7

The Poisson process

The Poisson process is a prominent stochastic process, mainly because it frequently appears in a wealth of physical phenomena and because it is relatively simple to analyze. Therefore, we will first treat the Poisson process before considering the more general Markov processes.

7.1 A stochastic process

7.1.1 Introduction and definitions

A *stochastic*[1] *process*, formally denoted as $\{X(t), t \in T\}$, is a sequence of random variables $X(t)$, where the parameter t – most often the time – runs over an index set T. The *state space* of the stochastic process is the set of all possible values for the random variables $X(t)$ and each of these possible values is called the *state* of the process. If the index set T is a countable set, $X[k]$ is a discrete stochastic process. Often k is the discrete time or a time slot in computer systems. If T is a continuum, $X(t)$ is a continuous stochastic process. For example, the outcome of n tosses of a coin is a discrete stochastic process with state space {heads, tails} and the index set $T = \{0, 1, 2, \ldots, n\}$. The number of arrivals of packets in a router during a certain time interval $[a, b]$ is a continuous stochastic process because $t \in [a, b]$. Any realization of a stochastic process is called a *sample path.* For example, a sample path of the outcome of n tosses of a coin is {heads, tails, tails, ..., heads}, while a sample path of the number of arrivals in $[a, b]$ is $1_{a \le t < a+h}, 3 \times 1_{a+h \le t < a+4h}, 8 \times 1_{a+4h \le t < a+5h}, \ldots, 13 \times 1_{a+(k-1)h \le t < b}$, where $h = \frac{b-a}{k}$. Other examples are the measurement of the temperature each day, the notation of the value of a stock each minute or rolling a die and recording its value, which is illustrated in Fig. 7.1.

---

[1] The word "stochastic" is derived from $\sigma\tau o\chi\alpha\zeta\epsilon\sigma\theta\alpha i$ in Greek which means "to aim at, try to hit".

Especially in continuous stochastic processes, it is convenient to define increments as the difference $X(t) - X(u)$. The continuous time stochastic process $X(t)$ has *independent increments* if changes in the value of the process in different time intervals are independent, or, if for all $t_0 < t_1 < \cdots < t_n$, the random variables $X(t_1) - X(t_0), X(t_2) - X(t_1), \ldots, X(t_n) - X(t_{n-1})$ are independent. The continuous (time) stochastic process has *stationary increments* if $X(t+s) - X(s)$ possesses the same distribution for all s. Hence, changes in the value of the process only dependent on the distance t between process events, not on the time point s.

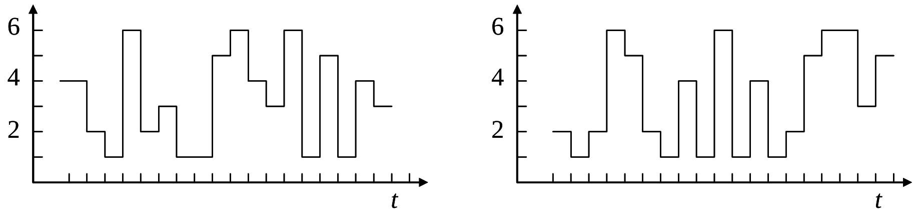

Fig. 7.1. Two different sample paths of the experiment: roll a die and record the outcome. The total number of different sample paths is 6^T where T is the number of times an outcome is recorded. The state space only contains 6 possible outcomes $\{1, 2, 3, 4, 5, 6\}$.

Stochastic processes are distinguished by (a) their state space, (b) the index set T and (c) by the dependence relations between random variables $X(t)$. For example, a standard Brownian motion (or Wiener process)[2] is defined as a stochastic process $X(t)$ having continuous sample paths, stationary independent increments and $X(t)$ has a normal distribution $N(0, t)$. A Poisson process, defined in more detail in Section 7.2, is a stochastic process $X(t)$ having discontinuous sample paths, stationary independent increments and $X(t)$ has a Poisson distribution. A generalization of the Poisson process is a counting process. A counting process is defined as a stochastic process $N(t) \geq 0$ with discontinuous sample paths, stationary independent increments, but with arbitrary distribution. A counting process $N(t)$ represents the total number of events that have occurred in a time interval $[0, t]$. Examples of a counting process are the number of telephone calls at a local exchange during an interval, the number of failures in a telecommunication network, the number of corrupted bits after transmission due to channel errors, etc.

---

[2] Harrison (1990) shows that the converse is also true: if Y is a continuous process with stationary independent increments, then Y is a Brownian motion.

7.1.2 Modeling a stochastic process from measurements

In practice, understanding observed phenomena often asks for a stochastic model that captures the main characteristics of the studied phenomena and that enables computations of diverse quantities of interest. Examples in the field of data communications networks are the determination of the arrival process at a switch or router in order to dimension the number of buffer (memory) places, the modeling of the graph of the Internet, the distribution of the duration of a telephone call or web browsing session, the number of visits to certain websites, the number of links that refer to a web page, the amount of downloaded information, the number of traversed routers by an email, etc. Accurate modeling is in general difficult and often trades off complexity against accuracy of the model.

Let us illustrate some aspects of *modeling* by considering Internet delay measurements. A motivation for obtaining an end-to-end delay model for (a part of) the Internet is the question whether massive service deployment of voice over IP (VoIP) can substitute classical telephony with a comparable quality. Specifically, classical telephony requires that the end-to-end delay of an arbitrary telephone conversation hardly exceeds 100 ms.

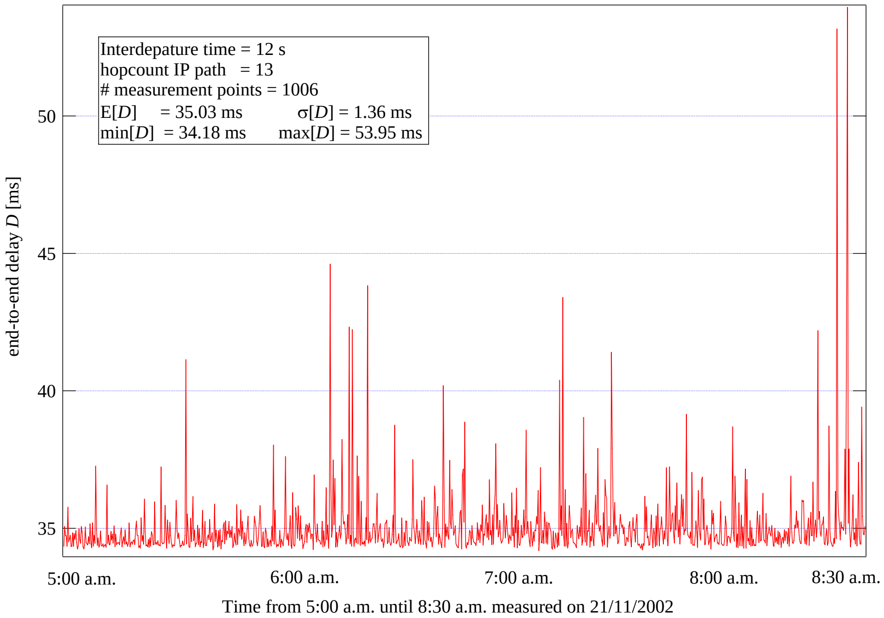

Fig. 7.2. The raw data of the end-to-end delay of IP test packets along a same path of 13 hops in the Internet measured during 3.5 hours.

The end-to-end delay along a fixed path between source and destination measured during some interval is an example of a continuous time stochastic process. We have received data of the delay measured at RIPE-NCC as illustrated in Fig. 7.2. Figure 7.2 shows a *sample path* of this continuous stochastic process. The precise details of the measurement configuration are for the present purpose not relevant. It suffices to add that Figure 7.2 shows the time difference between the departure of an IP test packet of 100 byte at the sending box and its arrival at the destination box accurate within 10 μs. The average sending rate of IP test packets is $\frac{1}{12}$ packets per second. Each IP test packet is assumed to follow the same path from sending to receiving box. The steadiness of the path is checked by trace-route measurements every 6 minutes.

Usually, in the next step, the histogram of the raw data is made. A histogram counts the number of data points that lie in an interval of ΔD ms, which is often called the bin size. Most graphical packages allow to choose the bin size. Figure 7.3 shows two different histograms with bin size $\Delta D = 0.5$ ms and $\Delta D = 0.1$ ms. In general, there is no universal rule to choose the bin size ΔD. Clearly, the bin size is bounded below by the measurement accuracy, in our case $\Delta D > 10$ μs. A finer bin size provides more detail, but the resulting histogram exhibits also more stochastic variations because there are fewer data points in a small bin and adjacent bins may possess a significantly different amount of data points. Hence, compared to one larger bin that covers a same interval, less averaging or smoothing occurs in a set of smaller bins. The normalized histogram obtained by dividing the counts per bin by the total number of data points provides a first approximation to the probability density function of D. However, it is still discrete and approximates $\Pr[k < D \leq k + \Delta D]$. A more precise description of constructing a histogram is given in Section C.1.

The histogram is generally better suited to decide whether outliners in de data points may be due to measurement errors or not. Figure 7.3 suggests to either neglect the data points with $D > 40$ ms or to measure at a higher sending rate of IP test packets in order to have more details in the intervals exceeding 38 or 40 ms. If there existed a good[3] stochastic model for the end-to-end delay along fixed Internet paths, a normal procedure[4] in engineering and physics would be to fit the histogram with that stochastic model to obtain the parameters of that stochastic model. The accuracy of the fit can

---

[3] Which is still lacking at the time of writing.

[4] Other more difficult methods in the realm of statistics must be invoked in case the measurement data are so precious and rare that any additional measurement point has a far larger cost than the cost of extensive additional computations.

be expressed in terms of the correlation coefficient ρ explained in Section 2.5.3. The closer ρ tends to 1, the better the fit, which gives confidence that the stochastic model corresponds with the real phenomenon.

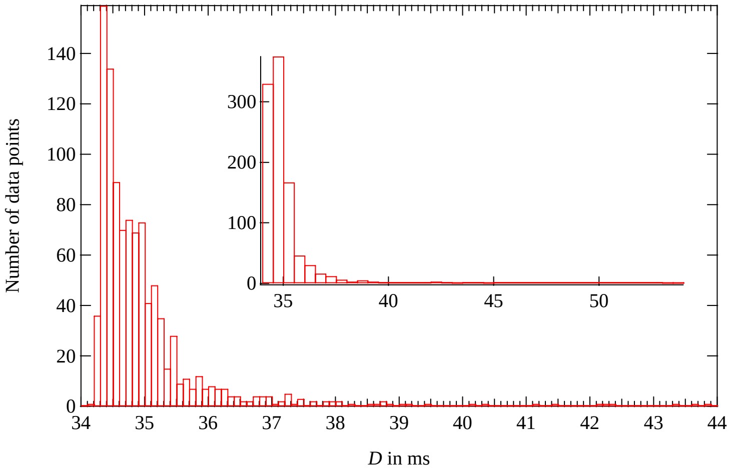

Fig. 7.3. The histogram of the end-to-end delay with a bin size of 0.1 ms (the insert has bin size of 0.5 ms).

Assuming that the presented measurement is a typical measurement along a fixed Internet path (which is true for about 80% of the investigated different paths), it demonstrates that there is a clear minimum at about 34 ms due to the propagation delay of electromagnetic waves. In addition, the end-to-end delay lies for 99% between 34 and 38 ms. However, there is insufficient data to pronounce claims in the tail behavior ($\Pr[D > x]$ for $x > 40$ ms). Just this region is of interest to compute the quality of service expressed as the probability that the end-to-end delay exceeds x ms is smaller than 10^{-a} where a specifies the stringency on the quality requirement. Toll quality in classical telephony sets x at 100 ms and a in the range of 4 to 5. The existence of a good stochastic model covering the whole possible range of the end-to-end delay D would enable us to compute tail probabilities based on the parameters that can be fitted from the measurements.

The histogram is in fact a projection of the raw measurement data onto the ordinate (end-to-end delay axis). All time information (the abscissa in Fig. 7.2) is lost. Usually, the time evolution and the dependencies or correlations over time of a stochastic phenomenon are difficult and most analyses are only tractable under certain simplifying conditions. For example, often only a steady state analysis is possible and the increments $X(t_k) - X(t_{k-1})$

of the process for all $t_0 < \cdots < t_{k-1} < t_k < \cdots < t_n$ are assumed to be independent or weakly dependent. The study of Markov processes (Chapters 9–11) basically tries to compute and analyze the process in steady state. Figure 7.2 is measured over a relatively long period of time and indicates that after 8.00 a.m. the background traffic increases. The background traffic interferes with the IP test packets and causes them to queue longer in routers such that larger variations are observed. However, it is in general difficult to ascertain that (a part of) the measurement is performed while the system operates in a certain stable regime (or steady state).

We have touched upon some aspects in the art of modeling to motivate the importance of studying stochastic processes. In the sequel of this chapter, one of the most basic and simplest stochastic processes is investigated.

7.2 The Poisson process

A Poisson process with parameter or rate $\lambda > 0$ is an integer-valued, continuous time stochastic process $\{X(t), t \geq 0\}$ satisfying

(i) $X(0) = 0$

(ii) for all $t_0 = 0 < t_1 < \cdots < t_n$, the increments $X(t_1) - X(t_0), X(t_2) - X(t_1), \ldots, X(t_n) - X(t_{n-1})$ are independent random variables

(iii) for $t \geq 0$, $s > 0$ and non-negative integers k, the increments have the Poisson distribution

$$\Pr\left[X(t+s) - X(s) = k\right] = \frac{(\lambda t)^k e^{-\lambda t}}{k!} \tag{7.1}$$

It is convenient to view the Poisson process $X(t)$ as a special counting process, where the number of events in any interval of length t is specified via condition (iii). From this definition, a number of properties can be derived:

(a) Condition (iii) implies that the increments are stationary because the right-hand side does not dependent on s. In other words, the increments only depend on the length of the interval t and not on the time s when the interval begins. Further, with (3.11), the mean $E\left[X(t+s) - X(s)\right] = \lambda t$ and because the increments are stationary, this holds for any value of s. In particular with $s = 0$ and condition 1, the expected number of events in a time interval with length t is

$$E\left[X(t)\right] = \lambda t \tag{7.2}$$

Relation (7.2) explains why λ is called the rate of the Poisson process, namely, the derivative over time t or the number of events per time unit.

(b) The probability that exactly one event occurs in an arbitrarily small time interval of length h follows from condition (iii) as

$$\Pr[X(h+s)-X(s)=1]=\lambda h e^{-\lambda h}=\lambda h+o(h)$$

while the probability that no event occurs in an arbitrarily small time interval of length h is

$$\Pr[X(h+s)-X(s)=0]=e^{-\lambda h}=1-\lambda h+o(h)$$

Similarly, the probability that more than one event occurs in an arbitrarily small time interval of length h is

$$\Pr[X(h+s)-X(s)>1]=o(h)$$

Example 1 A conversation in a wireless ad-hoc network is severely disturbed by interference signals according to a Poisson process of rate $\lambda=0.1$ per minute. (a) What is the probability that no interference signals occur within the first two minutes of the conversation? (b) Given that the first two minutes are free of disturbing effects, what is the probability that in the next minute precisely 1 interfering signal disturbs the conversation?

(a) Let $X(t)$ denote the Poisson interference process, then $\Pr[X(2)=0]$ needs to be computed. Since $X(0)=0$ and with (7.1), we can write $\Pr[X(2)=0]=\Pr[X(2)-X(0)=0]=e^{-2\lambda}$, which equals $\Pr[X(2)=0]=e^{-0.2}=0.8187$.

(b) The events during two non-overlapping intervals of a Poisson process are independent. Thus the event $\{X(2)-X(0)=0\}$ is independent from the event $\{X(3)-X(2)=1\}$ which means that the asked conditional probability $\Pr[X(3)-X(2)=1|X(2)-X(0)=0]=\Pr[X(3)-X(2)=1]$. From (7.1), we obtain $\Pr[X(3)-X(2)=1]=0.1e^{-0.1}=0.0905$.

Example 2 During a certain time interval $[t_1,t_1+10\text{ s}]$, the number of IP packets that arrive at a router is on average 40/s. A service provider asks us to compute the probability that there arrive 20 packets in the period $[t_1,t_1+1\text{ s}]$ and 30 IP packets in $[t_1,t_1+3\text{ s}]$. We may regard the arrival process as a Poisson process.

We are asked to compute $\Pr[X(1)=20,X(3)=30]$ knowing that $\lambda=40$ s^{-1}. Using the independence of increments and (7.1), we rewrite

$$\begin{aligned}\Pr[X(1)=20,X(3)=30]&=\Pr[X(1)-X(0)=20,X(3)-X(1)=10]\\&=\Pr[X(1)-X(0)=20]\Pr[X(3)-X(1)=10]\\&=\frac{(\lambda)^{20}e^{-\lambda}}{20!}\frac{(2\lambda)^{10}e^{-2\lambda}}{10!}=10^{-26}\approx 0\end{aligned}$$

which means that the request of the service provider does not occur in practice.

7.3 Properties of the Poisson process

The first theorem is the converse of the above property (b) that immediately followed from the definition. The Theorems presented here reveal the methodology of how stochastic processes are studied.

Theorem 7.3.1 *A counting process* $N(t)$ *that satisfies the conditions (i)* $N(0) = 0$*, (ii) the process* $N(t)$ *has stationary and independent increments, (iii)* $\Pr[N(h) = 1] = \lambda h + o(h)$ *and (iv)* $\Pr[N(h) > 1] = o(h)$ *is a Poisson process with rate* $\lambda > 0$.

Proof: We must show that conditions (iii) and (iv) are equivalent to condition (iii) in the definition of the Poisson process. Denote $P_n(t) = \Pr[N(t) = n]$ and consider first the case $n = 0$, then

$$P_0(t+h) = \Pr[N(t+h) = 0] = \Pr[N(t+h) - N(t) = 0, N(t) = 0]$$

Invoking independence via (ii)

$$P_0(t+h) = \Pr[N(t+h) - N(t) = 0] \Pr[N(t) = 0]$$

By definition, $P_0(t) = \Pr[N(t) = 0]$ and from (iii), (iv) and the fact that $\sum_{k=0}^{\infty} \Pr[N(h) = k] = 1$, it follows that

$$\Pr[N(h) = 0] = 1 - \lambda h + o(h) \tag{v}$$

Combining these with the stationarity in (ii), we obtain

$$P_0(t+h) = P_0(t)\,(1 - \lambda h + o(h))$$

or

$$\frac{P_0(t+h) - P_0(t)}{h} = -\lambda P_0(t) + \frac{o(h)}{h}$$

from which, in the limit $h \to 0$, the differential equation

$$P_0'(t) = -\lambda P_0(t)$$

is immediate. The solution is $P_0(t) = Ce^{-\lambda t}$ and the integration constant C follows from (i) and $P_0(0) = \Pr[N(0) = 0] = 1$ as $C = 1$. This establishes condition (iii) in the definition of the Poisson process for $k = 0$.

The verification for $n > 0$ is more involved. Applying the law of total probability (2.46),

$$\begin{aligned} P_n(t+h) &= \Pr[N(t+h) = n] \\ &= \sum_{j=0}^{n} \Pr[N(t+h) - N(t) = j | N(t) = n-j] \Pr[N(t) = n-j] \end{aligned}$$

By independence (ii),

$$\Pr[N(t+h) - N(t) = j | N(t) = n-j] = \Pr[N(t+h) - N(t) = j]$$

and by definition $\Pr[N(t) = n-j] = P_{n-j}(t)$, we have

$$P_n(t+h) = \sum_{j=0}^{n} \Pr[N(t+h) - N(t) = j] \, P_{n-j}(t)$$

By the stationarity (ii)

$$\Pr[N(t+h) - N(t) = j] = \Pr[N(h) - N(0) = j]$$

we obtain using (i)

$$P_n(t+h) = \sum_{j=0}^{n} \Pr[N(h) = j] \, P_{n-j}(t)$$

while (v) and (iii) suggest to write the sum as

$$\begin{aligned} P_n(t+h) = P_n(t) \Pr[N(h) = 0] + P_{n-1}(t) \Pr[N(h) = 1] \\ + \sum_{j=2}^{n} P_{n-j}(t) \Pr[N(h) = j] \end{aligned}$$

Since $P_n(t) \leq 1$ and using (iv),

$$\sum_{j=2}^{n} P_{n-j}(t) \Pr[N(h) = j] \leq \sum_{j=2}^{n} \Pr[N(h) = j] = \Pr[N(h) > 1] = o(h)$$

we arrive with (v), (iii) at

$$P_n(t+h) = P_n(t) (1 - \lambda h + o(h)) + P_{n-1}(t) (\lambda h + o(h)) + o(h)$$

or

$$\frac{P_n(t+h) - P_n(t)}{h} = -\lambda P_n(t) + \lambda P_{n-1}(t) + \frac{o(h)}{h}$$

which leads, after taking the limit $h \to 0$, to the differential equation

$$P_n'(t) = -\lambda P_n(t) + \lambda P_{n-1}(t)$$

with initial condition $P_n(0) = \Pr[N(0) = n] = 1_{\{n=0\}}$. This differential equation is rewritten as

$$\frac{d}{dt}\left(e^{\lambda t}P_n(t)\right) = \lambda e^{\lambda t}P_{n-1}(t) \tag{7.3}$$

In case $n = 1$, the differential equation reduces with $P_0(t) = e^{-\lambda t}$ to $\frac{d}{dt}\left(e^{\lambda t}P_1(t)\right) = \lambda$. The general solution is $e^{\lambda t}P_1(t) = \lambda t + C$ and, from the initial condition $P_1(0) = 0$, we have $C = 0$ and $P_1(t) = \lambda t e^{-\lambda t}$. The general solution to (7.3) is proved by induction. Assume that $P_n(t) = \frac{(\lambda t)^n e^{-\lambda t}}{n!}$ holds for n, then the case $n+1$ follows from (7.3) as

$$\frac{d}{dt}\left(e^{\lambda t}P_{n+1}(t)\right) = \lambda\frac{(\lambda t)^n}{n!}$$

and integrating from 0 to t using $P_{n+1}(0) = 0$, yields $P_{n+1}(t) = \frac{(\lambda t)^{n+1}e^{-\lambda t}}{(n+1)!}$ which establishes the induction and finalizes the proof of the theorem. □

The second theorem has very important applications since it relates the number of events in non-overlapping intervals to the interarrival time between these events.

Theorem 7.3.2 *Let $\{X(t); t \geq 0\}$ be a Poisson process with rate $\lambda > 0$ and denote by $t_0 = 0 < t_1 < t_2 < \cdots$ the successive occurrence times of events. Then the interarrival times $\tau_n = t_n - t_{n-1}$ are independent identically distributed exponential random variables with mean $\frac{1}{\lambda}$.*

Proof: For any $s \geq 0$ and any $n \geq 1$, the event $\{\tau_n > s\}$ is equivalent to the event $\{X(t_{n-1}+s) - X(t_{n-1}) = 0\}$. Indeed, the n-th interarrival time τ_n can only be longer than s time units if and only if the n-th event has not yet occurred s time units after the occurrence of the $(n-1)$-th event at t_{n-1}. Since the Poisson process has independent increments (condition (ii) in the definition of the Poisson process), changes in the value of the process in non-overlapping time intervals are independent. By the equivalence in events, this implies that the set of interarrival times τ_n are independent random variables. Further, by the stationarity of the Poisson process (deduced from condition (iii) in the definition of the Poisson process),

$$\Pr[\tau_n > s] = \Pr[X(t_{n-1}+s) - X(t_{n-1}) = 0] = e^{-\lambda s}$$

which implies that any interarrival time has an identical, exponential distribution,

$$F_{\tau_n}(x) = \Pr[\tau_n \leq x] = 1 - e^{-\lambda x}$$

This proves the theorem. □

The converse of Theorem 7.3.2 also holds: if the interarrival times $\{\tau_n\}$ of a counting process $\{N(t), t \geq 0\}$ are i.i.d. exponential random variables with mean $\frac{1}{\lambda}$, then $\{N(t), t \geq 0\}$ is a Poisson process with rate λ.

An association to the exponential distribution is the memoryless property,

$$\Pr[\tau_n > s + t | \tau_n > s] = \Pr[\tau_n > t]$$

By the equivalence of the events, for any $t, s \geq 0$,

$$\begin{aligned}\Pr[\tau_n > s + t | \tau_n > s] &= \Pr[X(t_{n-1} + s + t) - X(t_{n-1}) = 0 | X(t_{n-1} + s) - X(t_{n-1}) = 0] \\ &= \Pr[X(t_{n-1} + s + t) - X(t_{n-1} + s) = 0 | X(t_{n-1} + s) - X(t_{n-1}) = 0]\end{aligned}$$

By the independence of increments (in non-overlapping intervals),

$$\Pr[\tau_n > s + t | \tau_n > s] = \Pr[X(t_{n-1} + s + t) - X(t_{n-1} + s) = 0]$$

and by the stationarity of the increments, the memoryless property is established,

$$\Pr[\tau_n > s + t | \tau_n > s] = \Pr[X(t_{n-1} + t) - X(t_{n-1}) = 0] = \Pr[\tau_n > t]$$

Hence, the assumption of stationary and independent increments is equivalent to asserting that, at any time s, the process probabilistically restarts again with the same distribution and is independent of occurrences in the past (before s). Thus, the process has no memory and, since the only continuous distribution that satisfies the memoryless property is the exponential distribution, exponential interarrival times τ_n are a natural consequence.

The arrival time of the n-th event or the waiting time until the n-event is $W_n = \sum_{k=1}^{n} \tau_k$. In Section 3.3.1, it is shown that the probability distribution of the sum of independent exponential random variables has a Gamma distribution or Erlang distribution (3.24). Alternatively, the equivalence of the events, $\{W_n \leq t\} \Longleftrightarrow \{N(t) \geq n\}$, directly leads to the Erlang distribution,

$$F_{W_n}(t) = \Pr[W_n \leq t] = \Pr[N(t) \geq n] = \sum_{k=n}^{\infty} \frac{(\lambda t)^k e^{-\lambda t}}{k!}$$

The equivalence of the events, $\{W_n \leq t\} \Longleftrightarrow \{N(t) \geq n\}$, is a general relation and a fundamental part of the theory of renewal processes, which we will study in the next Chapter 8.

Theorem 7.3.3 *Given that exactly one event of a Poisson process* $\{X(t); t \geq 0\}$ *has occurred during the interval* $[0, t]$*, the time of occurrence of this event is uniformly distributed over* $[0, t]$.

Proof: Immediate application of the conditional probability (2.44) yields for $0 \leq s \leq t$,

$$\Pr\left[\tau_1 \leq s | X(t) = 1\right] = \frac{\Pr\left[\{\tau_1 \leq s\} \cap \{X(t) = 1\}\right]}{\Pr\left[X(t) = 1\right]}$$

Using the equivalence $\{\tau_1 \leq s\} \Longleftrightarrow \{X(t_0 + s) - X(t_0) = 1\}$ and the fact that $\{X(t_0 + s) - X(t_0) = 1\} = \{X(s) = 1\}$ by the stationarity of the Poisson process gives

$$\begin{aligned}\{\tau_1 \leq s\} \cap \{X(t) = 1\} &= \{X(s) = 1\} \cap \{X(t) = 1\} \\ &= \{X(s) = 1\} \cap \{X(t) - X(s) = 0\}\end{aligned}$$

Applying the independence of increments over non-overlapping intervals and (7.1) yields

$$\begin{aligned}\Pr\left[\tau_1 \leq s | X(t) = 1\right] &= \frac{\Pr\left[X(s) = 1\right] \Pr\left[X(t) - X(s) = 0\right]}{\Pr\left[X(t) = 1\right]} \\ &= \frac{(\lambda s)\, e^{-\lambda s} e^{-\lambda(t-s)}}{(\lambda t)\, e^{-\lambda t}} = \frac{s}{t}\end{aligned}$$

which completes the proof. □

Theorem 7.3.3 is immediately generalized to n events. For any set of real variables s_j satisfying $0 = s_0 < s_1 < s_2 < \cdots < s_n < t$ and given that n events of a Poisson process $\{X(t); t \geq 0\}$ have occurred during the interval $[0, t]$, the probability of the successive occurrence times $0 < t_1 < t_2 < \cdots < t_n < t$ of these n Poisson events is

$$\Pr\left[t_1 \leq s_1, \ldots, t_n < s_n | X(t) = n\right] = \frac{\Pr\left[\{t_1 \leq s_1, \ldots, t_n < s_n\} \cap \{X(t) = n\}\right]}{\Pr\left[X(t) = n\right]}$$

Using a similar argument as in the proof of Theorem 7.3.3,

$$\begin{aligned}p &= \Pr\left[\{t_1 \leq s_1, t_2 \leq s_2, \ldots, t_n < s_n\} \cap \{X(t) = n\}\right] \\ &= \Pr\left[X(s_1) - X(s_0) = 1, \ldots, X(s_n) - X(s_{n-1}) = 1, X(t) - X(s_n) = 0\right] \\ &= \left(\prod_{j=1}^{n} \Pr\left[X(s_j) - X(s_{j-1}) = 1\right]\right) \Pr[X(t) - X(s_n) = 0] \\ &= \left(\prod_{j=1}^{n} e^{-\lambda(s_j - s_{j-1})} \lambda\,(s_j - s_{j-1})\right) e^{-\lambda(t - s_n)} \\ &= \lambda^n \prod_{j=1}^{n} (s_j - s_{j-1})\, e^{-\lambda \sum_{j=1}^{n}(s_j - s_{j-1}) - \lambda(t - s_n)} = \lambda^n \prod_{j=1}^{n} (s_j - s_{j-1})\, e^{-\lambda t}\end{aligned}$$

Thus,

$$\begin{aligned}\Pr\left[t_1 \le s_1, t_2 \le s_2, \ldots, t_n < s_n | X(t) = n\right] &= \frac{\lambda^n \prod_{j=1}^{n} \left(s_j - s_{j-1}\right) e^{-\lambda t}}{\frac{(\lambda t)^n e^{-\lambda t}}{n!}} \\ &= \frac{n!}{t^n} \prod_{j=1}^{n} \left(s_j - s_{j-1}\right)\end{aligned}$$

from which the density function

$$f_{\{t_j\}}\left(s_1, \ldots, s_n | X(t) = n\right) = \frac{\partial^n}{\partial s_1 \ldots \partial s_n} \Pr\left[t_1 \le s_1, \ldots, t_n < s_n | X(t) = n\right]$$

follows as

$$f_{\{t_j\}}\left(s_1, s_2, \ldots, s_n | X(t) = n\right) = \frac{n!}{t^n}$$

which is independent of the rate λ. If $0 < t_1 < t_2 < \cdots < t_n < t$ are the successive occurrence times of n Poisson events in the interval $[0, t]$, then the random variables $t_1, t_2, \ldots, t_n$ are distributed as a set of order statistics, defined in Section 3.4.2, of n uniform random variables in $[0, t]$. In other words, if n i.i.d. uniform random variables on $[0, t]$ are assorted in increasing order, they may represent n successive occurrence times of a Poisson process. The average spacing between these n ordered i.i.d. uniform random variables is $\frac{t}{n}$ as computed in Problem (ii) of Section 3.7.

A related example is the conditional probability where $0 < s < t$ and $0 \le k \le n$,

$$\begin{aligned}\Pr\left[X(s) = k | X(t) = n\right] &= \frac{\Pr\left[\{X(s) = k\} \cap \{X(t) = n\}\right]}{\Pr\left[X(t) = n\right]} \\ &= \frac{\Pr\left[\{X(s) = k\} \cap \{X(t) - X(s) = n - k\}\right]}{\Pr\left[X(t) = n\right]} \\ &= \frac{\Pr\left[X(s) = k\right] \Pr[X(t) - X(s) = n - k]}{\Pr\left[X(t) = n\right]} \\ &= \frac{n!(\lambda s)^k e^{-\lambda s}}{k!(\lambda t)^n e^{-\lambda t}} \frac{(\lambda(t - s))^{n-k} e^{-\lambda(t-s)}}{(n - k)!} \\ &= \binom{n}{k} \frac{s^k}{t^n} (t - s)^{n-k}\end{aligned}$$

Hence, if $p = \frac{s}{t}$, the conditional probability becomes

$$\Pr\left[X(s) = k | X(t) = n\right] = \binom{n}{k} p^k (1 - p)^{n-k}$$

Given that a total number of n Poisson events have occurred in time interval $[0,t]$, the chance that precisely k events have taken place in the sub-interval $[0,s]$ is binomially distributed with parameter n and $p=\frac{s}{t}$. Observe that also this conditional probability is independent of the rate λ. In addition, since $\lim_{t\to\infty} X(t)=\infty$ such that $n\to\infty$, applying the law of rare events results in

$$\lim_{t\to\infty} \Pr\left[X(s)=k|X(t)=n\right]=\frac{s^k}{k!}e^{-s}$$

Given an everlasting Poisson process, the chance that precisely k events occur in the interval $[0,s]$ is Poisson distributed with mean equal to the length of the interval.

Application The arrival process of most real-time applications (such as telephony calls, interactive-video, ...) in a network is well approximated by a Poisson process. Suppose a measurement configuration is built to collect statistics of the arrival process of telephony calls in some region. During a period $[0,T]$, precisely 1 telephony call has been measured. What can be said of the time $x\in[0,T]$ at which the telephony call has arrived at the measurement device? Theorem 7.3.3 tells us that any time in that interval is equally probable.

Theorem 7.3.4 *If $X(t)$ and $Y(t)$ are two independent Poisson processes with rates λ_x and λ_y, then is $Z(t)=X(t)+Y(t)$ also a Poisson process with rate $\lambda_x+\lambda_y$.*

Proof: It suffices to demonstrate that the counting process $N_Z(t)=N_X(t)+N_Y(t)$ has exponentially distributed interarrival times τ_Z. Suppose that $N_Z(t_n)=n$, it remains to compute the next arrival at time $t_{n+1}=t_n+s$ for which $N_Z(t_n+s)=n+1$. Due to the memoryless property of the Poisson process, the occurrence of an event from t_n on for each random variable X and Y is again exponentially distributed with parameter λ_x and λ_y, respectively. In other words, it is irrelevant which process X or Y has previously caused the arrival at time t_n. Further, the event that the interarrival time of the sum processes $\{\tau_Z>s\}$ is equivalent to $\{\tau_X>s\}\cap\{\tau_Y>s\}$ or

$$\Pr\left[\tau_Z>s\right]=\Pr\left[\tau_X>s,\tau_Y>s\right]=\Pr\left[\tau_X>s\right]\Pr\left[\tau_Y>s\right]=e^{-(\lambda_x+\lambda_y)s}$$

where the independence of $X(t)$ and $Y(t)$ has been used. This proves the theorem. $\square$

A direct consequence is that any sum of independent Poisson processes is also a Poisson process with aggregate rate equal to the sum of the individual rates. This theorem is in correspondence with the sum property of the Poisson distribution.

7.4 The nonhomogeneous Poisson process

As will be shown later in Section 11.3.2, the Poisson process is a special case of a birth-and-death process, which is in turn a special case of a Markov process. Hence, it seems more instructive to discuss these special processes as applications of the Markov process. Therefore, only associations to the Poisson process are treated here. In many cases, the rate is a time variant function $\lambda(t)$ and such process is termed *a nonhomogeneous or nonstationary Poisson process*. For example, the arrival rate of a large number m of individual IP-flows at a router is well approximated by a nonhomogeneous Poisson process, where the rate $\lambda(t)$ varies over the day depending on the number m and the individual rate of each flow of packets. Since the sum of independent Poisson random variables is again a Poisson random variable, we have $\lambda(t) = \sum_{j=1}^{m(t)} \lambda_j(t)$.

If $X(t)$ is a nonhomogeneous Poisson process with rate $\lambda(t)$, the increment $X(t)-X(s)$ reflects the number of events in an interval $(s,t]$ and increments of non-overlapping intervals are still independent.

Theorem 7.4.1 *If* $\Lambda(t) = \int_0^t \lambda(u)du$ *and* $s < t$, *then* $X(t)-X(s)$ *is Poisson distributed with mean* $\Lambda(t) - \Lambda(s)$.

The demonstration is analogous to the proof of Theorem 7.3.1.

Proof (partly): Denote by $P_n(t) = \Pr[N(t) - N(s) = n]$, then

$$\begin{aligned} P_0(t+h) &= \Pr[N(t+h) - N(s) = 0] \\ &= \Pr[N(t+h) - N(t) = 0, N(t) - N(s) = 0] \end{aligned}$$

Invoking independence of the increments,

$$\begin{aligned} P_0(t+h) &= \Pr[N(t+h) - N(t) = 0] \Pr[N(t) - N(s) = 0] \\ &= P_0(t)(1 - \lambda(t)h + o(h)) \end{aligned}$$

or

$$\frac{P_0(t+h) - P_0(t)}{h} = -\lambda(t)P_0(t) + \frac{o(h)}{h}$$

from which, in the limit $h \to 0$, the differential equation

$$P_0'(t) = -\lambda(t)P_0(t)$$

is immediate. Rewritten as $\frac{d}{dt}\log P_0(t) = -\lambda(t)$, after integration over $(s, t]$, we find $\log P_0(t) = -(\Lambda(t) - \Lambda(s))$ since $P_0(s) = \Pr[N(s) - N(s) = 0] = 1$. Thus, for the case $n = 0$, we find $P_0(t) = \exp[-(\Lambda(t) - \Lambda(s))\, t]$, which proves the theorem for $n = 0$.

The remainder of the proof ($n > 0$) uses the same ingredients as the proof of Theorem 7.3.1 and is omitted. □

A nonhomogeneous Poisson process $X(t)$ with rate $\lambda(t)$ can be transformed to a homogeneous Poisson process $Y(u)$ with rate 1 by the time transform $u = \Lambda(t)$. For, $Y(u) = Y(\Lambda(t)) = X(t)$, and $Y(u + \Delta u) = Y(\Lambda(t) + \Delta\Lambda(t)) = X(t + \Delta t)$ because $\Delta\Lambda(t) = \Lambda(\Delta t)$ for small Δt such that

$$\begin{aligned}\Pr[Y(u + \Delta u) - Y(u) = 1] &= \Pr[X(t + \Delta t) - X(t) = 1]\\ &= \lambda(t)\Delta t + o(\Delta t)\\ &= \Delta u + o(\Delta u)\end{aligned}$$

because $\Delta u = \lambda(t)\Delta t + o(\Delta t)$. Hence, all problems concerning nonhomogeneous Poisson processes can be reduced to the homogeneous case treated above.

7.5 The failure rate function

Previous sections have shown that the Poisson process is specified by a rate function $\lambda(t)$. In this section, we consider the failure rate function of some object or system. Often it is interesting to know the probability that an object will fail in the interval $[t, t + \Delta t]$ given that the object was still functioning well up to time t. Let X denote the lifetime of an object[5], then this probability can be written with (2.44) as

$$\begin{aligned}\Pr[t \le X \le t + \Delta t | X > t] &= \frac{\Pr[\{t \le X \le t + \Delta t\} \cap \{X > t\}]}{\Pr[X > t]}\\ &= \frac{\Pr[t < X \le t + \Delta t]}{\Pr[X > t]}\end{aligned}$$

If $f_X(t)$ is the probability density function of X and $F_X(t) = \Pr[X \le t]$, then for small Δt and assuming that $f_X(t)$ is well behaved[6] such that

---

[5] In medical sciences, X can represent in general the time for a certain event to occur. For example, the time it takes for an organism to die, the time to recover from illness, the time for a patient to respond to a therapy and so on.

[6] Recall the discussion in Section 2.3.

$\Pr[t < X \leq t + \Delta t] = f_X(t)\Delta t$,

$$\Pr[t \leq X \leq t + \Delta t | X > t] = \frac{f_X(t)}{1 - F_X(t)} \Delta t$$

This expression shows that

$$r(t) = \frac{f_X(t)}{1 - F_X(t)} \tag{7.4}$$

can be interpreted as the intensity or rate that a t-year old object will fail. It is called *the failure rate* $r(t)$ and

$$R(t) = 1 - F_X(t) = \Pr[X > t] \tag{7.5}$$

is usually termed[7] *the reliability function*. Since $r(t) = \frac{\Pr[t \leq X \leq t+\Delta t | X > t]}{\Delta t}$ for small Δt, the failure rate $r(t) > 0$ because $r(t) = 0$ would imply an infinite lifetime X. Using the definition (2.30) of a probability density function, we observe that

$$r(t) = -\frac{\frac{dR(t)}{dt}}{R(t)} = -\frac{d \ln R(t)}{dt} \tag{7.6}$$

Or, since $R(0) = 1$, the corresponding integrated relation is

$$R(t) = \exp\left[-\int_0^t r(u) du\right] \tag{7.7}$$

The expressions (7.6) and (7.7) are inverse relations that specify $r(t)$ as function of $R(t)$ and vice versa. The reliability function $R(t)$ is non-increasing with maximum at $t = 0$ since it is a probability distribution function. On the other hand, the failure rate $r(t)$ being a probability density function can take any positive real value. From (7.4) we obtain the density function of the lifetime X in terms of failure rate $r(t)$ as

$$f_X(t) = r(t)R(t) = r(t)\exp\left[-\int_0^t r(u) du\right]$$

with $f_X(0) = r(0)$. Using the tail relation (2.35) for the expectation of the lifetime X immediately gives the *mean time to failure*,

$$E[X] = \int_0^\infty R(t) dt \tag{7.8}$$

In case $F_X(T) = 1$ and $f_X(T) \neq 0$ for a finite time T, which is the maximum lifetime, the definition (7.4) demonstrates that $r(t)$ has a pole at

---

[7] In biology, medical sciences and physics, $R(t)$ is called the survival function and $r(t)$ is the corresponding mortality rate or hazard rate.

$t = T$. In practice, the failure rate $r(t)$ is relatively high for small t due to initial imperfections that cause a number of objects to fail early and $r(t)$ is increasing towards the maximum life time T due to aging or wear and tear. This shape of $r(t)$ as illustrated in Fig. 7.4 is called a "bath-tub" curve, which is convex.

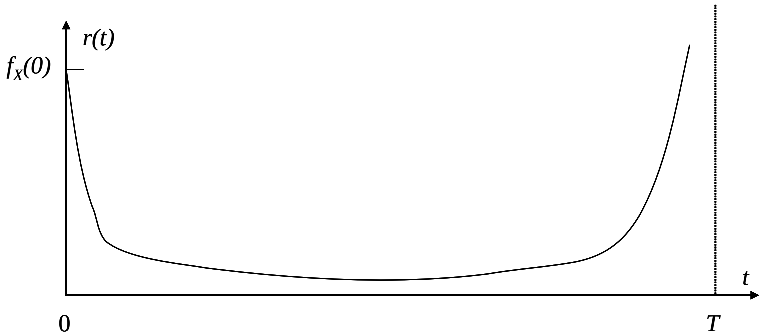

Fig. 7.4. Example of a "bath-tub" shaped failure rate function $r(t)$.

An often used model for the failure rate is $r(t) = a\lambda t^{a-1}$ with corresponding reliability function $R(t) = \exp[-\lambda t^a]$ and where the lifetime X has a Weibull distribution function $F_X(t) = 1 - R(t)$ as in (3.40). In case $a = 1$, the failure rate $r(t) = \lambda$ is constant over time, while $a > 1$ ($a < 1$) reflects an increasing (decreasing) failure rate over time. Hence, a "bath-tub" shaped (realistic) failure function as in Fig. 7.4 can be modeled by a Weibull model for $r(t)$ with $a < 1$ in the beginning, $a = 1$ in the middle and $a > 1$ at the end of the life time.

For an exponential lifetime where $f_X(t) = \lambda e^{-\lambda t}$, the failure rate (7.4) equals $r(t) = \lambda$ and is independent of time. This means that the failure rate for a t-year-old object is the same as for a new object, which is a manifestation of the memoryless property of the exponential distribution. It also explains why λ in both the exponential as Poisson process is often called a 'rate'.

7.6 Problems

(i) A series of test strings each with a variable number N of bits all equal to 1 are transmitted over a channel. Due to transmission errors, each 1-bit can be effected independently from the others and only arrives non-corrupted with probability p. The length N of the test strings (words) is a Poisson random variable with mean length λ bits. In this test, the sum Y of the bits in the arriving words is investigated to determine the channel quality via p. Compute the pdf of Y.

(ii) At a router, four QoS classes are supported and for each class packets arrive according to a Poisson process with rate λ_j for $j = 1, 2, 3, 4$. Suppose that the router had a failure at time t_1 that lasted T time units. What is the probability density function of the total number of packets of the four classes that has arrived during that period?

(iii) Let $N(t) = N_1(t) + N_2(t)$ be the sum of two independent Poisson processes with rates λ_1 and λ_2. Given that the process $N(t)$ had an arrival, what is the probability that that arrival came from the process $N_1(t)$?

(iv) Peter has been monitoring the highway for nearly his entire life and found that the cars pass his house according to a Poisson process. Moreover, he discovered that the Poisson process in one lane is independent from that in the other lanes. The rate of these independent processes differs per lane and is denoted by $\lambda_1, \lambda_2, \lambda_3$, where λ_j is expressed in the number of cars on lane j per hour.

  (a) Given that one car passed Peter, what is the probability that it passed in lane 1?
  (b) What is the probability that n cars pass Peter in 1 hour ?
  (c) What is the probability that in 1 hour n cars have passed and that they all have used lane 1?

(v) In a game, audio signals arrive in the interval $(0, T)$ according to a Poisson process with rate λ, where $T > 1/\lambda$. The player wins only if at least one audio signal arrives in that interval, and if he or she pushes a button (only one push allowed) upon the last of the signals. The player uses the following strategy: he or she pushes the button upon the arrival of the first signal (if any) after a fixed time $s \leq T$.

  (a) What is the probability that the player wins?
  (b) Which value of s maximizes the probability of winning, and what is the probability in that case?

(vi) The arrivals of voice over IP (VoIP) packets to a router is close to a Poisson process with rate $\lambda = 0.1$ packets per minute. Due to an upgrade to install weighted fair queueing as priority scheduling rule, the router is switched off for 10 minutes.

  (a) What is the probability of receiving no VoIP packets when switched off?
  (b) What is the probability that more than ten VoIP packets will arrive during this upgrade?

   (c) If there was one VoIP in the meantime, what is the most probable minute of the arrival?

(vii) A link of a packet network carries on average ten packets per second. The packets arrive according to a Poisson process. A packet has a probability of 30 % to be an acknowledgment (ACK) packet independent of the others. The link is monitored during an interval of 1 second.

   (a) What is the probability that at least one ACK packet has been observed?
   (b) What is the expected number of all packets given that five ACK packets have been spotted on the link?
   (c) Given that eight packets have been observed in total, what is the probability that two of them are ACK packets?

(viii) An ADSL helpdesk treats exclusively customer requests of one of three types: (i) login-problems, (ii) ADSL hardware and (iii) ADSL software problems. The opening hours of the helpdesk are from 8:00 until 16:00. All requests are arriving at the helpdesk according to a Poisson process with different rates: $\lambda_1 = 8$ requests with login problems/hour, $\lambda_2 = 6$ requests with hardware problems/hour, and $\lambda_3 = 6$ requests with software problems/hour. The Poisson arrival processes for different types of requests are independent.

   (a) What is the expected number of requests in one day?
   (b) What is the probability that in 20 minutes exactly three requests arrive, and that all of them have hardware problems?
   (c) What is the probability that no requests will arrive in the last 15 minutes of the opening hours?
   (d) What is the probability that one request arrives between 10:00 and 10:12 and two requests arrive between 10:06 and 10:30?
   (e) If at the moment $t + s$ there are $k + m$ requests, what is the probability that there were k requests at the moment t?

(ix) Arrival of virus attacks to a PC can be modeled by a Poisson process with rate $\lambda = 6$ attacks per hour.

   (a) What is the probability that exactly one attack will arrive between 1 p.m. and 2 p.m.?
   (b) Suppose that at the moment the PC is turned on there were no attacks on PC, but at the shut-down time precisely 60 attacks have been observed. What is the expected amount of time that the PC has been on?

(c) Given that six attacks arrive between 1 p.m. and 2 p.m., what is the probability that the fifth attack will arrive between 1:30 p.m. and 2 p.m.?

(d) What is the expected arrival time of that fifth attack?

(x) Consider a system S consisting of n subsystems in series as shown in Fig. 7.5. The system S operates correctly only if all subsystems operate correctly. Assume that the probability that a failure in a subsystem S_i occurs is independent of that in subsystem S_j. Given the reliability functions $R_j(t)$ or each subsystem S_j, compute the reliability function $R(t)$ of the system S.

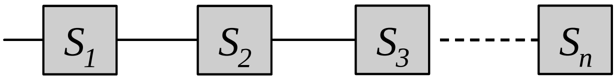

Fig. 7.5. A system consisting of n subsystems in series.

(xi) Same question as in previous exercise but applied to a system S consisting of n subsystem in parallel as shown in Fig. 7.6.

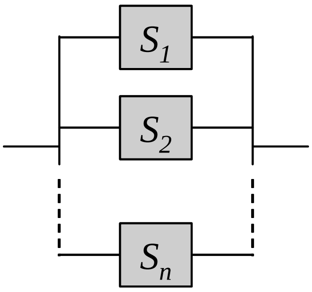

Fig. 7.6. A system consisting of n subsystems in parallel.