint.

c) Assume that n is odd, and consider $i = (n+1)/2$. What is the relationship between S_i , conditional on $N(t)=n$, and the sample median of n IID uniform random variables.

d) Give a weak law of large numbers for the above median.

Exercise 2.20. Suppose cars enter a one-way infinite length, infinite lane highway at a Poisson rate λ . The i th car to enter chooses a velocity V_i and travels at this velocity. Assume that the V_i 's are independent positive rv's having a common distribution F . Derive the distribution of the number of cars that are located in an interval $(0, a)$ at time t .

Exercise 2.21. Consider an M/G/ ∞ queue, i.e., a queue with Poisson arrivals of rate λ in which each arrival i , independent of other arrivals, remains in the system for a time X_i , where $\{X_i; i \geq 1\}$ is a set of IID rv's with some given CDF $F(x)$.

You may assume that the number of arrivals in any interval $(t, t + \epsilon)$ that are still in the system at some later time $\tau \geq t + \epsilon$ is statistically independent of the number of arrivals in that same interval $(t, t + \epsilon)$ that have departed from the system by time τ .

a) Let $N(\tau)$ be the number of customers in the system at time τ . Find the mean, $m(\tau)$, of $N(\tau)$ and find $\Pr\{N(\tau) = n\}$.

b) Let $D(\tau)$ be the number of customers that have departed from the system by time τ . Find the mean, $E[D(\tau)]$, and find $\Pr\{D(\tau) = d\}$.

c) Find $\Pr\{N(\tau) = n, D(\tau) = d\}$.

d) Let $A(\tau)$ be the total number of arrivals up to time τ . Find $\Pr\{N(\tau) = n | A(\tau) = a\}$.

e) Find $\Pr\{D(\tau + \epsilon) - D(\tau) = d\}$.

Exercise 2.22. The voters in a given town arrive at the place of voting according to a Poisson process of rate $\lambda = 100$ voters per hour. The voters independently vote for candidate A and candidate B each with probability $1/2$. Assume that the voting starts at time 0 and continues indefinitely.

a) Conditional on 1000 voters arriving during the first 10 hours of voting, find the probability that candidate A receives n of those votes.

b) Again conditional on 1000 voters during the first 10 hours, find the probability that candidate A receives n votes in the first 4 hours of voting.

c) Let T be the epoch of the arrival of the first voter voting for candidate A . Find the density of T .

d) Find the PMF of the number of voters for candidate B who arrive before the first voter for A .

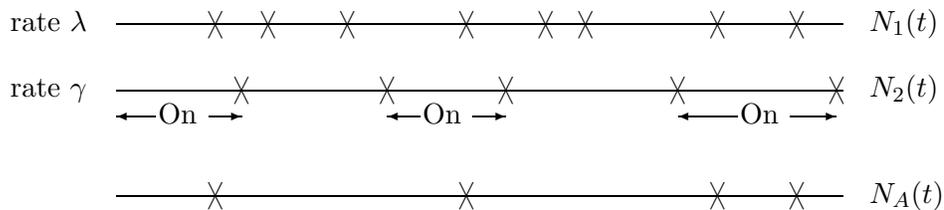
e) Define the n th voter as a *reversal* if the n th voter votes for a different candidate than the $n - 1$ st. For example, in the sequence of votes $AABAABB$, the third, fourth, and sixth voters are reversals; the third and sixth are A to B reversals and the fourth is a B to A reversal. Let $N(t)$ be the number of reversals up to time t (t in hours). Is $\{N(t); t > 0\}$ a Poisson process? Explain.

f) Find the expected time (in hours) between reversals.

g) Find the probability density of the time between reversals.

h) Find the density of the time from one A to B reversal to the next A to B reversal.

Exercise 2.23. Let $\{N_1(t); t > 0\}$ be a Poisson counting process of rate λ . Assume that the arrivals from this process are switched on and off by arrivals from a second independent Poisson process $\{N_2(t); t > 0\}$ of rate γ .



Let $\{N_A(t); t \geq 0\}$ be the switched process; that is $N_A(t)$ includes the arrivals from $\{N_1(t); t > 0\}$ during periods when $N_2(t)$ is even and excludes the arrivals from $\{N_1(t); t > 0\}$ while $N_2(t)$ is odd.

a) Find the PMF for the number of arrivals of the first process, $\{N_1(t); t > 0\}$, during the n th period when the switch is on.

b) Given that the first arrival for the second process occurs at epoch τ , find the conditional PMF for the number of arrivals of the first process up to τ .

c) Given that the number of arrivals of the first process, up to the first arrival for the second process, is n , find the density for the epoch of the first arrival from the second process.

d) Find the density of the interarrival time for $\{N_A(t); t \geq 0\}$. Note: This part is quite messy and is done most easily via Laplace transforms.

Exercise 2.24. Let us model the chess tournament between Fisher and Spassky as a stochastic process. Let X_i , for $i \geq 1$, be the duration of the i th game and assume that $\{X_i; i \geq 1\}$ is a set of IID exponentially distributed rv's each with density $f_X(x) = \lambda e^{-\lambda x}$. Suppose that each game (independently of all other games, and independently of the length of the games) is won by Fisher with probability p , by Spassky with probability q , and is a draw with probability $1 - p - q$. The first player to win n games is defined to be the winner, but we consider the match up to the point of winning as being embedded in an unending sequence of games.

a) Find the distribution of time, from the beginning of the match, until the completion of the first game that is won (i.e., that is not a draw). Characterize the process of the number $\{N(t); t > 0\}$ of games won up to and including time t . Characterize the process of the number $\{N_F(t); t \geq 0\}$ of games won by Fisher and the number $\{N_S(t); t \geq 0\}$ won by Spassky.

b) For the remainder of the problem, assume that the probability of a draw is zero; i.e., that $p + q = 1$. How many of the first $2n - 1$ games must be won by Fisher in order to win the match?

c) What is the probability that Fisher wins the match? Your answer should not involve any integrals. Hint: consider the unending sequence of games and use (b).

d) Let T be the epoch at which the match is completed (i.e., either Fisher or Spassky wins). Find the CDF of T .

e) Find the probability that Fisher wins and that T lies in the interval $(t, t + \delta)$ for arbitrarily small δ .

Exercise 2.25. a) Find the conditional density of S_{i+1} , conditional on $N(t) = n$ and $S_i = s_i$.

b) Use part a) to find the joint density of S_1, \dots, S_n conditional on $N(t) = n$. Verify that your answer agrees with (2.38).

Exercise 2.26. A two-dimensional Poisson process is a process of randomly occurring special points in the plane such that (i) for any region of area A the number of special points in that region has a Poisson distribution with mean λA , and (ii) the number of special points in nonoverlapping regions is independent. For such a process consider an arbitrary location in the plane and let X denote its distance from its nearest special point (where distance is measured in the usual Euclidean manner). Show that

a) $\Pr\{X > t\} = \exp(-\lambda\pi t^2)$

b) $E[X] = 1/(2\sqrt{\lambda})$.

Exercise 2.27. This problem is intended to show that one can analyze the long term behavior of queueing problems by using just notions of means and variances, but that such analysis is awkward, justifying understanding the strong law of large numbers. Consider an M/G/1 queue. The arrival process is Poisson with $\lambda = 1$. The expected service time, $E[Y]$, is $1/2$ and the variance of the service time is given to be 1.

a) Consider S_n , the time of the n th arrival, for $n = 10^{12}$. With high probability, S_n will lie within 3 standard deviations of its mean. Find and compare this mean and the 3σ range.

b) Let V_n be the total amount of time during which the server is busy with these n arrivals (i.e., the sum of 10^{12} service times). Find the mean and 3σ range of V_n .

c) Find the mean and 3σ range of I_n , the total amount of time the server is idle up until S_n (take I_n as $S_n - V_n$, thus ignoring any service time after S_n).

- d)** An idle period starts when the server completes a service and there are no waiting arrivals; it ends on the next arrival. Find the mean and variance of an idle period. Are successive idle periods IID?
- e)** Combine (c) and (d) to estimate the total number of idle periods up to time S_n . Use this to estimate the total number of busy periods.
- f)** Combine (e) and (b) to estimate the expected length of a busy period.

Exercise 2.28. The purpose of this problem is to illustrate that for an arrival process with independent but not identically distributed interarrival intervals, X_1, X_2, \dots the number of arrivals $N(t)$ in the interval $(0, t]$ can be a defective rv. In other words, the ‘counting process’ is not necessarily a stochastic process according to our definitions. This also suggests that it might be necessary to prove that the counting rv’s for a renewal process are actually rv’s.

- a)** Let the CDF of the i th interarrival interval for an arrival process be $F_{X_i}(x_i) = 1 - \exp(-\alpha^{-i}x_i)$ for some fixed $\alpha \in (0, 1)$. Let $S_n = X_1 + \dots + X_n$ and show that

$$E[S_n] = \frac{\alpha(1 - \alpha^n)}{1 - \alpha}.$$

- b)** Sketch a ‘reasonable’ sample path for $N(t)$.
- c)** Use the Markov inequality on $\Pr\{S_n \geq t\}$ to find an upper bound on $\Pr\{N(t) < n\}$ that is smaller than 1 for all n and for large enough t . Use this to show that $N(t)$ is defective for large enough t .
- d)** (For those looking for a challenge) Show that $N(t)$ is defective for all $t > 0$. Hint: Use the Markov inequality to find an upper bound on $\Pr\{S_m - S_n \leq t/2\}$ for all $m > n$ for any fixed n . Show that, for any $t > 0$, this is bounded below 1 for large enough n . Then show that S_n has a density that is positive for all $t > 0$.