Chapter 2
Lévy Processes and Lévy-Driven Queues

In classical queueing systems, there is the notion of customers (or work) arriving, and subsequently being processed by the server. The class of Lévy processes, being defined as processes with stationary and independent increments, covers processes with highly non-regular trajectories (think for instance of Brownian motion). As a consequence, it is not immediately clear how one should define a *queue with Lévy input*. One of the goals of the present chapter is to introduce a sound notion of Lévy-driven queues.

We do so by first providing an explicit definition of Lévy processes, and then extending the classical definition of a queue to a notion that can be used for general input processes as well (i.e. in principle *any* real-valued stochastic process can serve as input). For more background, we refer the reader e.g. to Applebaum [11], Asmussen [19], Kyprianou [146], and Sato [193].

In Section 2.1, as a first step we introduce notation, to be used throughout this book, together with a number of fundamental properties. As mentioned in Chapter 1, for the special case of one-sided jumps, the results are more explicit. Notation related to such *spectrally one-sided* Lévy processes is given in Section 2.2; this section also includes a number of frequently used Lévy processes. Another important class of Lévy processes, that is, α-stable Lévy motions, is covered by Section 2.3. Finally, in Section 2.4 we present the definition of Lévy-driven queues.

2.1 Infinitely Divisible Distributions, Lévy Processes

We say that a continuous-time real-valued stochastic process $(X_t)_t$ is a Lévy process if it has stationary and independent increments, with $X_0 = 0$ and càdlàg sample paths a.s. (*càdlàg* meaning 'continuous from right, limits from left'). The stationary increments property entails that for given s the distribution of $X_{t+s} - X_t$ is the same irrespective of the value of t, whereas the independent increments property means

© Springer International Publishing Switzerland 2015
K. Dębicki, M. Mandjes, *Queues and Lévy Fluctuation Theory*, Universitext,
DOI 10.1007/978-3-319-20693-6_2

that, for $t \geq 0$, the increment $X_{t+s} - X_s$ is independent of the history of the Lévy process, that is, $(X_u)_{u \leq s}$.

The initial condition $X_0 = 0$ together with the stationary increments property leads, for each $t > 0$, to the equation

$$X_t = \sum_{i=1}^{n} \left(X_{it/n} - X_{(i-1)t/n} \right),$$

in which the increments $X_{it/n} - X_{(i-1)t/n}$ are all distributed as $X_{t/n}$. Moreover, by virtue of the independent increments property, it follows that these increments are also independent. We thus arrive at the following distributional equality, with $X_t^{(i)}$ i.i.d. copies of X_t:

$$X_t \stackrel{\mathrm{d}}{=} \sum_{i=1}^{n} X_{t/n}^{(i)}, \tag{2.1}$$

for any $n \in \mathbb{N}$. In this way we see that, for any t, X_t has an *infinitely divisible* distribution. Indeed, let us recall that a random variable Z is infinitely divisible if for any $n \in \mathbb{N}$ there exist independent and identically distributed (i.i.d.) random variables $Z_{1,n}, \dots, Z_{n,n}$ such that Z is distributed as $\sum_{m=1}^{n} Z_{m,n}$; see e.g. De Finetti [70]. Conversely, for each infinitely divisible random variable Z there exists a Lévy process $(X_t)_t$ such that $X_1 \stackrel{\mathrm{d}}{=} Z$. This, for example, straightforwardly implies the existence of a Lévy process with Poisson marginals: if Z has a Poisson distribution with mean λ, it is distributed as the sum of n independent Poisson random variables with mean λ/n. Other examples of infinitely divisible distributions are the normal distribution, the negative binomial distribution, and the gamma distribution, as is readily verified.

One can alternatively say that, for any value of t,

$$\xi_t(s) := \log \mathbb{E} e^{\mathrm{i}sX_t} = t \log \mathbb{E} e^{\mathrm{i}sX_1} = t\xi(s),$$

for $s \in \mathbb{R}$, where $\xi(s) := \log \mathbb{E} e^{\mathrm{i}sX_1}$ is referred to as the so-called *Lévy exponent*. This equality is a direct consequence of (2.1), as can be seen as follows. Fixing an $s \in \mathbb{R}$, we find for any two integers m and n that $\xi_m(s) = n\xi_{m/n}(s)$ and $\xi_m(s) = m\xi_1(s)$. Combining these relations, we obtain $\xi_{m/n}(s) = (m/n)\,\xi_1(s) = (m/n)\,\xi(s)$, and hence for all $t \in \mathbb{Q}$ it follows that $\xi_t(s) = t\xi(s)$. By using a limiting argument, it follows immediately that the right continuity of the Lévy process implies that $\xi_t(s) = t\xi(s)$ for any $t \in \mathbb{R}$. As a result, one could informally say that each Lévy process can be associated with an infinitely divisible distribution, and vice versa.

It is immediately seen that the class of Lévy processes contains a number of canonical stochastic processes. In the first place it can be concluded that the *Poisson process* is Lévy. A Poisson process $(X_t)_t$ can be defined as follows: with Y_m i.i.d. exponential random variables with mean $\lambda^{-1} \in (0, \infty)$, we let X_t have the value n

if at the same time $\sum_{m=1}^{n} Y_m \leq t$ and $\sum_{m=1}^{n+1} Y_m > t$. It is well known that X_t has a Poisson distribution with mean λt, and as a consequence,

$$\log \mathbb{E} e^{\mathrm{i}sX_t} = \log \left(\sum_{n=0}^{\infty} e^{\mathrm{i}sn} e^{-\lambda t} \frac{(\lambda t)^n}{n!} \right) = \lambda t \left(e^{\mathrm{i}s} - 1 \right),$$

and hence $(X_t)_t$ is indeed Lévy (with Lévy exponent $\xi(s) = \lambda(e^{\mathrm{i}s} - 1)$ for $\lambda > 0$). Likewise, we can show that *Brownian motion* without drift is Lévy; here $\xi(s) = -\frac{1}{2}\sigma^2 s^2$ for $\sigma^2 > 0$. In Sections 2.2 and 2.3 we mention various other examples.

It is possible to characterize Lévy processes more specifically: it can be shown that the Lévy exponent $\xi(s)$ is necessarily of the form

$$\xi(s) = \mathrm{i}sd - \frac{1}{2}s^2\sigma^2 + \int_{-\infty}^{\infty} (e^{\mathrm{i}sx} - 1 - \mathrm{i}sx 1_{\{|x|<1\}}) \Pi(\mathrm{d}x), \tag{2.2}$$

where $d \in \mathbb{R}$, $\sigma \geq 0$, and the spectral measure (or *Lévy measure*) $\Pi(\cdot)$, concentrated on $\mathbb{R} \setminus \{0\}$, satisfies

$$\int_{\mathbb{R}} \min\{x^2, 1\} \Pi(\mathrm{d}x) < \infty.$$

For a proof of this fundamental representation of Lévy processes (or, in fact, a stronger version of it), called in the literature the *Lévy–Khintchine formula*, we refer e.g. to Kyprianou [146, Chapter II].

The triplet (d, σ^2, Π) is commonly referred to as the *characteristic triplet*, as it uniquely defines the underlying Lévy process: every Lévy process has its own specific d, σ^2, and Π. It is noted that in some cases it is possible to extend the domain of $\xi(s)$ to (a subset of) $\mathbb{C}$; we return to this issue in greater detail in Section 2.2 when we speak about Lévy processes with one-sided jumps.

For obvious reasons, we call the first parameter of the characteristic triplet, d, the *deterministic drift*, whereas the term $\frac{1}{2}s^2\sigma^2$ is often referred to as the *Brownian term*. The third term in (2.2) corresponds to the jumps of the Lévy process by the relation that the jumps of size x occur at intensity $\Pi(\mathrm{d}x)$. More precisely, for any bounded interval M such that $0 \notin M$, the sum of the jumps of size within M in the time interval $[0, t)$ is distributed as a compound Poisson random variable with intensity $t \int_M \Pi(\mathrm{d}y)$ and the jump-size distribution

$$\frac{\Pi(\mathrm{d}x) 1_{\{x \in M\}}}{\int_M \Pi(\mathrm{d}y)}.$$

Thus, if the jumps are only in the upward (respectively, downward) direction, then the support of Π is concentrated in $(0, \infty)$ (respectively, $(-\infty, 0)$). The process $(X_t)_t$ is of bounded variation if and only if both $\sigma = 0$ and $\int_{-1}^{1} |x| \Pi(\mathrm{d}x) < \infty$; we do not provide details on this, but refer to Kyprianou [146, Section 2.6.1].

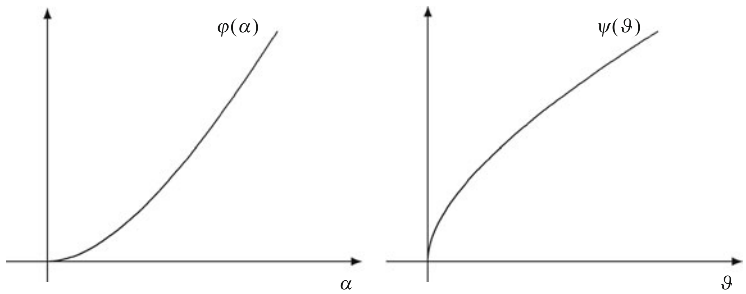

Fig. 2.1 Spectrally positive case: Laplace exponent and its inverse

2.2 Spectrally One-Sided Lévy Processes

Let $(X_t)_{t\geq 0}$ be a Lévy process, as introduced in Section 2.1. Unless stated otherwise, we assume throughout the book that the 'mean drift' $\mathbb{E}X_1$ of the Lévy process is negative, so as to make sure that the corresponding workload process (to be formally introduced in Section 2.4) is stable, thus guaranteeing the existence of a proper stationary workload distribution.

In this monograph, two special cases will often be considered in great detail, that is, spectrally positive and spectrally negative Lévy processes.

The Lévy process has no negative jumps—Here the Lévy process $(X_t)_{t\geq 0}$ has no negative jumps, or is *spectrally positive*; in the sequel this is denoted by $X \in \mathscr{S}_+$. In this case the spectral measure $\Pi(\cdot)$ is concentrated on $(0, \infty)$.

It turns out, in this case, to be convenient to work with the *Laplace exponent*, given by the function $\varphi(\alpha) := \log \mathbb{E}e^{-\alpha X_1}$, rather than the Lévy exponent $\xi(s)$. It is a consequence of the fact that there are only positive jumps that this Laplace exponent is well defined for all $\alpha \geq 0$.

It follows immediately from Hölder's inequality that the Laplace exponent $\varphi(\cdot)$ is convex on $[0, \infty)$; due to the assumption $\mathbb{E}X_1 < 0$, and observing that $\varphi(\cdot)$ has slope $\varphi'(0) = -\mathbb{E}X_1$ at the origin, we conclude that $\varphi(\cdot)$ is increasing on $[0, \infty)$, and hence the inverse $\psi(\cdot)$ of $\varphi(\cdot)$ is well defined on $[0, \infty)$; see Fig. 2.1. In the sequel we also require that X_t is not a *subordinator*, that is, a monotone process; this means that X_1 has probability mass on the negative half-line, which implies that $\lim_{\alpha\to\infty} \varphi(\alpha) = \infty$.

The Lévy process has no positive jumps—In this case the Lévy process $(X_t)_{t\geq 0}$ has no positive jumps, or is *spectrally negative*; throughout this book we denote this by $X \in \mathscr{S}_-$. Now the spectral measure $\Pi(\cdot)$ is concentrated on $(-\infty, 0)$. In this case, we define the *cumulant* $\Phi(\beta) := \log \mathbb{E}e^{\beta X_1}$. This function is well defined and finite for any $\beta \geq 0$ due to the fact that there are no positive jumps. We now rule out that $(X_t)_t$ has decreasing sample paths a.s. Recalling that $\Phi'(0) = \mathbb{E}X_1 < 0$, we see that $\Phi(\beta)$ is *not* a bijection on $[0, \infty)$; we define the *right* inverse through

$$\Psi(q) := \sup\{\beta \geq 0 : \Phi(\beta) = q\}.$$

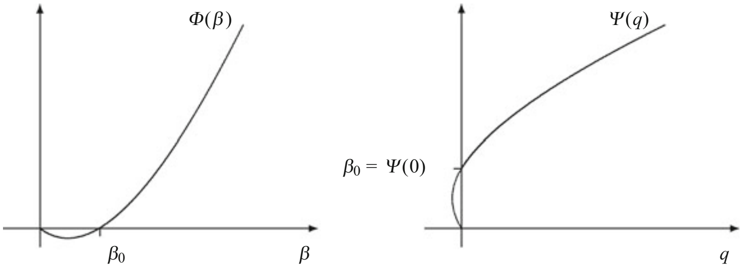

Fig. 2.2 Spectrally negative case: the cumulant and its right inverse

Note that $\beta_0 := \Psi(0) > 0$; this parameter plays a crucial role when analyzing queues with spectrally negative input; see Fig. 2.2.

The Lévy exponent (or the Laplace exponent for $X \in \mathscr{S}_+$, or cumulant for $X \in \mathscr{S}_-$) contains all information about X_1, and hence, due to the infinite divisibility, also about the whole process $(X_t)_t$. For instance, it enables the computation of all moments (provided they exist), as follows. For example, for $X \in \mathscr{S}_+$, we have $\mathbb{E}X_t = -\varphi'(0)\,t$ and $\mathbb{V}\mathrm{ar}\,X_t = \varphi''(0)\,t$ (given that these derivatives are well defined). It is also noted that

$$\varphi'(0) = -d - \int_{[1,\infty)} x\Pi(\mathrm{d}x), \qquad \varphi''(0) = \sigma^2 + \int_{(0,\infty)} x^2\Pi(\mathrm{d}x),$$

whereas, for $n = 3, 4, \ldots,$

$$\varphi^{(n)}(0) = (-1)^n \int_{(0,\infty)} x^n\Pi(\mathrm{d}x).$$

We now treat in greater detail a number of examples of spectrally one-sided Lévy processes.

(1) *Brownian motion with drift*. This process has sample paths that are continuous a.s., and is therefore both spectrally positive and spectrally negative. In this case X_t has a normal distribution with mean dt and variance $\sigma^2 t$. It is readily verified that, with U denoting a standard normal random variable, $\mathbb{E}e^{-\alpha X_t} = e^{-\alpha dt}\mathbb{E}e^{-\alpha\sqrt{t}\sigma U}$, and

$$\mathbb{E}e^{-\alpha U} = \int_{-\infty}^{\infty} e^{-\alpha u}\frac{1}{\sqrt{2\pi}}e^{-u^2/2}\mathrm{d}u = e^{\alpha^2/2}\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}}e^{-(u+\alpha)^2/2}\mathrm{d}u = e^{\alpha^2/2}.$$

It is concluded that $\log \mathbb{E}e^{-\alpha X_t} = t(-\alpha d + \frac{1}{2}\alpha^2\sigma^2)$. We write $X \in \mathbb{B}\mathrm{m}(d, \sigma^2)$ when $\varphi(\alpha) = -\alpha d + \frac{1}{2}\alpha^2\sigma^2$. The mean drift of this process is d, which is assumed to be negative (to make sure that $\mathbb{E}X_1 < 0$).

(2) *Compound Poisson with drift*. This process corresponds to i.i.d. jobs arriving according to a Poisson process, from which a deterministic drift is subtracted. More concretely, we let the jobs $B_1, B_2, \ldots$ be i.i.d. positive-valued random variables with Laplace transform $b(\alpha) := \mathbb{E}e^{-\alpha B}$ and $(N_t)_t$ be a Poisson process of rate λ (independent of the job sizes). Then the time-changed random walk, with the parameter r assumed to be positive,

$$X_t = \sum_{i=1}^{N_t} B_i - rt$$

(following the convention that $\sum_{i=1}^{0} B_i := 0$) is a spectrally positive Lévy process which we call a compound Poisson process with drift. We write $X \in \mathbb{C}\mathrm{P}(r, \lambda, b(\cdot))$.

It can be verified that

$$\mathbb{E}e^{-\alpha X_t} = e^{r\alpha t} \sum_{n=0}^{\infty} (b(\alpha))^n e^{-\lambda t} \frac{(\lambda t)^n}{n!} = \exp\left(t(r\alpha - \lambda + \lambda b(\alpha))\right).$$

As a consequence, $\varphi(\alpha) = r\alpha - \lambda + \lambda b(\alpha)$. The mean drift of this process is $\mathbb{E}X_1 = \lambda\, \mathbb{E}B - r$, which we assume to be negative to ensure stability.

Clearly, if the depletion rate r were negative, and the jobs were i.i.d. samples from a non-positive distribution (i.e. the jumps were downward), then the resulting process would be spectrally negative.

It is instructive to express the compound Poisson process in terms of a triplet (d, σ^2, Π). Obviously, because of the lack of a Brownian term, $\sigma^2 = 0$. In addition, for the Lévy measure we have $\Pi(\mathrm{d}x) = \lambda \mathbb{P}(B \in \mathrm{d}x)$. It is then readily verified that

$$d = -r + \lambda \int_0^1 x \Pi(\mathrm{d}x).$$

(3) *Gamma process*. This process is characterized by the characteristic triplet (d, σ^2, Π), where $\sigma^2 = 0$ and

$$\Pi(\mathrm{d}x) = \frac{\beta}{x} e^{-\gamma x} \mathrm{d}x \quad \text{for } x > 0, \qquad d = \int_0^1 x \Pi(\mathrm{d}x),$$

for $\gamma, \beta > 0$. From the above formulation it is clear that the jumps of this process are non-negative, that is, the gamma process is spectrally positive. In fact its sample paths are non-decreasing a.s.; we return to this property below.

The Laplace exponent corresponding to the gamma process can be evaluated explicitly, but this requires some non-standard computations. These rely on the well-known *Frullani integral*: for $z \in \mathbb{C}$ with non-positive real part,

$$\beta \log\left(1 - \frac{z}{\gamma}\right) = \int_0^\infty (1 - e^{zx}) \frac{\beta}{x} e^{-\gamma x} \mathrm{d}x; \tag{2.3}$$

see e.g. Kyprianou [146, Lemma I.1.7]. The validity of Eqn. (2.3) is a direct consequence of the identity (given that appropriate regularity conditions are imposed on the function $f(\cdot)$)

$$\int_0^\infty \frac{f(ax) - f(bx)}{x} \mathrm{d}x = -\int_0^\infty \int_a^b f'(xy) \mathrm{d}y \, \mathrm{d}x = -\int_a^b \int_0^\infty f'(xy) \mathrm{d}x \, \mathrm{d}y$$
$$= \int_a^b \frac{f(0) - f(\infty)}{y} \mathrm{d}y = (f(0) - f(\infty)) \log \frac{b}{a}$$

by picking $f(x) := e^{-x}$, $a = \gamma$, and $b = \gamma - z$.

As a consequence of the above computations, it follows that the corresponding Laplace exponent

$$\varphi(\alpha) = \log \mathbb{E} e^{-\alpha X_1} = -\alpha \int_0^1 x \Pi(\mathrm{d}x) + \int_{-\infty}^\infty (e^{-\alpha x} - 1 + \alpha x 1_{[0,1)}(|x|)) \Pi(\mathrm{d}x),$$

can now be rewritten as

$$\int_0^\infty (e^{-\alpha x} - 1) \frac{\beta}{x} e^{-\gamma x} \mathrm{d}x = \beta \log\left(\frac{\gamma}{\gamma + \alpha}\right).$$

From the equation

$$\int_0^\infty \left(\gamma e^{-\gamma x} \frac{(\gamma x)^{\beta t - 1}}{\Gamma(\beta t)}\right) e^{-\alpha x} \mathrm{d}x$$
$$= \left(\frac{\gamma}{\gamma + \alpha}\right)^{\beta t} \int_0^\infty (\gamma + \alpha) e^{-(\gamma+\alpha)x} \frac{((\gamma + \alpha)x)^{\beta t - 1}}{\Gamma(\beta t)} \mathrm{d}x = \left(\frac{\gamma}{\gamma + \alpha}\right)^{\beta t},$$

where $\Gamma(z) := \int_0^\infty e^{-x} x^{z-1} \mathrm{d}x$ denotes the gamma function, it follows that the marginals X_t have a gamma distribution with parameters γ and βt. We write throughout this monograph $X \in \mathbb{G}(\gamma, \beta)$.

The gamma process has interesting qualitative properties. Observe that X_t has the same distribution as the sum of X_s and X_{t-s} (with $s \in (0, t)$), with the latter two random variables being sampled independently, which are both non-negative random variables. From this we conclude that $(X_t)_t$ is a non-decreasing process.

In the second place, it is observed that the gamma process is *not* compound Poisson. This is a consequence of the fact that we cannot write $\Pi(\mathrm{d}x)$ as $\lambda\,\mathbb{P}(B \in \mathrm{d}x)$. To see this, realize that, as a consequence of $\beta/x \cdot e^{-\gamma x}$ being roughly β/x for x close to 0,

$$\int_0^\infty \Pi(\mathrm{d}x) = \int_0^\infty \frac{\beta}{x} e^{-\gamma x} \mathrm{d}x = \infty,$$

and hence it is not possible to properly define a (finite) jump intensity λ. Indeed, the gamma process is a Lévy process with the remarkable property that it has infinitely many jumps (almost surely) in any finite amount of time. We refer to this phenomenon by saying that the gamma process has *small jumps*, or, equivalently, *infinite activity*.
As mentioned above, the gamma process is increasing; to make sure that $\mathbb{E}X_1 < 0$ (so as to guarantee that the corresponding workload process is stable) a negative drift has to be added.

(4) *Inverse Gaussian process.* Like the gamma process, this process is increasing. It is defined as follows. For any $X \in \mathscr{S}_+$, we define the *first passage time*,

$$\tau(x) := \inf\{t \geq 0 : X_t < -x\};$$

this is a notion that will play an important role later in this book. It is straightforward to observe that $e^{-\varphi(\alpha)t}\,e^{-\alpha X_t}$ is a mean-1 martingale [220]: for all $s \leq t$, using the properties of Lévy processes,

$$\begin{aligned}
&\mathbb{E}\left(e^{-\varphi(\alpha)t}\,e^{-\alpha X_t} \mid \{e^{-\varphi(\alpha)u}\,e^{-\alpha X_u} : u \leq s\}\right)\\
&= \mathbb{E}\left(e^{-\varphi(\alpha)t}\,e^{-\alpha X_t} \mid \{X_u : u \leq s\}\right)\\
&= e^{-\varphi(\alpha)s}\,e^{-\alpha X_s}\,\mathbb{E}\left(e^{-\varphi(\alpha)(t-s)}\,e^{-\alpha X_{t-s}}\right) = e^{-\varphi(\alpha)s}\,e^{-\alpha X_s}.
\end{aligned}$$

Considering $X \in \mathbb{B}\mathrm{m}(d, \sigma^2)$, clearly $d < 0$ implies $\tau(x) < \infty$ almost surely. The a.s. continuous sample paths imply that $X_{\tau(x)} = -x$, which, together with 'optional sampling' [220, Chapter A14], leads to

$$\mathbb{E}e^{-\varphi(\alpha)\tau(x)} = e^{-\alpha x}.$$

As a consequence, replacing $\varphi(\alpha)$ by ϑ (and hence α by $\psi(\vartheta)$),

$$\mathbb{E}e^{-\vartheta\tau(x)} = \exp\left(-\left(\frac{d}{\sigma^2} + \sqrt{\left(\frac{d}{\sigma^2}\right)^2 + 2\frac{\vartheta}{\sigma^2}}\right)x\right).$$

Conclude that $\tau(x)$ is an increasing Lévy process (in x); the class of these processes we call *inverse Gaussian*, and we denote it by $\mathbb{IG}(d, \sigma^2)$. Again, to have $\mathbb{E}X_1 < 0$, a negative drift is added. The identification of the spectral measure $\Pi(\cdot)$ is the subject of one of the exercises. The inverse Gaussian process has 'small jumps', too: it experiences an infinite number of jumps (almost surely) over any time interval of finite length.

2.3 α-Stable Lévy Motions

This section focuses on a subclass of Lévy processes that has attracted substantial attention in the literature: α-stable Lévy motions. This class of processes is particularly suitable when modeling various sorts of heavy-tailed phenomena [192].

To introduce α-stable Lévy motions, we first define the class of stable distributions. Here we follow the exposition in Samorodnitsky and Taqqu [192], but various other parameterizations are possible [213]. We say that a random variable Y has a stable distribution if for any $a, b > 0$ there exist $c > 0$ and $d \in \mathbb{R}$ such that

$$aY' + bY'' \stackrel{\mathrm{d}}{=} cY + d,$$

where Y' and Y'' are independent copies of Y. Due e.g. to Bingham et al. [47, Thm. 8.3.2], it turns out that the characteristic function of Y can be written in the form

$$\log \mathbb{E}e^{\mathrm{i}\theta Y} = \begin{cases} -\sigma^\alpha|\theta|^\alpha(1 - \mathrm{i}\beta \mathrm{sign}(\theta)\tan(\pi\alpha/2)) + \mathrm{i}m\theta, & \alpha \neq 1; \\ -\sigma|\theta|(1 + \mathrm{i}\beta\pi/2 \mathrm{sign}(\theta)\log|\theta|) + \mathrm{i}m\theta, & \alpha = 1; \end{cases}$$

where $\alpha \in (0, 2]$, $\beta \in [-1, 1]$, $\sigma \in [0, \infty)$, $m \in \mathbb{R}$, and $\mathrm{sign}(x) := 1_{(0,\infty)}(x) - 1_{(-\infty,0)}(x)$. We write that Y is distributed $S_\alpha(\sigma, \beta, m)$.

Let us consider the meaning of the parameters in more detail.

- The parameter α is commonly referred to as the *index of stability*. Later we will observe that α is directly related to the 'heaviness' of the tail distribution. In particular, if $\alpha \in (0, 1]$, then $\mathbb{E}|Y| = \infty$ (for $\alpha = 1$ we have the Cauchy distribution). For $\alpha = 2$ we obtain the normal distribution.
- The parameter β is known as the *skewness*. The extreme cases are $\beta = 1$, corresponding to a *totally skewed to the right* distribution, and $\beta = -1$, which corresponds to a *totally skewed to the left* distribution. For $\alpha < 1$, $m = 0$, and $\beta = 1$ (respectively, $\beta = -1$), the support of the distribution is the positive (respectively, negative) half-line, but this is no longer true for $\alpha \geq 1$. The choice of $m = 0$ and $\beta = 0$ leads to a symmetric distribution.
- For obvious reasons, σ is called the *scale parameter*.
- For $\alpha \in (1, 2]$, we have $\mathbb{E}Y = m$. This explains why m is called the *shift parameter*.

The following useful property, describing the distribution's tail asymptotics, can be found in e.g. Samorodnitsky and Taqqu [192, p. 16]. As before, $\Gamma(z) := \int_0^\infty e^{-x}x^{z-1}\mathrm{d}x$ denotes the gamma function. Also, $f(x) \sim g(x)$ as $x \to \infty$ means that $f(x)/g(x) \to 1$ as $x \to \infty$.

Proposition 2.1 *Let $Y \stackrel{\mathrm{d}}{=} S_\alpha(\sigma, \beta, m)$ with $\beta \in (-1, 1]$. Then, as $u \to \infty$,*

$$\mathbb{P}(Y > u) \sim C_{\alpha,\sigma}\left(\frac{1+\beta}{2}\right)u^{-\alpha},$$

where

$$C_{\alpha,\sigma} := \begin{cases} \sigma^\alpha(1-\alpha)/\left(\Gamma(2-\alpha)\cos(\pi\alpha/2)\right) & \text{if } \alpha \neq 1; \\ 2\sigma/\pi & \text{if } \alpha = 1. \end{cases}$$

The case $\beta = -1$ has to be treated separately; see e.g. [192, pp. 17–18].

Proposition 2.2 *Let $Y \stackrel{\mathrm{d}}{=} S_\alpha(\sigma, -1, 0)$.*

(i) If $\alpha = 1$, then as $u \to \infty$,

$$\mathbb{P}(Y > u) \sim \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{(\pi/2\sigma)u - 1}{2} - e^{(\pi/2\sigma)u-1}\right).$$

(ii) If $\alpha > 1$, then as $u \to \infty$,

$$\mathbb{P}(Y > u) \sim \frac{1}{\sqrt{2\pi\alpha(\alpha-1)}}\left(\frac{\alpha\hat{\sigma}_\alpha}{u}\right)^{\alpha/(2(\alpha-1))}\exp\left(-(\alpha-1)\left(\frac{u}{\alpha\hat{\sigma}_\alpha}\right)^{\alpha/(\alpha-1)}\right),$$

where

$$\hat{\sigma}_\alpha := \sigma\left(\cos\left(\frac{\pi(2-\alpha)}{2}\right)\right)^{-1/\alpha}.$$

Having defined stable distributions, we can now introduce α-stable Lévy motions, as follows. We say that $(X_t)_t$ is an α-stable Lévy motion if $(X_t)_t$ has stationary and independent increments such that the marginals obey

$$X_t \stackrel{\mathrm{d}}{=} S_\alpha(t^{1/\alpha}, \beta, mt);$$

we write $X \in \mathbb{S}(\alpha, \beta, m)$. From the above we conclude that if $\beta = \pm 1$, then $X \in \mathscr{S}_\pm$.

For given $\alpha \in (0, 2]$ the Lévy measure has the form, for $A, B > 0$,

$$\Pi(\mathrm{d}x) = \left(\frac{A}{(-x)^{\alpha+1}}1_{\{x<0\}} + \frac{B}{x^{\alpha+1}}1_{\{x>0\}}\right)\mathrm{d}x;$$

it is verified that we again have the property of infinitely many jumps in any finite time interval, almost surely.

One could say that α-stable Lévy motions are *self-similar*: picking $m = 0$, and writing $(X_t^{(\alpha)})_t$ to stress the dependence on α, one has that, for $M > 0$,

$$\left(X_{Mt}^{(\alpha)}\right)_t \stackrel{\mathrm{d}}{=} \left(M^{1/\alpha} X_t^{(\alpha)}\right)_t$$

(unless $\alpha = 1$, $\beta \neq 0$). In other words, when zooming in, one sees essentially the same pattern, given that one adjusts the axes in a suitable fashion.

2.4 Lévy-Driven Queues

Having defined Lévy processes, in this section we introduce the notion of queues with Lévy input (or *Lévy-driven queues*). It is noticed, however, that these definitions are by no means restricted to the Lévy framework; based on the formalism defined below, one can define for *any* real-valued stochastic process the corresponding workload process. We provide two types of characterizations.

In the first approach, we define the Lévy-driven queue as the continuous-time counterpart of the classical discrete-time queue. In discrete time, a queue can be described through the well-known Lindley recursion: we have that the workload process (Q_n) satisfies

$$Q_{n+1} = \max\{Q_n + Y_n, 0\},$$

where Y_n is the *net* input to the queue in slot n (i.e. the input minus the amount that can potentially be served). Iterating this recursion, we obtain $Q_{n+1} = \max\{Q_{n-1} + Y_{n-1} + Y_n, Y_n, 0\}$. With $X_n := \sum_{i=0}^n Y_i$, and with $Q_0 = x$ for $x \geq 0$, this eventually leads to

$$Q_n = X_n + \max\left\{x, \max_{0 \leq i \leq n} -X_i\right\}.$$

In this way we have written the workload process $(Q_n)_n$ as a functional of the cumulative net input process $(X_n)_n$, and now the idea is to use the very same functional to define the workload in continuous time.

More concretely, a queue in continuous time can be defined by just taking the continuous-time analogue of the above, so that we obtain

$$Q_t = X_t + \max\{x, L_t\}, \;\; t \geq 0, \tag{2.4}$$

with

$$L_t := \sup_{0 \leq s \leq t} -X_s = - \inf_{0 \leq s \leq t} X_s;$$

this increasing (and therefore of bounded variation) process L_t is often referred to as *local time* (*at zero*) or a *regulator process*; see e.g. Harrison [108]. Assuming the queue has been running from $-\infty$, one can alternatively write

$$Q_t = \sup_{s \le t}(X_t - X_s).$$

To ensure the existence of a stationary distribution, it is evident that a stability condition needs to be fulfilled. In the case of input processes $(X_t)_t$ with stationary increments (as is the case in our Lévy context) it needs to be assumed that $\mathbb{E}X_1 < 0$ for the workload process to be stable (which we do throughout this book). If the input process X_t is *reversible*, that is, $(X_{(s-t)_-} - X_s)_t \stackrel{\mathrm{d}}{=} (-X_t)_t$ for each given $s > 0$ (which is true in the Lévy case), then we have the following distributional equality for the stationary workload Q, commonly attributed to Reich [182]:

$$Q \stackrel{\mathrm{d}}{=} \sup_{t \ge 0} X_t. \tag{2.5}$$

Above we constructed the Lévy-driven queue in continuous time analogously to its discrete-time counterpart. An alternative way of introducing Lévy-driven queues is by defining them as the solution of a so-called *Skorokhod problem*, as introduced by Skorokhod in [201, 202]; then one commonly says that $(Q_t)_t$ is *the reflection of* $(X_t)_t$ *at* 0. This is done as follows. Let $(L_t^\star)_t$ be a non-decreasing right-continuous process such that the following two requirements are fulfilled.

(A) The workload process $(Q_t)_t$, defined through $Q_0 := x$ and $Q_t := X_t + L_t^\star$, is non-negative for all $t \ge 0$.

(B) $L_t^\star$ can only increase when $Q_t = 0$, that is,

$$\int_0^T Q_t \mathrm{d}L_t^\star = 0, \quad \text{for all } T > 0.$$

Observe that it is natural to impose these conditions on a queueing process. The process $(L_t^\star)_t$ can be informally thought of as the *cumulative idle time process*; then (A) indicates by how much X_t should be increased to obtain Q_t (to account for the effect of the boundary at 0), and (B) entails that it is not possible that at the same time the queue is non-empty and the cumulative idle time grows.

Importantly, it can be proved that the only process satisfying these two conditions is $L_t^\star = \max\{x, L_t\}$, so that $Q_t = X_t + \max\{x, L_t\}$ for $t \ge 0$, where L_t is defined as above; see e.g. Asmussen [19, Prop. IX.2.2] and Robert [185, p. 375]. We conclude that the expression found in this way coincides with the one obtained when taking the continuous counterpart of the discrete-time definition, as in (2.4). For the sake of completeness we include the proof here.

Proposition 2.3 *The process $(L_t^\star)_t$, defined by $L_t^\star := \max\{x, L_t\}$, is the unique solution to the Skorokhod problem (A)–(B).*

Proof There are several ways to prove the statement; we follow the proof of [19, Prop. IX.2.2]. Let $(\bar{L}_t^\star)_t$ be another solution to (A)–(B), and $(\bar{Q}_t)_t$ be the corresponding workload process. Defining $D_t := \bar{L}_t^\star - L_t^\star$, it is our goal to verify that necessarily $D_t \equiv 0$. By applying integration by parts for right-continuous processes of bounded variation, and defining $\Delta D_s := D_s - D_{s-}$,

$$
\begin{aligned}
D_t^2 &= 2\int_0^t D_s \mathrm{d}D_s - \sum_{s\le t} (\Delta D_s)^2 \\
&= 2\int_0^t (\bar{L}_s^\star - L_s^\star)\mathrm{d}\bar{L}_s^\star - 2\int_0^t (\bar{L}_s^\star - L_s^\star)\mathrm{d}L_s^\star - \sum_{s\le t} (\Delta D_s)^2 \\
&= 2\int_0^t (\bar{Q}_s - Q_s)\mathrm{d}\bar{L}_s^\star - 2\int_0^t (\bar{Q}_s - Q_s)\mathrm{d}L_s^\star - \sum_{s\le t} (\Delta D_s)^2,
\end{aligned}
$$

where the last step is due to $X_t = Q_t - L_t^\star = \bar{Q}_t - \bar{L}_t^\star$. Realizing that

$$
\int_0^t \bar{Q}_s \mathrm{d}\bar{L}_s^\star = \int_0^t Q^\star \mathrm{d}L_s^\star = 0,
$$

it follows that

$$
D_t^2 = -2\int_0^t Q_s \mathrm{d}\bar{L}_s^\star - \int_0^t \bar{Q}_s \mathrm{d}L_s^\star - \sum_{s\le t} (\Delta D_s)^2.
$$

As Q_s and $\bar{Q}_s$ are non-negative, we conclude that $D_t^2 \le 0$, and therefore $D_t = 0$. □

In the case $X \in \mathbb{C}\mathrm{P}(r, \lambda, b(\cdot))$, the queue under study is the well-known M/G/1 queue. We refer to Fig. 2.3 for a pictorial illustration of the evolution of the workload in time, jointly with the $(X_t)_t$ process (where we consider for ease the special case of $Q_0 = 0$ and $r = 1$). It is elementary to verify that in the case that

$$
\arg \inf_{0\le s\le t} (X_t - X_s)
$$

is smaller than t, this time epoch can be interpreted as the start of the busy period in which t is contained; if it equals t (meaning that X_t is the 'all-time low' of the process so far), then the workload is 0 at time t. It also follows that in this context, the process $L_t^\star$ is the queue's cumulative idle time up to time t.

Importantly, however, we would like to stress that this general notion of a queueing system can be used in settings beyond traditional queues: the process $(X_t)_t$ does not need necessarily to relate to positive quantities of work arriving. In this sense, we now have developed the concept of a queue fed by for instance Brownian

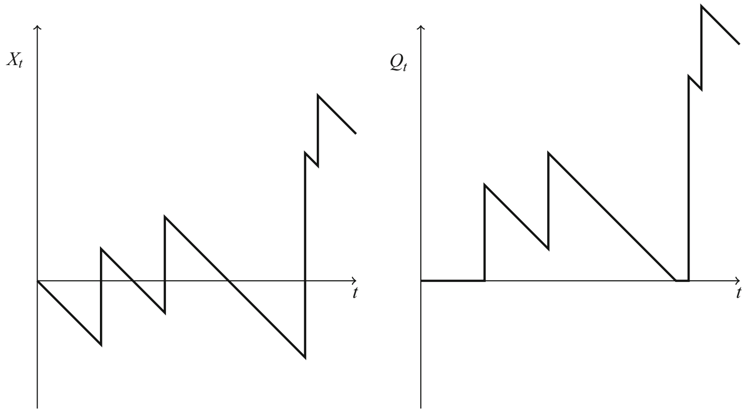

Fig. 2.3 Net input process and workload process for a compound Poisson process

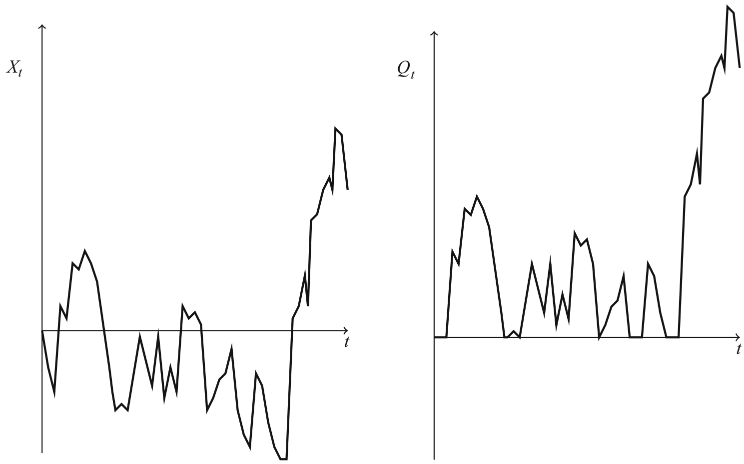

Fig. 2.4 Net input process and workload process for an erratic, 'Brownian-like' process

motion, or any other real-valued continuous-time stochastic process. In the case $X \in \mathbb{B}\mathrm{m}(d, \sigma^2)$, the resulting workload process is often referred to as *reflected* (or *regulated*) *Brownian motion*. We refer to Fig. 2.4 for an illustrative example of such a workload process.

One of the main objectives in this book is the identification of the distribution of the transient workload Q_t and its stationary counterpart $Q := \lim_{t\to\infty} Q_t$. Note that due to (2.5), as $t \uparrow \infty$,

$$\bar{X}_t := \sup_{0\leq s\leq t} X_s \uparrow \sup_{s\geq 0} X_s \stackrel{\mathrm{d}}{=} Q.$$

Likewise, $(Q_0 \mid Q_{-t} = 0)$ increases to Q as t goes to ∞. In operations research the steady-state workload is the natural performance metric when studying queueing systems that are in operation over long periods of time.

A second frequently used performance measure is the so-called *busy period*, to be denoted by τ, being the time it takes for the queue to drain (starting from time 0):

$$\tau := \inf\{t \geq 0 : Q_t = 0\}.$$

In this book we study the busy period in detail, where we typically assume that the workload is in stationarity at time 0. Several other metrics are analyzed as well, such as the workload correlation function $\mathbb{C}\text{orr}(Q_0, Q_t)$ and the infimum attained by the workload process over a time interval of length t, that is, $\inf_{s\in[0,t]} Q_s$, in both cases assuming the workload is in stationarity at time 0.

Exercises

Exercise 2.1 Prove the Frullani integral equality, Eqn. (2.3), for $z \in \mathbb{C}$ with non-positive real part.

Hint: In the text a rough sketch was provided. Consider first $z \leq 0$. Use that

$$\frac{e^{-\gamma x} - e^{-(\gamma - z)x}}{x} = \int_{\gamma}^{\gamma - z} e^{-yx} \mathrm{d}y,$$

and then change the order of integration. Finally, by analytic extension, show that the formula is valid for $z \in \mathbb{C}$ with non-positive real part.

Exercise 2.2 Consider $X \in \mathbb{I}\mathrm{G}(-1, 1)$. Prove that

$$\Pi(\mathrm{d}x) = \frac{1}{\sqrt{2\pi x^3}} e^{-x/2}.$$

Exercise 2.3 Let $X \in \mathbb{S}(\alpha_1, \beta_1, m)$ and $Y \in \mathbb{S}(\alpha_2, \beta_2, m_2)$ be independent.

(a) Check that X_1 is infinitely divisible.
(b) Characterize when $Z_t = X_t + Y_t$ has a stable distribution. Find the parameters of Z_t.
(c) Assume that $m_1 = 0$ and check that X is self-similar, that is, show that

$$(X_{Mt})_t \stackrel{\mathrm{d}}{=} (M^{1/\alpha_1} X_t)_t.$$

(d) Characterize γ for which $\mathbb{E}(X_1)^{\gamma} < \infty$.

Exercise 2.4 Let $X \stackrel{d}{=} S_\alpha(\sigma, \beta, m)$ with $\alpha \in (1, 2)$. Check that

(a) $aX \stackrel{d}{=} S_\alpha(|a|\sigma, \text{sign}(a)\beta, am)$, for $a \neq 0$;
(b) $-X \stackrel{d}{=} S_\alpha(\sigma, -\beta, -m)$;
(c) X is symmetric if and only if $\beta, m = 0$.

Exercise 2.5 Let $X \stackrel{d}{=} S_\alpha(\sigma, \beta, 0)$ with $\alpha \in (1, 2)$. In addition, we have the processes $X^{(1)} \stackrel{d}{=} S_\alpha(\sigma, 1, 0)$, $X^{(2)} \stackrel{d}{=} S_\alpha(\sigma, -1, 0)$, which we assume to be mutually independent. Check that

$$X \stackrel{d}{=} \left(\frac{1+\beta}{2}\right)^{1/\alpha} X^{(1)} + \left(\frac{1-\beta}{2}\right)^{1/\alpha} X^{(2)}.$$

Exercise 2.6 Prove that the sum of independent compound Poisson processes is a compound Poisson process. Find its parameters.

Exercise 2.7 Let X and Y be two independent Lévy processes; assume Y is increasing.

(a) Show that $(X_{Y_t})_{t \geq 0}$ is a Lévy process as well.
(b) Let X be a (standard) Brownian motion, and $Y \in \mathbb{G}(\beta, \gamma)$. Determine the Lévy exponent of $(X_{Y_t})_{t \geq 0}$.

(*Note*: With a specific choice of the parameters, this process is called a *variance gamma process*; see also Chapter 15.)

Exercise 2.8 Prove Prop. 2.1.

Exercise 2.9 For a given Lévy process X with $\mathbb{E}\, X_1 < 0$, let Q_0 obey the stationary workload distribution, and let L be the regulator process, with

$$L_t := \sup_{0 \leq s \leq t} -X_s = -\inf_{0 \leq s \leq t} X_s.$$

Then, according to the definition of the workload process, for $t \geq 0$,

$$Q_t = X_t + \max\{Q_0, L_t\}.$$

Show that $Q_t = \sup_{-\infty < s \leq t}(X_t - X_s)$, and that Q_t is stationary.

http://www.springer.com/978-3-319-20692-9

Queues and Lévy Fluctuation Theory
Debicki, K.; Mandjes, M.
2015, XI, 255 p. 12 illus., Softcover
ISBN: 978-3-319-20692-9