

Sojourn Time Asymptotics in Processor-Sharing Queues

Sem Borst^{†,*‡}, Rudesindo Núñez-Queija^{†,*}, Bert Zwart^{†,*}

[†]CWI

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

*Department of Mathematics & Computer Science

Eindhoven University of Technology

P.O. Box 513, 5600 MB Eindhoven, The Netherlands

[‡]Bell Laboratories, Lucent Technologies

P.O. Box 636, Murray Hill, NJ 07974, USA

Abstract

Over the past few decades, the Processor-Sharing (PS) discipline has attracted a great deal of attention in the queueing literature. While the PS paradigm emerged in the sixties as an idealization of round-robin scheduling in time-shared computer systems, it has recently captured renewed interest as a useful concept for modeling the flow-level performance of bandwidth-sharing protocols in communication networks. In contrast to the simple geometric queue length distribution, the sojourn time lacks such a nice closed-form characterization, even for exponential service requirements. In case of heavy-tailed service requirements however, there exists a simple asymptotic equivalence between the sojourn time and the service requirement distribution, which is commonly referred to as a reduced service rate approximation. In the present survey paper, we give an overview of several methods that have been developed to obtain such an asymptotic equivalence under various distributional assumptions. We outline the differences and similarities between the various approaches, discuss some connections, and present necessary and sufficient conditions for an asymptotic equivalence to hold. We also consider the generalization of the reduced service rate approximation to several extensions of the M/G/1 PS queue. In addition, we identify a relationship between the reduced service rate approximation and a queue length distribution with a geometrically decaying tail, and extend it to so-called bandwidth-sharing networks. The state-of-the art with regard to sojourn time asymptotics in PS queues with light-tailed service requirements is also briefly described. Last, we reflect on some possible avenues for further research.

2000 AMS Subject Classification: 60K25 (primary), 60F10, 68M20, 90B18, 90B22 (secondary).

Keywords: large deviations, tail asymptotics, processor sharing, regular variation, subexponential distributions, heavy-tailed distributions, reduced service rate approximation, light-tailed distributions, bandwidth-sharing networks.

1 Introduction

Over the past few decades, the Processor-Sharing (PS) discipline has gained a prominent role in queueing theory. In a PS system, the total service rate is equally shared among all users present. Thus, when there are $n \geq 1$ users present, each of them receives service at rate $1/n$. Originally, the PS paradigm emerged as an idealization of round-robin scheduling mechanisms in time-shared computer systems [34, 35]. In recent years, the PS discipline has attracted renewed interest as a convenient abstraction for modeling the flow-level performance of bandwidth-sharing protocols in packet-switched communication networks [29, 37, 42, 43].

From a practical point of view, one of the most appealing properties of fair-sharing strategies such as PS is that they prevent large jobs from hogging the server, and thus avoid that small jobs get stuck behind large ones. As a result, fair-sharing strategies are particularly suitable for dealing with the negative impact of high variability in job sizes. These features are reflected in the basic theoretical properties of the PS discipline. In particular, in case of Poisson arrivals, the stationary queue length distribution is known to have a simple geometric distribution that only depends on the service requirement distribution through its mean, and not through any higher-order statistics [52, 53]. Because of Little's law, the insensitivity of the queue length distribution translates into insensitivity of the mean sojourn time. In addition, the *conditional* mean sojourn time is known to be proportional to the actual service requirement of a user [34, 35, 53], embodying a certain fairness principle.

Besides the mean sojourn time, the distribution of the sojourn time also constitutes a highly relevant performance measure. In contrast to the simple geometric distribution of the queue length however, the sojourn time distribution does not have any simple characterization. Initiated by Kleinrock's analysis of the M/M/1 PS queue [34, 35], many studies in the literature have focused on the analysis of the sojourn time conditioned on the service requirement. Extensions to generally distributed service requirements, multiple servers, and more general sharing disciplines were pursued in [20, 52, 53]. However, determining the sojourn time *distribution* in PS queues turned out to be a rather challenging problem. For the M/M/1 PS queue, Coffman *et al.* [19] first derived a closed-form expression for the LST of the sojourn time distribution conditioned on the service requirement and the number of customers seen upon arrival. Sengupta & Jagerman [54] found an alternative expression for the LST of the distribution of the sojourn time conditioned only on the number of customers seen upon arrival. Building on [19], Morrison [40] established an expression for the *distribution* function of the sojourn time. For results on the sojourn time distribution in M/G/1 PS queues, we refer to the survey papers [55, 56].

The sojourn time distribution in G/M/1 PS queues has received less attention in the literature. Ramaswami [51] characterized the LST of the sojourn time distribution by a differential equation and determined the first two moments of the distribution. Jagerman & Sengupta [30] gave explicit expressions for the LST, and derived a heavy-traffic limit distribution under proper scaling, showing that, in the limit, the sojourn time is distributed as the product of two independent exponentially distributed random variables. The sojourn time in the 'repair' node (with PS discipline) of the machine-repairman model was examined by Mitra [38]. Extensions to multiple customer classes, both in the moderate and in the heavy-traffic regime, were considered by Mitra & Morrison [39, 41].

As mentioned above, fair-sharing strategies as modeled by the PS discipline are particularly

attractive for alleviating the negative impact of high variability in job sizes. This capability is especially critical as traffic measurements indicate that file transfers in the Internet and document sizes on Web servers show extreme variability and commonly exhibit heavy-tailed characteristics [21]. These findings have triggered a strong interest in the delay characteristics of PS queues with heavy-tailed service requirements. In view of the poor tractability, even for exponential service requirements, most of the studies have focused on an asymptotic characterization of the tail distribution. Using different techniques and under various distributional assumptions, several papers have established the following asymptotic equivalence:

$$\mathbb{P}\{V > x\} \sim \mathbb{P}\{B > (1 - \rho)x\}, \quad (1)$$

where B denotes a generic service requirement, V denotes a generic sojourn time, and \sim denotes that the ratio of both sides converges to 1 as x tends to ∞ . The asymptotic equivalence (1) was first proved by Zwart & Boxma [59], assuming Poisson arrivals and regularly varying service requirements. Under practically the same conditions (allowing for intermediately regularly varying service distributions), a probabilistic proof was given by Núñez Queija [43, 44] allowing for scenarios with random service interruptions and other service disciplines. An important extension to the class of *subexponential concave distribution functions* was provided by Jelenković & Momčilović [32] by means of a sample-path proof technique. Notably, they showed that the result does not hold for subexponential distribution functions that are not *square-root insensitive*. For regularly varying distributions, Guillemin *et al.* [28] demonstrated that the asymptotic equivalence remains true for several model extensions including admission control and impatience.

The asymptotic equivalence (1) may be heuristically explained as follows. Consider a tagged customer with a large service requirement. The sojourn time V of the tagged customer consists of its own service requirement B plus the amount of service provided to other customers during its sojourn time. Because of the PS discipline, virtually all the work that arrives over the course of its sojourn time must also be served during its sojourn time. Thus, the amount of service provided to other customers over the course of its sojourn time will be roughly ρV , so that $V \approx B + \rho V$, or equivalently, $V \approx B/(1 - \rho)$. In particular, $B > (1 - \rho)x$ “implies” that $V > x$. Put differently, the mean service rate received by the tagged customer is approximately $1 - \rho$. Hence, the asymptotic equivalence (1) is sometimes called a *reduced service rate approximation*, in analogy with the term *reduced load equivalence* that is often used to refer to similar types of results for workload asymptotics. Observe that the asymptotic equivalence indirectly shows that the above scenario is in fact the only plausible way for a long sojourn time to occur, i.e., with overwhelming probability a long sojourn time is due to a large service requirement of the customer itself.

In the present paper, we give an overview of the above-mentioned methods that have been devised to obtain sojourn time asymptotics under various distributional assumptions. We describe the differences and similarities between the various approaches and present some new results and insights. We also consider the generalization of the reduced service rate approximation to several extensions of the M/G/1 PS queue, with features such as (i) time-varying or state-dependent service rates; (ii) multi-class versions such as Discriminatory Processor Sharing queues; (iii) renewal arrival processes. In addition, we identify a relationship between the reduced service rate approximation and a queue length distribution with a geometric tail, and extend it to so-called monotone PS networks and networks

with balanced fairness.

The remainder of the paper is organized as follows. In Section 2 we review four different methods that have been developed to derive sojourn time asymptotics in PS queues with heavy-tailed service requirements. In Section 3 we establish general necessary and sufficient conditions for the asymptotic equivalence (1) to hold. As we will show, these conditions provide a unifying framework connecting some of the methods described in Section 2. We then proceed to discuss several extensions to models with a varying service rate and multi-class settings in Section 4. In Section 5 we turn the attention to models with Discriminatory (non-egalitarian) Processor Sharing (DPS), which involve major difficulties and require fundamentally different proof techniques. In Section 6 we demonstrate an intimate relationship between the asymptotic equivalence (1) and a geometrically bounded queue length, which is illustrated in the context of a DPS system. Section 7 then examines the existence of a reduced service rate approximation for the sojourn time in bandwidth-sharing networks. In Section 8 we briefly discuss the state-of-the-art with regard to sojourn time asymptotics for PS queues with light-tailed service requirements, which turn out to be considerably more delicate than their heavy-tailed counterparts. In Section 9 we make some concluding remarks and sketch possible directions for further research.

2 Methods

In this section we review various methods that have been developed to derive sojourn time asymptotics in PS queues with heavy-tailed service requirements. Before we do so, we first provide some background on heavy-tailed distributions. A random variable X , or its distribution function, is called long-tailed ($X \in \mathcal{L}$) if $\mathbb{P}\{X > x\} \sim \mathbb{P}\{X > x - y\}$ as $x \rightarrow \infty$ for any $y > 0$. Here, and in the remainder of the paper, we write $f(x) \sim g(x)$, $x \rightarrow \infty$ (or simply $f(x) \sim g(x)$ if no ambiguity arises) whenever $\lim_{x \rightarrow \infty} \frac{f(x)}{g(x)} = 1$. X is said to be *subexponential* if $\mathbb{P}\{X_1 + \dots + X_n > x\} \sim n\mathbb{P}\{X > x\}$ with X_i , $i = 1, \dots, n$, i.i.d. copies of X . X is called regularly varying of index $\alpha \geq 0$ ($X \in \mathcal{R}_\alpha$) if $\mathbb{P}\{X > x\} = L(x)x^{-\alpha}$, with $L(ax)/L(x) \rightarrow 1$ for any $a > 0$. The function $L(\cdot)$ is called slowly varying. We set $\mathcal{R} = \cup_\alpha \mathcal{R}_\alpha$, and note that $\mathcal{R} \subseteq \mathcal{S} \subseteq \mathcal{L}$. A minor extension of the class \mathcal{R} is the class of intermediately regularly varying distributions, denoted by \mathcal{IR} . We have $X \in \mathcal{IR}$ if

$$\lim_{\epsilon \downarrow 0} \limsup_{x \rightarrow \infty} \frac{\mathbb{P}\{X > x(1 - \epsilon)\}}{\mathbb{P}\{X > x\}} = 1.$$

This definition is far from natural, but interchanging the limits w.r.t. ϵ and x is sometimes exactly what is required in technical proofs. A more natural description of \mathcal{IR} is due to D.A. Korshunov (personal communication): $X \in \mathcal{IR}$ if and only if $\mathbb{P}\{B > x\} \sim \mathbb{P}\{B > x - o(x)\}$ for any function $o(x)$ satisfying $o(x)/x \rightarrow 0$.

All definitions and classes given here are standard; further background on heavy-tailed distributions can be found in the monograph of Embrechts *et al.* [24]. Other classes of distributions which are not considered to be standard are introduced when necessary.

2.1 Tauberian approach

The asymptotic equivalence (1) was first established by Zwart & Boxma [59] for the case of a single-server PS queue with Poisson arrivals and regularly varying service requirements.

Zwart [57] used a similar method to derive the sojourn time asymptotics in a multi-class M/G/1 PS queue.

Theorem 2.1 (Zwart & Boxma [59])

Let $\nu > 1$ be non-integer and $L(\cdot)$ be a slowly varying function. Then $\mathbb{P}\{V > x\} \sim (1 - \rho)^{-\nu} x^{-\nu} L(x)$ if and only if $\mathbb{P}\{B > x\} \sim x^{-\nu} L(x)$.

Note that the asymptotic equivalence (1) immediately follows from the above theorem.

We now present a brief outline of the proof of Theorem 2.1. The approach relies on a Tauberian theorem which relates the behavior of a distribution function at infinity to the behavior of its Laplace Stieljes Transform (LST) near the origin. To be specific, let X be a random variable with distribution function $F(\cdot)$, $F(x) := \mathbb{P}\{X < x\}$, first n moments $\mu_i := \mathbb{E}\{X^i\} = \int_{x=0}^{\infty} x^i dF(x)$, $i = 0, \dots, n$, LST $\phi(s) := \mathbb{E}\{e^{-sX}\} = \int_{x=0}^{\infty} e^{-sx} dF(x)$, and

$$\phi_n(s) := (-1)^{n+1} \left[\phi(s) - \sum_{i=1}^n \mu_i \frac{(-s)^i}{i!} \right].$$

The next lemma links the behavior of $1 - F(x)$ for $x \rightarrow \infty$ to the behavior of $\phi(s)$ for $s \downarrow 0$. It provides a systematic recipe that reduces the derivation of the tail behavior of the sojourn time distribution to the characterization of the behavior of its LST around zero.

Lemma 2.2 Let $n < \nu < n + 1$ and $C \geq 0$. Then $\phi_n(s) \sim (C + o(1))s^\nu L(1/s)$ as $s \downarrow 0$ for s real if and only if

$$1 - F(x) \sim (C + o(1)) \frac{(-1)^n}{\Gamma(1 - \nu)} x^{-\nu} L(x).$$

The above lemma was originally established in [6] for the case $C > 0$. The case $C = 0$ is treated in [17]. The more complicated case when ν is integer is covered in [7].

In the analysis of sojourn times in the M/G/1 PS queue, the conditional sojourn time $V(\tau)$ of a customer with service requirement τ has played a central role. Various expressions have been obtained for its LST $v(s, \tau) := \mathbb{E}\{e^{-sV(\tau)}\}$ using different techniques. Note that the LST of the unconditional sojourn time distribution readily follows as $v(s) = \int_{\tau=0}^{\infty} v(s, \tau) dB(\tau)$. The derivation in [59] starts from the expression obtained by Ott [48]:

$$v(s, \tau) = \frac{1 - \rho}{(1 - \rho)H_1(s, \tau) + sH_2(s, \tau)},$$

with

$$\int_{\tau=0}^{\infty} e^{-r\tau} dH_1(s, \tau) = \frac{r - \lambda(1 - \beta(r))}{r - s - \lambda(1 - \beta(r))},$$

$$\int_{\tau=0}^{\infty} e^{-r\tau} dH_2(s, \tau) = \frac{\rho r - \lambda(1 - \beta(r))}{r(r - s - \lambda(1 - \beta(r)))},$$

and $\operatorname{Re} r > 0$. In order to apply Lemma 2.2, Zwart & Boxma rewrite the above expression in a more manageable power-series form:

$$v(s, \tau)^{-1} = \sum_{k=0}^{\infty} \frac{s^k}{k!} \alpha_k(\tau),$$

where $\alpha_0(\tau) = 1$, $\alpha_1(\tau) := \tau/(1 - \rho)$, and for $k \geq 2$,

$$\alpha_k(\tau) := \frac{k}{(1 - \rho)^k} \int_{x=0}^{\tau} (\tau - x)^{k-1} R^{(k-1)*}(x) dx,$$

with

$$R^{k*}(x) := (1 - \rho)^k \sum_{n=0}^{\infty} \binom{n + k - 1}{k - 1} \rho^n \tilde{B}^{n*}(x)$$

representing the k -fold convolution of the waiting-time distribution in the M/G/1 queue with the same traffic characteristics but with the First-Come First-Served discipline. The above situation is quite characteristic, in the sense that the analysis of the behavior of the LST around zero is highly problem-specific and far from straightforward in general. Also, the Tauberian approach is inherently restricted to models where the LST is available in the first place, and may not apply in case of minor modifications, even when these do not alter the qualitative characteristics of the system.

As a by-product of the derivation, Zwart & Boxma obtain the asymptotic behavior of the conditional moments of the sojourn time:

$$\mathbb{E}\{V(\tau)^k\} = (\mathbb{E}\{V(\tau)\})^k + \frac{\beta_2}{2\beta_1} \frac{\rho}{1 - \rho} \frac{k(k-1)}{(1 - \rho)^k} \tau^{k-1} + o(\tau^{k-1}), \quad (2)$$

which is of independent interest and in fact useful in the alternative proof method discussed in Section 2.3.

2.2 Sample-path large-deviations approach

Jelenković & Momčilović [32] devised a proof technique which enabled the extension of the asymptotic equivalence (1) to the class \mathcal{SC} of so-called subexponential concave distribution functions. A non-negative random variable X , or its distribution function, belongs to the class \mathcal{SC} if its hazard function $Q(x) := -\log \mathbb{P}\{X > x\}$ is eventually concave such that $Q(x)/\log x \rightarrow \infty$ as $x \rightarrow \infty$ and

$$\frac{Q(x) - Q(u)}{Q(x)} \leq \alpha \frac{x - u}{x}$$

for $x \geq x_0$, $\beta x \leq u \leq x$, where $0 < \alpha < 1$, $0 < \beta < 1$. Note that $Q(x)/\log x \rightarrow \infty$ as $x \rightarrow \infty$ implies that all moments of X are finite (which, in particular, rules out regularly varying distributions).

Examples of random variables that belong to the class \mathcal{SC} include distributions with hazard functions of the form (i) $Q(x) = c(\log x)^\gamma$, $\gamma > 1$, and (ii) $Q(x) = c(\log x)^\gamma x^\alpha$, $\gamma \geq 0$, $0 < \alpha < 1$. In particular, lognormal and Weibull distributions belong to the class \mathcal{SC} .

The main result of Jelenković & Momčilović [32] is stated in the next theorem, where B^r denotes a random variable with the distribution of the residual lifetime of B , i.e.,

$$\mathbb{P}\{B^r > x\} := \frac{1}{\mathbb{E}\{B\}} \int_{y=x}^{\infty} \mathbb{P}\{B > y\} dy.$$

Theorem 2.3 (*Jelenković & Momčilović [32]*)

Let B belong to the class \mathcal{SC} with $\alpha < 1/2$ and

$$\limsup_{x \rightarrow \infty} \frac{\mathbb{P}\{B^r > x\}}{x\mathbb{P}\{B > x\}} < \infty. \quad (3)$$

Then $\mathbb{P}\{V > x\} \sim \mathbb{P}\{B > (1 - \rho)x\}$.

The condition (3) is not very restrictive since it is satisfied as long as $x^{1+\delta}\mathbb{P}\{B > x\}$ is eventually monotonically decreasing for some $\delta > 0$.

The condition $\alpha < 1/2$ is essential. In fact, Jelenković & Momčilović [32] show that if $\mathbb{P}\{B > x\} = e^{-x^\alpha}$, $\alpha > 1/2$, then $\mathbb{P}\{B > x\} = o(\mathbb{P}\{V(1 - \rho) > x\})$ as $x \rightarrow \infty$. Loosely speaking, the asymptotic relation (1) only holds when the service requirement distribution has a tail heavier than $e^{-\sqrt{x}}$, and does not extend to the entire class of subexponential functions. The criticality of $\alpha = 1/2$ may be informally explained from the Central Limit Theorem (CLT). The reduced service rate equivalence entails that the fluctuations in the amount of work arriving during the sojourn time average out, which requires that $\alpha < 1/2$. This is related to work on sampling at subexponential times in [3, 27]. In particular, the above statement implies for the busy period P that $\mathbb{P}\{B > x\} = o(\mathbb{P}\{P(1 - \rho) > x\})$ as $x \rightarrow \infty$, if $\mathbb{P}\{B > x\} = e^{-x^\alpha}$, $\alpha > 1/2$, as had been shown before in [3].

The proof of Theorem 2.3 is quite appealing but technical. It relies on a powerful large deviations bound for sums of random variables in \mathcal{SC} , which is also developed in [32]. Instead, we provide a short outline of a simpler proof provided in [32] for the special case of intermediate regular variation with $\mathbb{E}\{B^\alpha\} < \infty$ for some $\alpha > 1$, which avoids some minor technical conditions in [44, 59].

Theorem 2.4 (*Jelenković & Momčilović [32]*)

Let the service requirement distribution be intermediately regularly varying with $\mathbb{E}\{B^\alpha\} < \infty$ for some $\alpha > 1$. Then $\mathbb{P}\{V > x\} \sim \mathbb{P}\{B > (1 - \rho)x\}$.

The proof of Theorem 2.4 consists of lower and upper bounds which asymptotically coincide. The proof in fact closely follows the heuristic arguments sketched in the introduction. Let B_0 and V_0 be the service requirement and the sojourn time, respectively, of a tagged customer arriving at time $t = 0$. Let B_i and T_i denote the service requirement and the arrival time of the i -th customer arriving after time $t = 0$. Let $L(0)$ be the number of customers in the system just before time $t = 0$, and let B_l^r denote the remaining service requirement of the l -th customer at time 0. We use the *sample-path* representation

$$V_0 = B_0 + \sum_{l=1}^{L(0)} \min\{B_l^r, B_0\} + \sum_{i=1}^{N(0, V_0)} \min\{B_i, R_0(T_i)\}, \quad (4)$$

with $N(0, t)$ denoting the number of customers arriving during the time interval $(0, t)$, and $R_0(t)$ representing the remaining service requirement of the tagged customer at time t . The above representation gives implies the following (sample-path) lower and upper bounds for the sojourn time:

Lower bound: For any $\delta > 0$,

$$(1 - \rho + \delta)V_0 \geq B_0 - U^{\rho-\delta} - Z(V_0), \quad (5)$$

with $Z(t) := \sum_{i=1}^{N(0,t)} \max\{B_i - R_0(T_i), 0\}$ and $U^{\rho-\delta} := -\inf_{t \geq 0} \left\{ \sum_{i=1}^{N(0,t)} B_i - (\rho - \delta)t \right\}$.

Upper bound: For any $\delta > 0$,

$$(1 - \rho - \delta)V_0 \leq B_0 + \sum_{l=1}^{L(0)} \min\{B_l^r, B_0\} + W_{B_0}^{\rho+\delta}, \quad (6)$$

where, for any $y > 0$ and $c > \rho$, we define $W_y^c := \sup_{t \geq 0} \left\{ \sum_{i=1}^{N(0,t)} \min\{B_i, y\} - ct \right\}$.

These two bounds provide the necessary ingredients for the proof of Theorem 2.4.

Proof of Theorem 2.4

(Lower bound) From (5) we obtain

$$\mathbb{P}\{V_0 > x\} \geq \mathbb{P}\{B_0 > (1 - \rho + \delta + 2\epsilon)x\} \left(\mathbb{P}\{Z(V_0) \leq \epsilon x\} - \mathbb{P}\{U^{\rho-\delta} > \epsilon x\} \right).$$

Because of the law of large numbers, $\mathbb{P}\{U^{\rho-\delta} > \epsilon x\} \rightarrow 0$ as $x \rightarrow \infty$ and as observed in [32], $\mathbb{P}\{Z(V_0) \leq \epsilon x\} \rightarrow 1$ as $x \rightarrow \infty$.

(Upper bound) Using (6) we find

$$\mathbb{P}\{V_0 > x\} \leq \mathbb{P}\{B_0 + \sum_{l=1}^{L(0)} \min\{B_l^r, B_0\} > (1 - \rho - \delta - \epsilon)x\} + \mathbb{P}\{W_{B_0}^{\rho+\delta} > \epsilon x\}.$$

As demonstrated in [32], $\mathbb{P}\{B_0 + \sum_{l=1}^{L(0)} \min\{B_l^r, B_0\} > x\} \sim \mathbb{P}\{B > x\}$ and $\mathbb{P}\{W_{B_0}^{\rho+\delta} > \epsilon x\} = o(\mathbb{P}\{B > x\})$. Letting $\delta, \epsilon \downarrow 0$ and using that B is intermediately regularly varying then completes the proof. \square

2.3 Probabilistic approach using the conditional sojourn time

Núñez-Queija [44] developed a probabilistic proof technique based on properties of the sojourn time conditional on the customer's service requirement. The approach allowed an extension of the asymptotic equivalence (1) to the case of intermediately regularly varying service requirements. A convenient property of any intermediately regularly varying distribution function $F(\cdot)$ is the existence of $\zeta \in (0, \infty)$, $\eta \in (0, 1)$ and $x_0 \in (0, \infty)$ such that $(1 - F(x_2))/(1 - F(x_1)) \geq \eta(x_2/x_1)^{-\zeta}$ for all $x_2 \geq x_1 \geq x_0$. This property states that asymptotically $F(\cdot)$ is bounded from below by a regularly varying function with index ζ . We may also write $x^{-\zeta} = O(1 - F(x))$.

The main result in Núñez Queija [44] provides sufficient conditions in terms of the conditional sojourn time for an asymptotic equivalence of the form (1) to hold. The result facilitated a proof of the tail equivalence for other service disciplines as well, including FBPS (*Foreground-Background Processor Sharing*; also known as *Least Attained Service first*) and SRPT (*Shortest Remaining Service Time first*), as well as PS queues with an unreliable server [43, 44].

For the formulation of the theorem, $V(\tau)$ can be thought of as a stochastic process for which we are interested in the tail distribution at a random stopping time $\tau = B$. (A similar starting point is taken in the approach discussed in Section 2.4.) When applied to

sojourn times in the M/G/1 PS queue, it is not difficult to check that, until departure, the rate at which service is received is independent of the customer's service requirement. As a further remark, the third condition in the theorem is trivially satisfied by the conditional sojourn times in many queueing models.

Theorem 2.5 (Núñez-Queija [44])

Let $B \in \mathcal{IR}$. Assume $\inf\{\alpha \leq \infty : \mathbb{E}\{B^\alpha\} = \infty\} \neq 2$ and choose $\zeta > 0$ such that $x^{-\zeta} = O(\mathbb{P}\{B > x\})$, $x \rightarrow \infty$. Let the process $V(\tau)$ be independent of B and satisfy the following three conditions:

- (i) $\lim_{\tau \rightarrow \infty} \frac{\mathbb{E}\{V(\tau)\}}{\tau} = \frac{1}{\gamma}$, for some $\gamma > 0$;
- (ii) There exist $\kappa > \zeta$ and $\delta > 0$ such that $\mathbb{P}\{V(\tau) - \mathbb{E}\{V(\tau)\} > t\} \leq \frac{h(\tau)}{t^\kappa}$, for all $\tau \geq 0$ and $t > 0$, with $h(\tau) = o(\tau^{\kappa-\delta})$.
- (iii) $V(\tau)$ is stochastically non-decreasing in τ , i.e., for all $x \geq 0$, $\mathbb{P}\{V(\tau) > x\}$ is non-decreasing in $\tau \geq 0$.

Then $\mathbb{P}\{V > x\} \equiv \mathbb{P}\{V(B) > x\} \sim \mathbb{P}\{B > \gamma x\}$.

Remark. Notice that, in the theorem, we imposed separate conditions on B and $V(\cdot)$. For the application that we have in mind, any distributional condition on B also affects the statistical properties of the process $V(\tau)$, since the latter are partly determined by other customers' service requirements. However, for the proof of the theorem, the relation between B and $V(\tau)$ is not used. Therefore, the theorem also applies to a scenario where the customer for which we study its sojourn time, may have a different service requirement distribution than all other customers.

To prove (1) it suffices to check the conditions of Theorem 2.5 with $\gamma = 1 - \rho$. In [43, 44] it was shown that for PS, FBPS, SRPT and PS with random service interruptions, the conditions can be verified by using the Chebyshev-Markov inequality

$$\mathbb{P}\{V(\tau) - \mathbb{E}\{V(\tau)\} > t\} \leq \frac{\mathbb{E}\{|V(\tau) - \mathbb{E}\{V(\tau)\}|^\kappa\}}{t^\kappa}.$$

For the M/G/1 PS queue, the first condition follows from the well-known fact that the mean conditional sojourn time is proportional to the service requirement, with constant of proportionality $1/(1 - \rho)$. The second condition is ensured by the Chebyshev-Markov inequality and the asymptotic behavior of the conditional moments in (2). The third condition follows from an elementary sample-path argument.

Like the approach described in Section 2.2, the proof of Theorem 2.5 involves asymptotic lower and upper bounds which asymptotically coincide. In the context of the M/G/1 PS queue, the bounds correspond to the intuitive insight mentioned in the introduction that a long sojourn time is caused by a large service requirement of the customer itself. For the upper bound, write

$$\mathbb{P}\{V > x\} \leq \mathbb{P}\{V > x; B < (\gamma - \epsilon)x\} + \mathbb{P}\{B > (\gamma - \epsilon)x\}. \quad (7)$$

Under the conditions of Theorem 2.5, the first term on the right-hand side is asymptotically negligible compared to the second. Under the same conditions, it can be shown that

$$\mathbb{P}\{V(B) > x\} \geq \mathbb{P}\{V > x; B > (\gamma + \epsilon)x\} \sim \mathbb{P}\{B > (\gamma + \epsilon)x\}.$$

The proof is completed by letting $\epsilon \downarrow 0$ and using $B \in \mathcal{IR}$.

2.4 Probabilistic approach using the attained-service process

Guillemin *et al.* [28] constructed a proof framework that is closely related to the technique described in Section 2.3, which allowed the extension of the asymptotic equivalence (1) to models with features such as admission control and impatience. Their approach is based on the quantity

$$R(x) := \int_{t=0}^x \frac{1}{1 + N(t)} dt, \tag{8}$$

which equals the amount of service received during the time interval $[0, x]$ by a tagged customer arriving at time $t = 0$. Here, $N(t)$ denotes the number of other customers in the system at time t . Notice that $R(\cdot)$ and the conditional sojourn time $V(\cdot)$ are intimately related through $R(V(\tau)) \equiv \tau$. The main result in [28] is stated in the next theorem.

Theorem 2.6 (*Guillemin et al. [28]*)

Let the service requirement distribution be regularly varying of index $\nu > 1$. Suppose that the following two conditions are satisfied:

- (i) $R(x)/x \rightarrow \gamma$ a.s. as $x \rightarrow \infty$ with $0 < \gamma < 1$;
- (ii) *There exists a positive and finite constant $M \in (0, \infty)$ such that $\mathbb{P}\{R(x) \leq x/M\} = o(\mathbb{P}\{B > x\})$ as $x \rightarrow \infty$.*

Then $\mathbb{P}\{V > x\} \equiv \mathbb{P}\{B > R(x)\} \sim \mathbb{P}\{B > \gamma x\}$.

This theorem has recently been applied in Nuyens *et al.* [45] to derive the tail behavior of the sojourn time in GI/GI/1 queues for a class of service disciplines that includes SRPT and FBPS.

The reader may notice the strong similarity between Theorems 2.5 and 2.6. A similar remark as below Theorem 2.5 applies in this case: when applied to sojourn times in queues, any assumption regarding the service requirements B also affects the process $R(x)$. Like the method in [44], the above theorem is not restricted to the M/G/1 PS queue but extends to various other settings, and was in fact developed with that goal in mind. Note that although the first condition in Theorem 2.5 is very weak, $V(\tau)/\tau$ must converge to $1/\gamma$ in probability by the second condition. The flexibility in choosing M in (ii) makes Theorem 2.6 sometimes easy to apply. For example, that condition trivially holds in PS queues with an upper bound on the number of customers. The counterpart in Theorem 2.5 poses restrictions on *any* deviation from the mean, but does so for a *fixed* service requirement τ which makes that method convenient when information is available about the conditional sojourn times. Instead of comparing the pro's and con's of the two methods, we aim at a unifying discussion in Section 3.

Using Theorem 2.6, the asymptotic equivalence (1) is obtained by taking $\gamma = 1 - \rho$ and checking conditions (i) and (ii). For the M/G/1 PS queue it is shown in [28] that, if $a < 1 - \rho$ and $(1 - \rho)/a$ is not an integer, then

$$\mathbb{P}\{V(ax) > x\} = \mathbb{P}\{R(x) < ax\} = O((x\mathbb{P}\{B > x\})^{\ell(a)}),$$

with $\ell(a) = \lfloor \frac{1-\rho}{a} \rfloor$. The intuition behind this result is that the long-term service rate equals $(1 - \rho)/(\ell + 1)$ if there are ℓ additional permanent customers in the system. For the service rate to drop below a , one needs $\ell(a)$ large customers to be in the system simultaneously for $O(x)$ time.

At this stage we omit the proof details as presented in [28]. Instead, we choose this framework for our discussion in the next section, and in doing so provide a shorter proof of Theorem 2.6 (see Proposition 3.3).

3 Conditions for tail equivalence

In this section, we review conditions under which the tail equivalence

$$\mathbb{P}\{V > x\} \sim \mathbb{P}\{B > \gamma x\} \tag{9}$$

holds, as $x \rightarrow \infty$, for some $\gamma > 0$. The discussion in this section is motivated by the studies [44] and [28], which were reviewed in Subsections 2.3 and 2.4, respectively. For our purpose, we choose the framework of [28]. It is important to again draw the attention to the remarks below Theorems 2.5 and 2.6: We impose separate conditions on B and $R(x)$. This corresponds to a situation where the customer, for which we study the sojourn time, may have a service requirement distribution that is different from that of all other customers. The sojourn time of a customer with service time τ is $V(\tau) = R^{-1}(\tau) = \inf\{s : R(s) \geq \tau\}$. In what follows, we use both $V(\tau)$ and its inverse $R(x)$. Note that these two processes are independent of B , and that

$$\mathbb{P}\{V > x\} = \mathbb{P}\{V(B) > x\} = \mathbb{P}\{B > R(x)\}.$$

We further assume that the first condition of Theorem 2.6 is satisfied, which we restate for convenience:

Assumption 3.1 *There exists a constant $\gamma \in (0, \infty)$ such that $R(x)/x \rightarrow \gamma$ almost surely.*

We call γ the long-term service rate. Since the $R(x)$ -process in most queueing systems has a regenerative structure, Assumption 3.1 is not very restrictive in that context. However natural, it is not a necessary condition for (9) to hold, as the following example shows.

Example 3.1 Take $B \in \mathcal{R}_\alpha$, $R(x) = x/Y$ with Y a random variable independent of B , and suppose $\mathbb{E}\{Y^{\alpha+\epsilon}\} < \infty$. Then, by Breiman's theorem [18],

$$\mathbb{P}\{V > x\} = \mathbb{P}\{YB > x\} \sim \mathbb{E}\{Y^\alpha\}\mathbb{P}\{B > x\} \sim \mathbb{P}\{B > \gamma x\},$$

with $\gamma = \mathbb{E}\{Y^\alpha\}^{-1/\alpha}$. □

Assumption 3.1 is also far from sufficient. The goal of this section is to explore what kind of additional conditions on the distribution of B and the process $R(x), x \geq 0$ (or equivalently $V(\tau), \tau > 0$) need to be imposed such that (9) holds.

In Subsection 3.1, we first focus on a weaker form of (9). In particular, we consider conditions under which $\mathbb{P}\{V > x\}$ is at least as large as $\mathbb{P}\{B > \gamma x\}$ (as $x \rightarrow \infty$). This also leads to necessary conditions for (9) to hold. In Subsection 3.2 we describe two sets of conditions which ensure (9).

3.1 Lower bounds and necessary conditions

From Assumption 3.1 it follows that there exists a function $\sigma(x) = o(x)$ such that

$$S(u) = \liminf_{x \rightarrow \infty} \mathbb{P}\{R(x) \leq \gamma x + u\sigma(x)\}$$

has the property that $S(u) \rightarrow 1$ as $u \rightarrow \infty$. Although $\sigma(x)$ always exists, the convergence of $\sigma(x)/x \rightarrow 0$ can be arbitrarily slow. Note that, if $R(x)$ were to satisfy a CLT, then one could take $\sigma(x) = \sqrt{x}$. (This situation is discussed in more detail below.)

A key tool in this subsection is the following simple lower bound.

$$\begin{aligned} \mathbb{P}\{V > x\} = \mathbb{P}\{B > R(x)\} &\geq \mathbb{P}\{B > R(x); R(x) \leq \gamma x + u\sigma(x)\} \\ &\geq \mathbb{P}\{B > \gamma x + u\sigma(x)\} \mathbb{P}\{R(x) \leq \gamma x + u\sigma(x)\}. \end{aligned}$$

This lower bound and the above considerations yield the following result.

Proposition 3.1 *If Assumption 3.1 is satisfied and $B \in \mathcal{IR}$, then*

$$\liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{V > x\}}{\mathbb{P}\{B > \gamma x\}} \geq 1.$$

Below we use the earlier mentioned property that $B \in \mathcal{IR}$ if and only if $\mathbb{P}\{B > x\} \sim \mathbb{P}\{B > x - o(x)\}$ for any function $o(x)$ satisfying $o(x)/x \rightarrow 0$ (due to D.A. Korshunov). If B is not in \mathcal{IR} , then one needs additional regularity conditions as the following example shows.

Example 3.2 Suppose B is not in \mathcal{IR} . Then there exists a function $f(x) = o(x)$ such that $\mathbb{P}\{B > x + f(x)\}/\mathbb{P}\{B > x\} \rightarrow 0$. Let $0 < p < 1$. Take now $R(x) = \gamma x$ with probability p and $R(x) = \gamma x + f(\gamma x)$ with probability $1 - p$. Assumption 3.1 is satisfied and

$$\mathbb{P}\{V > x\} = p\mathbb{P}\{B > \gamma x\} + (1 - p)\mathbb{P}\{B > \gamma x + f(\gamma x)\}.$$

We see that

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}\{V > x\}}{\mathbb{P}\{B > \gamma x\}} = p,$$

which can take any value between 0 and 1. □

We now relax the condition in Proposition 3.1 regarding the distribution of B , while imposing additional conditions on $R(x)$. As in [33], we call B *square-root insensitive* if

$$\mathbb{P}\{B > x\} \sim \mathbb{P}\{B > x - \sqrt{x}\}.$$

Note that all distributions in \mathcal{IR} are square-root insensitive. It will also be convenient to define the following class of distributions, which have lighter tails than square-root insensitive distributions: B is called *moderately heavy-tailed* if $B \in \mathcal{L}$ and

$$\mathbb{P}\{B > x\} = o(\mathbb{P}\{B > x - \sqrt{x}\}).$$

If $\mathbb{P}\{B > x\} = e^{-x^\beta}$, then B is square-root insensitive for $\beta < 1/2$, and B is moderately heavy-tailed for $1/2 < \beta < 1$. These two concepts play a key role in the second result in this subsection, which gives asymptotic lower bounds for $\mathbb{P}\{V > x\}$, assuming that the $R(\cdot)$ -process satisfies a CLT.

The proposition implies that, if B is moderately heavy-tailed, then (9) does *not* hold (V has a heavier tail than B/γ), which agrees with [32].

Proposition 3.2 *Assume that, in addition to Assumption 3.1, $(R(x) - \gamma x)/\sqrt{x}$ converges to a normally distributed random variable U . In that case, if B is square-root insensitive, then*

$$\liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{V > x\}}{\mathbb{P}\{B > \gamma x\}} \geq 1,$$

whereas

$$\liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{V > x\}}{\mathbb{P}\{B > \gamma x\}} = \infty,$$

if B is moderately heavy-tailed.

Proof

Since

$$\mathbb{P}\{V > x\} \geq \mathbb{P}\{B > \gamma x + u\sqrt{x}\} \mathbb{P}\{R(x) > \gamma x + u\sqrt{x}\},$$

and $\mathbb{P}\{R(x) > \gamma x + u\sqrt{x}\} \rightarrow \mathbb{P}\{U > u\} > 0$, we see that

$$\liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{V > x\}}{\mathbb{P}\{B > \gamma x\}} \geq \mathbb{P}\{U > u\} \liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{B > \gamma x + u\sqrt{x}\}}{\mathbb{P}\{B > x\}}.$$

The value of the lim inf on the right-hand side is ∞ if B is moderately heavy-tailed, and is 1 for any u whenever B is square-root insensitive. Thus, in that case we get

$$\liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{V > x\}}{\mathbb{P}\{B > \gamma x\}} \geq \mathbb{P}\{U > u\}.$$

The proof now follows by letting $u \rightarrow -\infty$. □

3.2 Sufficient conditions

We now turn our attention to conditions which *ensure* (9). As noted earlier, since B is independent of $V(\cdot)$, the problem can be seen as sampling the stochastic process $V(\cdot)$ at a random time B . In the probability literature, this problem has been considered in many different settings.

Assumption 3.2 *In this subsection we assume that $B \in \mathcal{IR}$.*

This is a restrictive assumption, and examples of processes for $V(\cdot)$ have been considered which allow for more general distributions of B , see e.g. [3, 27, 33]. However, the classes of processes one can take for $V(\cdot)$ in these papers do not seem to apply to PS queues. Essentially, the large deviations of the process $V(\cdot)$ are assumed to be on a scale which differs significantly from that of B . A notable exception is the treatment of the M/G/1 PS queue by Jelenković & Momčilović [32], who show that (9) holds for a large class of square-root insensitive distributions, as reviewed in Subsection 2.2. However, their analysis is based on an exact representation of the sojourn time V , which is specific to the standard M/G/1 PS queue.

If Assumptions 3.1 and 3.2 are satisfied, then it is possible to give a sufficient condition for (9) which is close to necessary.

Proposition 3.3 *Suppose that Assumptions 3.1 and 3.2 are satisfied. Then (9) holds if*

$$\lim_{\epsilon \downarrow 0} \limsup_{x \rightarrow \infty} \frac{\mathbb{P}\{R(x) < B < \epsilon x\}}{\mathbb{P}\{B > \gamma x\}} = 0. \quad (10)$$

Furthermore, if

$$\lim_{\epsilon \downarrow 0} \liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{R(x) < B < \epsilon x\}}{\mathbb{P}\{B > \gamma x\}} > 0, \quad (11)$$

then (9) does not hold.

Proof

We start by proving the first part of the proposition, which is a simplification of the proof of Theorem 1 in [28]. Note first that the asymptotic lower bound is satisfied (cf. Proposition 3.1). To get an upper bound, write for $\delta > 0$,

$$\begin{aligned} \mathbb{P}\{V > x\} &= \mathbb{P}\{B > R(x); B > (\gamma - \delta)x\} + \mathbb{P}\{B > R(x); B \in (\epsilon x, (\gamma - \delta)x)\} \\ &\quad + \mathbb{P}\{\epsilon x > B > R(x)\} \end{aligned}$$

Denote the three terms on the right-hand side by I, II and III. Term I is less than $\mathbb{P}\{B > (\gamma - \delta)x\}$. To bound the second term, simply note that

$$\text{II} \leq \mathbb{P}\{R(x) < (\gamma - \delta)x\} \mathbb{P}\{B > \epsilon x\}.$$

This is of order $o(\mathbb{P}\{B > x\})$, since by Assumption 3.1, $\mathbb{P}\{R(x) < (\gamma - \delta)x\} \rightarrow 0$, and $B \in \mathcal{IR}$ implies that

$$\limsup_{x \rightarrow \infty} \frac{\mathbb{P}\{B > \epsilon x\}}{\mathbb{P}\{B > x\}} < \infty.$$

We conclude that

$$\limsup_{x \rightarrow \infty} \frac{\mathbb{P}\{V > x\}}{\mathbb{P}\{B > \gamma x\}} \leq \limsup_{x \rightarrow \infty} \frac{\mathbb{P}\{B > (\gamma - \delta)x\}}{\mathbb{P}\{B > \gamma x\}} + \limsup_{x \rightarrow \infty} \frac{\mathbb{P}\{R(x) < B < \epsilon x\}}{\mathbb{P}\{B > \gamma x\}}.$$

The first term converges to 1 as $\delta \downarrow 0$, by the defining property of \mathcal{IR} . The second term converges to 0 as $\epsilon \downarrow 0$ by (10). We conclude that (9) indeed holds.

To prove the second statement, note that

$$\text{I} \geq \mathbb{P}\{B > (\gamma + \delta)x; B > R(x)\} \geq \mathbb{P}\{B > (\gamma + \delta)x\} \mathbb{P}\{R(x) < (\gamma + \delta)x\},$$

and note that $\text{II} \geq 0$. Consequently, using the law of large numbers (Assumption 3.1), we obtain for any $\delta > 0$ and any $\epsilon \in (0, \gamma)$ that

$$\liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{V > x\}}{\mathbb{P}\{B > \gamma x\}} \geq \liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{B > (\gamma + \delta)x\}}{\mathbb{P}\{B > \gamma x\}} + \liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{\epsilon x > B > R(x)\}}{\mathbb{P}\{B > \gamma x\}}.$$

We can now let $\delta \downarrow 0$ to conclude that

$$\liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{V > x\}}{\mathbb{P}\{B > \gamma x\}} \geq 1 + \liminf_{x \rightarrow \infty} \frac{\mathbb{P}\{\epsilon x > B > R(x)\}}{\mathbb{P}\{B > \gamma x\}},$$

which proves the second part of the theorem after taking $\epsilon \downarrow 0$. \square

The above discussion indicates that whether or not the probability $\mathbb{P}\{\epsilon x > B > R(x)\}$ can be bounded in the sense of (10) and (11), is crucial for the asymptotic equivalence (1) to be valid. Note that the event $\{\epsilon x > B > R(x)\}$ corresponds to the simultaneous occurrence of a large sojourn time and small service requirement; this needs to be sufficiently unlikely in order for the tail equivalence to hold. In practice, the double limit condition (10) may not be convenient. We discuss two cases where this condition can be replaced by a simpler one.

Case A. Write

$$\mathbb{P}\{\epsilon x > B > R(x)\} = \int_0^{\epsilon x} \mathbb{P}\{V(\tau) > x\} d\mathbb{P}\{B \leq \tau\},$$

and assume that there exist $\kappa > \alpha, \delta > 0$, and a function $h(\tau) = o(\tau^\kappa)$ such that

$$\mathbb{P}\{V(\tau) > x\} \leq \frac{h(\tau)}{x^\kappa}, \quad (12)$$

for $\tau \leq \epsilon x$. Then it can be checked that (10) is satisfied. In fact, this parallels the approach in Subection 2.3 (cf. Equation (7)). In the current setting, however, we have already discarded the probability $\mathbb{P}\{R(x) < B; \epsilon x < B < (\gamma - \delta)x\} = \mathbb{P}\{V > x; \epsilon x < B < (\gamma - \delta)x\}$ for large x , because of Assumption 3.1. Therefore, we do not need to verify (12) for all $\epsilon < \gamma$ and all $\tau > 0$, as was the case in Subection 2.3.

Condition (12) is especially convenient when information is available regarding the centered moments of $V(\tau)$. As we saw in Subection 2.3, this allows an extension of the asymptotic relation to other disciplines and modifications of the standard M/G/1 PS queue as well.

Case B. Upper bound $\mathbb{P}\{\epsilon x > B > R(x)\}$ by $\mathbb{P}\{R(x) < \epsilon x\}$, i.e., assume that there exists some $\epsilon > 0$ such that

$$\mathbb{P}\{R(x) < \epsilon x\} = o(\mathbb{P}\{B > \gamma x\}). \quad (13)$$

This assumption is used in Guillemin *et al.* [28], and was originally motivated by an application to the M/G/1 PS queue with the additional feature that at most N customers are admitted into the system. In this case $R(x) \geq x/N$ and (13) is simply satisfied by taking $\epsilon = 1/N$. In [28], condition (13) was shown to hold in the ordinary M/G/1 PS queue, in PS queues with impatience, and in single-class PS queues with state-dependent service rates. In the next section we extend the latter case to PS networks.

As a final remark in this section, note that neither of the conditions (12) and (13) implies the other. Which condition is the most convenient to check, depends on the specific model situation. Sometimes neither of the two is easy to check, and one must resort to a different approach; see also Subection 5.2.

4 Extensions

In the previous sections we reviewed various methods that have been developed to derive sojourn time asymptotics in single-server PS queues with Poisson arrivals and heavy-tailed service requirements. As mentioned earlier, some of the methods and results in fact extend to a broader class of models. In the present section, we will discuss several such extensions.

4.1 Systems with varying service rate

Consider an M/G/1 PS queue where the total available service rate varies over time according to a stationary ergodic process $\{C(t), t \in \mathbb{R}\}$, taking values in some set \mathcal{C} . For stability, we assume $\rho < c := \mathbb{E}\{C(0)\}$ where, as before, ρ is the load of the queue (i.e., the arrival rate times the mean service requirement). Models with a time-varying service rate play a particularly crucial role in analyzing the performance of elastic traffic in an integrated system, sharing the bandwidth with higher-priority streaming traffic, see for instance [22, 43]. Because of specific performance requirements, streaming traffic typically receives a certain degree of priority. Elastic traffic then fairly shares the remaining capacity. From the perspective of elastic traffic, we can view the system as a PS queue with varying capacity. Three classes of models can be distinguished:

- If the streaming traffic receives strict priority, then the process $\{C(t), t \in \mathbb{R}\}$ will be completely exogenous and independent of the stochastic behavior of the elastic traffic (see for example Chapter 2 of [43]).
- For general sharing policies, there may be a complex interaction, as in [4, 14, 43] for example.
- A special situation, where the process $\{C(t), t \in \mathbb{R}\}$ depends on the traffic processes only and is not affected by any exogenous factors at all, arises when the service rate $C(t)$ varies with the number of customers $N(t)$ at time t as in [4, 28].

We again suppose that Assumption 3.1 is satisfied, i.e., $R(x)/x \rightarrow \gamma$ a.s. for $x \rightarrow \infty$ and $\gamma > 0$. If the process $\{C(t), t \in \mathbb{R}\}$ does not depend on the traffic processes (the first bullet above) and no customers are denied access to the system, then the ‘reduced service rate’ is $\gamma = c - \rho$. This value may be higher if the presence of a permanent customer causes the available service rate to increase above the long-term average c , as is the case in [4]. Note that, in such a scenario, the dynamics of $C(t)$ change with the presence of the tagged user for which we determine the sojourn time.

In general, let $C^*(t)$ denote the service process if we place a *permanent* user in the system, that shares in the capacity just as any other user, but never completes service. Note that the sojourn time of a tagged user in the original system has the same distribution as the time it takes the permanent user in the modified system to receive an amount of service equal to the tagged customer’s service requirement. Assuming $C^*(t)$ is stationary and ergodic, we may write

$$\gamma = \mathbb{E}\{C^*(0)\} - \rho.$$

Remark. Suppose new customers may be blocked from the system, for example to ensure a minimum service rate per user at times when $C(t)$ is low. Then the above relation can

be further generalized to $\gamma = \mathbb{E}\{C^*(0)\} - \rho^*$, where ρ^* is the load of customers that are *admitted into the system* with a permanent user, see for example [43]. Note that ρ^* does not involve the service requirement of the permanent customer. For clarity, we assume in the sequel that no user is denied access, i.e., $\rho^* = \rho$.

Define $c^{\min} := \inf\{c \in \mathcal{C}\}$. It is readily verified that the amount of service received by each customer is bounded from below, sample-path wise, by that in a PS queue with constant service rate c^{\min} . Thus, if the latter queue is stable, i.e., $\rho < c^{\min}$, then Condition (ii) of Theorem 2.6 can be seen to be satisfied for $M = 2/(c^{\min} - \rho)$. This implies the following result, regardless of the further statistical characteristics of the process $\{C(t), t \in \mathbb{R}\}$. The model in [14] falls in this category.

Proposition 4.1 *If $\rho < c^{\min}$ and $B \in \mathcal{R}$ then $\mathbb{P}\{V > x\} \sim \mathbb{P}\{B > \gamma x\}$ with $\gamma = \mathbb{E}\{C^*(0)\} - \rho$.*

In the opposite case ($\rho > c^{\min}$), the sojourn time asymptotics do strongly depend on the stochastic properties of the process $\{C(t), t \in \mathbb{R}\}$. In particular, denote

$$C(s, t) := \int_{u=s}^t C(u) du,$$

and let

$$Z^- := \sup_{t \geq 0} \{\rho t - C(-t, 0)\}, \quad \text{and} \quad Z^+ := \sup_{t \geq 0} \{\rho t - C(0, t)\}.$$

If $\mathbb{P}\{B > x\} = O(\mathbb{P}\{Z^- > x\})$ or $\mathbb{P}\{B > x\} = O(\mathbb{P}\{Z^+ > x\})$, then a large sojourn time may be likely to occur due to a long period with a low service rate and not necessarily as a result of a large service requirement of a customer itself. Thus, a minimal assumption for the asymptotic tail equivalence to prevail is that $\mathbb{P}\{Z^- > x\} = o(\mathbb{P}\{B > x\})$ and $\mathbb{P}\{Z^+ > x\} = o(\mathbb{P}\{B > x\})$. It was first proven in [43] that the latter assumption is in fact also sufficient in case $\{C(t), t \in \mathbb{R}\}$ is an independent On-Off process, i.e., $C = \{0, 1\}$, with exponentially distributed On-periods. The result in [4] belongs to this category as well.

4.2 Multi-class systems

Consider an $M_k/G_k/1$ PS queue which is offered traffic from K classes. Class- k customers arrive as a Poisson process of rate λ_k , and have generic service requirements B_k . Note that the system may equivalently be interpreted as a standard single-class PS queue with total arrival rate

$$\lambda := \sum_{k=1}^K \lambda_k,$$

where the generic service requirement is distributed as B_k with probability λ_k/λ . Thus, if we consider the sojourn time of an arbitrary customer, then the asymptotic tail equivalence will hold if the conditions of Theorems 2.3, 2.5 or 2.6 are satisfied.

However, if we consider the sojourn time of a customer of a specific class, then the situation is different. Note that each class in isolation may be viewed as a PS queue with

time-varying service rate as discussed above, but in this case that point of view does not seem particularly useful in deriving tail asymptotics. Using a similar method as outlined in Subsection 2.1, Zwart [57] shows that if the service requirement distribution of a particular class is regularly varying, then the asymptotic tail equivalence holds for that class, irrespective of the service requirement distributions of the other classes.

Theorem 4.2 (Zwart [57])

Consider the $M/G/1$ PS queue with multiple classes and total traffic load ρ . Suppose $\mathbb{E}\{B_k^\alpha\} < \infty$ for some $\alpha > 1$ and all classes k . If $B_i \in \mathcal{R}_\nu$ for some class i and some non-integer number $\nu > 1$, then

$$\mathbb{P}\{V_i > x\} \sim \mathbb{P}\{B_i > (1 - \rho)x\}, \quad x \rightarrow \infty,$$

where V_i denotes the steady-state sojourn time of class- i customers.

In particular, this result applies when some of the other classes have service requirement distributions that are much heavier, and in that sense the result suggests that a class is not significantly affected by other classes.

However, it is crucial that the service requirement of the class under consideration itself is regularly varying. An intriguing question is what occurs if the service requirement distribution of the class under consideration itself is light-tailed, but some of the other classes have heavy-tailed service requirements. In [15] a system is considered with $K = 2$ traffic classes. Class one has service requirements

$$\mathbb{P}\{B_1 > x\} = q_1(x)e^{-x^{\eta_1}},$$

with $q_1(x)$ a regularly varying function and $\eta_1 > 0$. Thus, service requirements of class 1 have a lighter tail than any distribution in \mathcal{R} . For class 2 it is assumed that $B_2 \in \mathcal{R}_{\nu_2}$, for some $\nu_2 > 1$. If the number of customers is limited to a maximum value of $M > 0$, it is shown that, for large x , $\mathbb{P}\{V > x\}$ behaves as $\mathbb{P}\{B_1 > x/M\}$ times a regularly varying function, as is shown in the next theorem. In Subection 8.3 we give an intuitive explanation for this result.

Theorem 4.3 (Borst et al. [15])

For any fixed value of $M < \infty$,

$$\mathbb{P}\{V_1 > x\} \sim \mathbb{P}\{B_1 > \frac{x}{M}\} \frac{(1 - \rho)\rho_2^{M-1}}{1 - \rho^{M+1}} \left(\mathbb{P}\{B_2^r > \frac{x}{M}\} \right)^{M-1},$$

where B_2^r represents the residual service requirement of class 2, so that $B_2^r \in \mathcal{R}_{\nu_2-1}$.

In particular, the above result implies that if B_1 is exponentially distributed, then the sojourn time has an exponentially bounded tail as well. However, in the absence of any admission control (i.e., $M = \infty$) it is demonstrated in the same paper that the sojourn time distribution has a heavier tail, and in fact it will be subexponential even if B_1 is exponentially distributed.

Theorem 4.4 (Borst et al. [15])

When $\mathbb{P}\{B_1 > x\} = e^{-\mu_1 x^{\eta_1}}$ and $M = \infty$, i.e., no admission control, the sojourn time distribution of class-1 customers satisfies, for $x \rightarrow \infty$,

$$\mathbb{P}\{V_1 > x\} \geq (1 + o(1))(1 - \rho)c_2 \sqrt{2\pi/c_3} \frac{x^{\frac{1}{2}r_1}}{(\ln x)^{1 - \frac{1}{2}r_1}} e^{-c_1(x \ln x)^{r_1}},$$

where $r_1 = \eta_1/(1 + \eta_1)$ and

$$\begin{aligned} c_0 &= (1 - r_1)(\nu_2 - 1)/(\eta_1 \mu_1), \\ c_1 &= \mu_1 (c_0)^{r_1} + (1 - r_1)(\nu_2 - 1)(1/c_0)^{1-r_1}, \\ c_2 &= (1/c_0)^{2(1-r_1)}, \\ c_3 &= \mu_1 \eta_1 (\eta_1 - 1) (c_0)^{3r_1 - 2} + 2(1 - r_1)(\nu_2 - 1)(1/c_0)^{3(1-r_1)}. \end{aligned}$$

In the multi-class systems discussed above all customers are treated equally, irrespective of their class. The negative impact of other customer types on the quality of service seen by a particular class can be reduced by giving the latter preferential treatment. A natural service differentiation mechanism is to assign different shares to customers of the various classes. This amounts to Discriminatory Processor Sharing, where each class- i customer receives a portion $w_i / (\sum_{k=1}^K w_k n_k)$ of the total service rate when the number of customers of class k is n_k , for all $k = 1, 2, \dots, K$. The positive constants w_k (chosen a priori) determine the level of differentiation between the various classes. Observe that the heuristic explanation provided below (1) does not rely on the shares of all customers being equal. This suggests that the asymptotic tail equivalence should hold for any given class with sufficiently heavy-tailed service requirements. As it turns out however, all the methods discussed so far rely in one way or the other on the fact that the total service rate is shared in an egalitarian manner. In the next section we describe two approaches that have been developed to overcome that limitation. As a by-product, these approaches also allow an extension of the tail equivalence to non-Poisson arrivals.

5 Discriminatory Processor Sharing queues

As illustrated in the previous section, some of the methods and results described in Section 2 extend beyond the setting of a standard single-class PS queue to a considerably wider set of models with time-varying service rates and several customer classes. However, all the methods presented so far do rely on the assumption that the total service rate is shared in an egalitarian manner, and do not easily extend to Discriminatory Processor Sharing (DPS) models, where the total service rate is shared among customers in proportion to class-dependent weight factors. Note that DPS shows some resemblance with the *Generalized* Processor-Sharing (GPS) discipline (or Generalized Head-Of-the-Line (HOL) PS), where the service rate is also shared in accordance with class-dependent weight factors. In GPS, the rate is not divided however among all customers present, but distributed across (non-empty) classes (e.g. the customers at the head-of-the-line of the various classes), irrespective of the actual number of customers present.

The analysis of the DPS discipline is extremely difficult compared to that of the ordinary egalitarian PS discipline, as the relative paucity of results suggests. For a comprehensive survey on results for DPS models in the literature, we refer to [1] in this special issue. Most

notably, the simple geometric queue length distribution for the standard PS discipline does not have any counterpart for DPS. In addition, there do not seem to be manageable transform results available for the sojourn time distribution. This circumstance considerably complicates the derivation of tail asymptotics, since the methods presented in Section 2 rely either on transform techniques or probabilistic approaches that exploit knowledge of the queue length distribution. The derivation of tail asymptotics for DPS thus requires a fundamentally different approach to circumvent these difficulties. We now describe two approaches from [16, 47] that partially fulfill these requirements, but involve additional distributional assumptions. In particular, we assume that customers arrive as a renewal process with mean interarrival time $1/\lambda$, and that an arriving customer is of class i with probability p_i . Let B_i be the service requirement of a class- i customer, and let B be a random variable with distribution

$$\mathbb{P}\{B > x\} := \sum_{i=1}^K p_i \mathbb{P}\{B_i > x\},$$

representing the service requirement of an arbitrary customer. The main result in [16, 47] is stated in the following theorem. We also refer to Theorems 6.1 and 7.2 for sufficient conditions under which the tail equivalence under DPS extends to classes with a lighter service requirement distribution than B .

Theorem 5.1 *If B_i and B are both regularly varying of index $\nu > 2$, then*

$$\mathbb{P}\{V_i > x\} \sim \mathbb{P}\{B_i > (1 - \rho)x\}.$$

Note that the above theorem involves the assumption that B_i and B are regularly varying of index $\nu > 2$, which in particular means that the service requirements have finite variance. While the assumption that B_i is regularly varying is natural (though possibly not strictly necessary), the intuitive explanation mentioned above suggests that the additional assumptions that $\nu > 2$ and that B is regularly varying of the same index, may not be essential for the result to hold: we conjecture that the tail equivalence continues to hold without the latter two assumptions, i.e., apply for any class i with a regularly varying service requirement distribution.

5.1 Workload-based method

The proof method in [16] involves lower and upper bounds which asymptotically coincide. The lower bound proceeds along similar lines as in [28, 44] for the ordinary PS queue. The upper bound however entails quite different arguments, which may be outlined as follows. Let V_0 be the sojourn time of a tagged customer arriving at time 0. First, the probability of interest is split into three terms by considering the amount of time P^r since the busy period containing time 0 started, and the total number of ‘large’ customers arriving during that period, where a customer is considered to be ‘large’ if its service requirement exceeds ϵx for some constant $\epsilon > 0$, independent of x . Specifically,

$$\begin{aligned} \mathbb{P}\{V_0 > x\} &= \mathbb{P}\{V_0 > x; P^r \leq x^\alpha\} + \mathbb{P}\{V_0 > x; P^r > x^\alpha\} \\ &= \sum_{k=0}^L \mathbb{P}\{V_0 > x; P^r \leq x^\alpha; N_{>\epsilon x}(-x^\alpha, 0) = k\} \\ &\quad + \mathbb{P}\{V_0 > x; P^r \leq x^\alpha; N_{>\epsilon x}(-x^\alpha, 0) \geq L + 1\} + \mathbb{P}\{V_0 > x; P^r > x^\alpha\} \\ &\leq \text{I} + \text{II} + \text{III}, \end{aligned}$$

with

$$\begin{aligned} \text{I} &= \sum_{k=0}^L \mathbb{P}\{V_0 > x \mid (P^r \leq x^\alpha; N_{>\epsilon x}(-x^\alpha, 0) = k)\} \mathbb{P}\{N_{>\epsilon x}(-x^\alpha, 0) = k\} \\ \text{II} &= \mathbb{P}\{N_{>\epsilon x}(-x^\alpha, 0) \geq L + 1\} \\ \text{III} &= \mathbb{P}\{P^r > x^\alpha\}, \end{aligned}$$

and $N_{>u}(0, t)$ denoting the total number of customers with a service requirement larger than u arriving during the time interval $[0, t]$. Next, each of the three terms I, II, III is considered separately.

Since B is regularly varying of index ν , P is regularly varying of index ν as well (according to Lemma 5.3.1 in Zwart [58]). Hence, P^r is regularly varying of index $\nu - 1$. So there exists a slowly varying function $l(\cdot)$ so that $\mathbb{P}\{P^r > x^\alpha\} = l(x)x^{\alpha(1-\nu)}$. If we assume $\nu > 2$, and take $\alpha := -(\nu + \delta)/(1 - \nu)$, then $\delta > 0$ can be chosen sufficiently small so that $\alpha < \nu$, so that

$$\text{III} = \mathbb{P}\{P^r > x^\alpha\} = l(x)x^{\alpha(1-\nu)} = l(x)x^{-(\nu+\delta)} = o(\mathbb{P}\{B > x\}).$$

Using the Elementary Renewal Theorem, it may be shown that $\text{II} = \mathbb{P}\{N_{>\epsilon x}(-x^\alpha, 0) \geq L + 1\} = O((x^\alpha \mathbb{P}\{B > x\})^{L+1}) = o(\mathbb{P}\{B_0 > x\})$ for $L = \lfloor \frac{\nu+\delta}{\nu-\alpha} \rfloor$.

Also, $\mathbb{P}\{N_{>\epsilon x}(-x^\alpha, 0) = k\} = o(1)$ for all $k = 1, \dots, L$.

Finally, using sample-path arguments and Theorem 2 from Jelenković [31], it may be shown that for any $\delta > 0$ there exists an $\epsilon > 0$ such that

$$\text{I} = \mathbb{P}\{V_0 > x \mid (P^r \leq x^\alpha; N_{>\epsilon x}(-x^\alpha, 0) = k)\} \leq \mathbb{P}\{B > \frac{1 - \rho - 3\delta}{kf_{DPS} + 1}x\}(1 + o(1)),$$

where $f_{DPS} := \frac{\max_{j=1, \dots, K} w_j}{w_i}$ is a coefficient that depends on the specific values of the DPS weights, with i the class index of the tagged customer.

The proof is then completed by letting $\delta \downarrow 0$ and using the fact that B is regularly varying. Note that the assumption $\nu > 2$ is needed to ensure that $\text{I} = o(\mathbb{P}\{B > x\})$.

5.2 Stochastic mean-value method

The proof method in [47] also involves lower and upper bounds which asymptotically coincide. The lower bound is identical to that in [16], but the upper bound proceeds along a novel path based on the stochastic mean-value theorem, see for instance Corollary 1.4 on p. 171 of Asmussen [2]. Specifically, the stochastic mean-value theorem is used to rewrite the probability of interest in terms of the expected fraction of class- i customers with sojourn time larger than x arriving during a busy period:

$$\mathbb{P}\{V_i > x\} = \frac{1}{\mathbb{E}\{N\}p_i} \mathbb{E}\left\{\sum_{j=1}^N \mathbb{I}_{\{V^{(j)} > x\}} \mathbb{I}_{\{j \in T_i\}}\right\},$$

where $\mathbb{I}_{\{A\}}$ denotes the indicator function of the event A , and the event $j \in T_i$ indicates whether or not the j -th customer arriving during the busy period belongs to class i . Denote by N the number of customers that arrive during a busy period of length P . Let $B^{(j)}$ be the service requirement of the j -th customer arriving during the busy period and let $V^{(j)}$ be its sojourn time. The variable $A(y)$ denotes the set of customers with service requirement

larger than y that arrive during the busy period, i.e., $A(y) := \{j \in \{1, \dots, N\} : B^{(j)} > y\}$, and $N_{>y}$ indicates the cardinality of this set. The above expression may then be divided into several terms by observing that $V^{(j)} > x$ for some $j \in \{1, \dots, N\}$ implies that $P > x$, and conditioning on the number of ‘large’ customers that arrive during the busy period, where as before a customer is said to be ‘large’ when its service requirement exceeds ϵx , for some constant $\epsilon > 0$ independent of x :

$$\mathbb{P}\{V_i > x\} = \frac{1}{\mathbb{E}\{N\}p_i} \mathbb{E}\left\{\sum_{i=1}^N \mathbb{I}_{\{V^{(j)} > x\}} \mathbb{I}_{\{P > x\}} \mathbb{I}_{\{j \in T_i\}}\right\} = \frac{1}{\mathbb{E}\{N\}p_i} \left(\sum_{k=0}^2 (\text{I}_k + \text{II}_k) + \text{III}\right),$$

where

$$\begin{aligned} \text{I}_k &= \mathbb{E}\left\{\sum_{j \notin A(\epsilon x)} \mathbb{I}_{\{V^{(j)} > x\}} \mathbb{I}_{\{j \in T_i\}} \mathbb{I}_{\{P > x\}} \mathbb{I}_{\{N_{>\epsilon x} = k\}}\right\}, \\ \text{II}_k &= \mathbb{E}\left\{\sum_{j \in A(\epsilon x)} \mathbb{I}_{\{V^{(j)} > x\}} \mathbb{I}_{\{j \in T_i\}} \mathbb{I}_{\{P > x\}} \mathbb{I}_{\{N_{>\epsilon x} = k\}}\right\}, \\ \text{III} &= \mathbb{E}\left\{\sum_{j=1}^N \mathbb{I}_{\{V^{(j)} > x\}} \mathbb{I}_{\{j \in T_i\}} \mathbb{I}_{\{P > x\}} \mathbb{I}_{\{N_{>\epsilon x} \geq 3\}}\right\}, \end{aligned}$$

Consider each of the above terms separately.

We start with the dominant term II_1 , and then proceed to show that all others can be asymptotically neglected. Clearly, $\text{II}_1 \leq \mathbb{E}\{\mathbb{I}_{\{P > x\}}\} = \mathbb{P}\{P > x\}$, and Theorem 5.3.1 in Zwart [58] implies

$$\mathbb{P}\{P > x\} \sim \mathbb{E}\{N\} \mathbb{P}\{B > (1 - \rho)x\} \sim p_i \mathbb{E}\{N\} \mathbb{P}\{B_i > (1 - \rho)x\}.$$

Turning to the term II_2 , it may be shown that

$$\text{II}_2 \leq \mathbb{E}\{2\mathbb{I}_{\{N_{>\epsilon x} = 2\}}\} = 2\mathbb{P}\{N_{>\epsilon x} = 2\} = o(\mathbb{P}\{B_i > (1 - \rho)x\}).$$

By definition of $A(\epsilon x)$, we have $\text{II}_0 = 0$.

In conclusion, $\frac{1}{\mathbb{E}\{N\}p_i} \sum_{k=0}^2 \text{II}_k \leq \mathbb{P}\{B_i > (1 - \rho)x\}(1 + o(1))$.

Next, we show that the terms I_k , $k = 0, 1, 2$, asymptotically vanish as well. The sojourn time of customers that arrive to a system with at most k large customers is smaller than the sojourn time of these same customers entering a system with k permanent customers. Define $N_{\leq \epsilon x, k}$ as the number of customers arriving during a busy period in a system with k permanent customers and all service requirements truncated at level ϵx . Denote by $V_{\leq \epsilon x, k}^{(j)}$ and $V_{\leq \epsilon x, k}$ the sojourn times of the j -th customer and an arbitrary customer entering this system, respectively. With these random variables, the above assertion may be formalized as follows:

$$\text{I}_k \leq \mathbb{E}\left\{\sum_{j=1}^{N_{\leq \epsilon x, k}} \mathbb{I}_{\{V_{\leq \epsilon x, k}^{(j)} > x\}}\right\} = \mathbb{E}\{N_{\leq \epsilon x, k}\} \mathbb{P}\{V_{\leq \epsilon x, k} > x\}.$$

Let $W_{\leq \epsilon x, k}(0^-)$ be the amount of work in the system just before the tagged customer arrives, and let $W_{\leq y}^c := \sup_{t \geq 0} [A_{\leq y}(0, t) - ct]$ be the stationary amount of work in a system

with constant service rate c where all service requirements are truncated at level y . Using sample-path arguments, it may be shown that

$$\begin{aligned} \mathbb{P}\{V_{\leq \epsilon x, k} > x\} &\leq \mathbb{P}\{W_{\leq \epsilon x, k}(0^-) + W_{\leq \epsilon f_{DPS} x}^{\rho+\delta} > (1 - (k+1)\epsilon)x\} \\ &\leq \mathbb{P}\{W_{\leq \epsilon x, k}(0^-) > \frac{1 - (k+1)\epsilon}{2}x\} + \mathbb{P}\{W_{\leq \epsilon f_{DPS} x}^{\rho+\delta} > \frac{1 - (k+1)\epsilon}{2}x\}, \end{aligned}$$

where $f_{DPS} := \frac{\max_{j=1, \dots, K} w_j}{w_i}$ is a coefficient that bounds the ratio between the weight of class- i customers and the weight of other customers.

Both of the above probabilities are $o(\mathbb{P}\{B > x\})$ and hence $o(\mathbb{P}\{B_i > (1 - \rho)x\})$. Since $\mathbb{E}\{N_{\leq \epsilon x, k}\}$ is finite, it follows that $I_k = o(\mathbb{P}\{B_i > (1 - \rho)x\})$ for all k .

It remains to prove that the term III is asymptotically negligible. Applying Hölder's inequality, we may write

$$\begin{aligned} \text{III} &\leq \mathbb{E}\{N \mathbb{I}_{\{N_{> \epsilon x} \geq 3\}}\} \\ &\leq (\mathbb{E}\{N^2\})^{1/2} (\mathbb{E}\{(\mathbb{I}_{\{N_{> \epsilon x}\}} \geq 3)^2\})^{1/2} \\ &= (\mathbb{E}\{N^2\})^{1/2} (\mathbb{P}\{N_{> \epsilon x} \geq 3\})^{1/2}. \end{aligned}$$

The proof is then completed by establishing that $\mathbb{E}\{N^2\} < \infty$ and $\mathbb{P}\{N_{> \epsilon x} \geq 3\}^{1/2} = o(\mathbb{P}\{B_i > (1 - \rho)x\})$.

The assumption $\nu > 2$ is required for the proof (though probably not necessary for the final result to hold) because otherwise the probability $\mathbb{P}\{N_{> \epsilon x} \geq 3\}$ is not sufficiently small. The reason is that when $1 < \nu < 2$ the most likely scenario for the event of interest to occur no longer involves the arrival of three or more large customers with service requirements of the order $O(x)$. A more likely scenario consists of the arrival of an extremely large customer with a service requirement of the order $O(x^\nu)$, after which it is virtually certain that several other customers will arrive with service requirements of the order $O(x)$. For further details we refer to [47].

6 Relation with geometric queue length distribution

In this section we identify an intimate relationship between the asymptotic tail equivalence (1) and a geometrically bounded queue length distribution. We illustrate this connection in the context of a DPS system for which the tail equivalence has only been established under restrictive distributional assumptions, as described in the previous section.

Using the methods in [11], it can actually be shown that the vector of remaining service requirements in a DPS system can be bounded from above sample-path wise by that in an ordinary PS system with a lower service rate. Specifically, the state vector in a DPS system with K classes with weights w_1, \dots, w_K is bounded from above by that in an ordinary PS system with service rate $s(w_1, \dots, w_K) := \min_{k=1, \dots, K} w_k / \max_{l=1, \dots, K} w_l$. It then follows that the queue length distribution in the original DPS system is geometrically bounded provided $\rho < s(w_1, \dots, w_K)$.

This relationship thus provides a novel proof of the tail equivalence for a DPS system which involves additional assumptions on the load and weight parameters, but holds for any class with a regularly varying service requirement distribution (not just the “heaviest” one), and no longer requires the restrictive assumption of finite variance. See also Theorem 7.2 for

yet another sufficient set of assumptions under which the tail equivalence under DPS is preserved.

Let N be the total number of customers in the DPS system under consideration.

Theorem 6.1 *If B_i is regularly varying of index $\nu_i > 1$ and $\mathbb{P}\{N > n\} \leq ab^n$, with $a < \infty$, $b < 1$, then $\mathbb{P}\{V_i > x\} \sim \mathbb{P}\{B_i > (1 - \rho)x\}$.*

Proof

The proof relies on lower and upper bounds which asymptotically coincide. As noted in the previous section, a lower bound may be established along similar lines as in [28, 44] for the ordinary PS queue. To obtain a matching upper bound, let us focus on a tagged customer of one particular class, say class 1, and write $\nu := \nu_1$. Let B_0 and V_0 be the service requirement and sojourn time of the tagged customer, respectively.

Using a similar sample-path representation as in (4), it may be shown that

$$V_0(1 - \rho - \delta) \leq B_0 + \sum_{k=1}^K \sum_{i=1}^{L_k(0)} \min\left\{\frac{w_k}{w_1} B_0, B_{ki}^r\right\} + W_{w_{\max} B_0}^{\rho+\delta},$$

with $w_{\max} := \max_{k=1, \dots, K} w_k/w_1$,

$$W_{w_{\max} B_0}^{\rho+\delta} := \sup_{t \geq 0} \left\{ \sum_{k=1}^K \sum_{i=1}^{N_k(0,t)} \min\{w_{\max} B_0, B_{ki}\} - (\rho + \delta)t \right\}$$

and that $\mathbb{P}\{W_{w_{\max} B_0}^{\rho+\delta} > \delta x\} = o(\mathbb{P}\{B_0 > x\})$.

Thus,

$$V_0(1 - \rho - \delta) \leq B_0 + W + W_{w_{\max} B_0}^{\rho+\delta}, \tag{14}$$

with $W := \sum_{k=1}^K \sum_{i=1}^{L_k(0)} B_{ki}^r$, and

$$V_0(1 - \rho - \delta) \leq (w_{\max} N + 1)B_0 + W_{w_{\max} B_0}^{\rho+\delta}, \tag{15}$$

with $N := \sum_{k=1}^K L_k(0)$.

We may write

$$\mathbb{P}\{V_0 > x\} = \text{I} + \text{II} + \text{III},$$

with

$$\begin{aligned} \text{I} &= \mathbb{P}\{V_0 > x; N \leq \log x^m; W \leq \epsilon x\} \\ \text{II} &= \mathbb{P}\{V_0 > x; N \leq \log x^m; W > \epsilon x\} \\ \text{III} &= \mathbb{P}\{V_0 > x; N > \log x^m\}. \end{aligned}$$

Using (14), term I may be estimated by

$$\begin{aligned} \text{I} &\leq \mathbb{P}\{B_0 + W + W_{w_{\max} B_0}^{\rho+\delta} > (1 - \rho - \delta)x; W \leq \epsilon x\} \\ &\leq \mathbb{P}\{B_0 + W_{w_{\max} B_0}^{\rho+\delta} > (1 - \rho - \delta - \epsilon)x\} \\ &\leq \mathbb{P}\{B_0 > (1 - \rho - \epsilon - 2\delta)x\} + \mathbb{P}\{W_{w_{\max} B_0}^{\rho+\delta} > \delta x\} \\ &= \mathbb{P}\{B_0 > (1 - \rho - \epsilon - 2\delta)x\}(1 + o(1)). \end{aligned}$$

Using (15), term II may be bounded by

$$\begin{aligned}
\text{II} &\leq \mathbb{P}\{(w_{\max}N + 1)B_0 + W_{w_{\max}B_0}^{\rho+\delta} > (1 - \rho - \delta)x; N \leq \log x^m; W > \epsilon x\} \\
&\leq \mathbb{P}\{B_0 > ((1 - \rho - \delta)x - W_{w_{\max}B_0}^{\rho+\delta})/(w_{\max} \log x^m + 1); W > \epsilon x\} \\
&\leq \mathbb{P}\{B_0 > (1 - \rho - 2\delta)x/(w_{\max} \log x^m + 1); W > \epsilon x\} + \mathbb{P}\{W_{w_{\max}B_0}^{\rho+\delta} > \delta x\} \\
&\leq o(x^{\eta-\nu})O(x^{1-\nu}) + o(\mathbb{P}\{B_0 > x\}) \\
&= o(x^{\eta+1-2\nu}) + o(x^{-\nu})
\end{aligned}$$

for any $m, \eta > 0$ and $\eta + 1 - 2\nu < -\nu$ for η sufficiently small, and thus $\text{II} = o(\mathbb{P}\{B_0 > x\})$. In order to control term III, note that

$$\text{III} \leq \mathbb{P}\{N > \log x^m\} \leq ab^{\log x^m} = ax^{m \log b} = o(x^{-\nu})$$

for m sufficiently large. □

Observe that the above proof does not use any knowledge of the remaining service requirements of individual customers. Instead, it exploits the fact that their contribution to the sojourn time of the tagged customer is simultaneously bounded by the total workload and the total number of customers times the service requirement of the tagged customer, as captured by inequalities (14) and (15), respectively. This observation also played a key role in the proof of tail equivalence in a PS queue with random service interruptions in [43, Chapter 5]. When information on the remaining service requirements is available, the proof may be simplified.

7 Bandwidth-sharing networks

In this section we consider a class of bandwidth-sharing models that has attracted much attention recently. We adopt in particular the framework of Bonald & Proutière [11]. Suppose there are K classes of customers which arrive according to Poisson processes of rate λ_i , and have service requirements B_i , $i = 1, \dots, K$. As before, the service requirement of a customer from an arbitrary class is denoted by B which, with probability $\lambda_i/(\lambda_1 + \lambda_2 + \dots + \lambda_K)$, is distributed as B_i . Set $\rho_i := \lambda_i \mathbb{E}\{B_i\}$ and $\rho := \rho_1 + \rho_2 + \dots + \rho_K$. The state of the system with x_i customers of classes $i = 1, \dots, K$, is denoted by the vector $x := (x_1, \dots, x_K)$. Throughout the section, we use the notation $|x| := x_1 + x_2 + \dots + x_K$. The service rate allocated to class i in state x is supposed to be given by a function $\phi_i(x)$, and each customer of class i receives an equal share $\phi_i(x)/x_i$ whenever $x_i > 0$.

Let V_i be the generic sojourn time of a class- i customer. Since Assumption 3.1 is satisfied as long as the network is stable, it is natural to conjecture that if $B_i \in \mathcal{IR}$, then there exists a constant γ_i such that

$$\mathbb{P}\{V_i > x\} \sim \mathbb{P}\{B_i \sim \gamma_i x\}. \tag{16}$$

The main result of this section establishes this asymptotic relation under a particular assumption on the allocation function ϕ . This extends the result of Theorem 3 in [28] where the single-class case was considered.

Assumption 7.1 *The allocation function satisfies*

$$\phi_i(x) \geq c_i x_i / (|x| + k),$$

for some integer $k \geq 0$ and constants c_i , $i = 1, 2, \dots, K$.

If the above assumption is satisfied, then we may in fact take $c_i = 1$, for all i , without loss of generality. To see this, note that for any constants $c_i > 0$, $i = 1, 2, \dots, K$, we can consider the following equivalent “normalized” system in terms of the queue length processes and the sojourn times of individual customers. The normalization consists in scaling all service requirements of class i by c_i , that is, we take $B'_i = B_i/c_i$ as the service requirements in the normalized system. In order to preserve the queue length dynamics and sojourn times of customers, we scale the allocation function correspondingly: $\phi'_i(x) = \phi_i(x)/c_i$. Clearly, we then have $\phi'_i(x) \geq x_i/(|x| + k)$. Note that in the normalized system the loads are given by $\rho'_i := \rho_i/c_i$.

We will show later (see the discussion preceding Theorem 7.2 below) that the additional restriction imposed in the next theorem is, in general, not necessary for (16) to hold. Note further that the constant γ_i is not determined by the theorem. For the special case of DPS and so called *balanced networks* we will determine the constant in closed form.

Theorem 7.1 *Assume that the bandwidth-sharing network is stable, that $B_i \in \mathcal{R}_\alpha$, $\alpha > 1$, and that $\mathbb{E}\{B^{1+\eta}\} < \infty$ for some $\eta > 0$. Suppose further that Assumption 7.1 is satisfied and that*

$$\rho_c = \sum_{i=1}^K \frac{\rho_i}{c_i} < 1, \tag{17}$$

then (16) holds for some constant γ_i .

Proof

The proof is based on that of Theorem 3 in [28] where the result was proved for the single-class case. The approach is similar to the stochastic-comparison method in [11]. By Assumption 3.1 and the law of large numbers for regenerative processes, it suffices to show that (13) holds, where $R(x) := R_i(x)$ stands for the attained-service process of an arbitrary customer of class i and B is distributed as the service requirement of an arbitrary customer. The key element in the proof is to bound sojourn times in the original queue from above by sojourn times in a multi-class M/G/1 PS queue with k permanent customers. Recall that we may assume that the service requirements and the allocation functions are normalized such that $c_j = 1$ for all classes j . To reflect this normalization in the notation, we denote the service requirements, the allocation functions and the individual traffic loads in the normalized system by B'_j , $\phi'_j(x)$ and ρ'_j , respectively. Note that $\rho_c := \rho'_1 + \rho'_2 + \dots + \rho'_K$ can be seen as the total traffic load in the normalized system. Let us compare the normalized system with a (multi-class) M/G/1 PS queue with k permanent customers. This queue is fed with the same arrivals and service requirements as the normalized system. The condition $\rho_c < 1$ ensures stability of this queue. Assumption 7.1 implies $\phi'_i(x) \geq x_i/(|x| + k)$. A straightforward sample-path comparison shows that, at all times, all customers receive at least as much service in the normalized system

as in the reference PS queue with k permanent customers. Denoting the attained-service process of a class- i customer in the normalized queue by $R'_i(x) = R_i(x)/c_i$, we have

$$R'_i(x) \geq R^k(x).$$

Here, $R^k(x)$ is the attained-service process in the PS queue with k permanent customers (which is independent of the class since all customers share equally in the service capacity). We conclude that, for any $\epsilon > 0$,

$$\mathbb{P}\{R'_i(x) < \epsilon x\} \leq \mathbb{P}\{R^k(x) < \epsilon x\} = \mathbb{P}\{V^k(\epsilon x) > x\}.$$

By the decomposition of conditional sojourn times established in Van den Berg & Boxma [5], we have ($\stackrel{d}{=}$ denotes equality in distribution)

$$V^k(\epsilon x) \stackrel{d}{=} V_1^0(\epsilon x) + \dots + V_{k+1}^0(\epsilon x),$$

where all terms are independent of each other and $V_i^0(\epsilon x)$ is distributed as the conditional sojourn time in an ordinary M/G/1 PS queue with service requirement ϵx . Consequently,

$$\mathbb{P}\{R^k(x) < \epsilon x\} \leq (k+1)\mathbb{P}\{V^0(\epsilon x) > x/(k+1)\}.$$

Thus, it is sufficient to verify condition (13) for the ordinary M/G/1 PS queue. Since $\mathbb{E}\{B^{1+\eta}\} < \infty$, we may apply Lemma 3 of [28] which ensures that $\mathbb{P}\{V^0(\epsilon x) > x/(k+1)\} = o(\mathbb{P}\{B > x\})$. \square

An example of a system that satisfies Assumption 7.1 is the DPS model discussed in Section 5:

$$\phi_i(x) = \frac{w_i x_i}{w_1 x_1 + w_2 x_2 + \dots + w_K x_K} \geq \frac{c_i x_i}{|x|},$$

with $c_i := w_i / \max\{w_1, w_2, \dots, w_K\}$. Therefore, Theorem 7.1 also shows that under the more restrictive “stability” assumption (17) the result of Theorem 5.1 is also true for classes whose service requirement distributions have lighter tails than B . Note the similarity with Theorem 6.1. It is worth emphasizing that, in this case, the constant γ_i is simply equal to $1 - \rho$, independent of the class index i . This class-independence of large sojourn times under DPS is proved in the following theorem; see [1] for a related discussion.

Theorem 7.2 *Consider the stable DPS model with Poisson arrivals and assume*

$$\max\{w_1, w_2, \dots, w_K\} \sum_{i=1}^K \rho_i / w_i < 1.$$

Suppose further that $B_i \in \mathcal{R}_\alpha$, $\alpha > 1$ and $\mathbb{E}\{B^{1+\eta}\} < \infty$ for some $\eta > 0$. Then

$$\mathbb{P}\{V_i > x\} \sim \mathbb{P}\{B_i > (1 - \rho)x\}.$$

Proof

By Theorem 7.1 it remains to be shown that $\gamma_i = 1 - \rho$ for all i . Note that the DPS system with one permanent customer is stable and regenerative. Thus, for any class, irrespective of the weight, it must be that $R_k(x)/x \rightarrow 1 - \rho$, as $x \rightarrow \infty$. \square

Balanced networks

Theorem 7.1 establishes the asymptotic tail equivalence, but does not provide closed-form expressions for condition (17) and the constants γ_i . These can be determined explicitly for the class of *balanced networks*, which we describe next, following the discussion in [11]. Let e_i be the unit vector with value 1 in the i -th component and 0 in all other components. The PS network is called balanced if

$$\phi_i(x)/\phi_i(x - e_j) = \phi_j(x)/\phi_j(x - e_i), \quad (18)$$

for all i, j and x . Define the *balance function* $\Phi(x)$ by

$$\Phi(x) = \Phi(x - e_i)/\phi_i(x).$$

If (18) holds, and if

$$\sum_x \Phi(x) \prod_{i=1}^K \rho_i^{x_i} < \infty,$$

then the PS network is stable, and the distribution of the steady-state population vector X has distribution

$$\mathbb{P}\{X = x\} = C_\Phi \Phi(x) \prod_{i=1}^K \rho_i^{x_i}.$$

The constant C_Φ follows from the common normalization. In what follows, we are interested in the sojourn time of a class- i customer. As before, we consider a PS network with one permanent class- i customer. The attained service process of the permanent customer is identical to that of an arbitrary class- i customer (until departure). Let $X^i(t)$ be the state vector at time t in the PS network with the permanent class- i customer. The state descriptor $X^i(t)$ does not include the permanent customer. It is easy to see that this network is balanced as well, with balance function $\Phi^i(x) = \Phi(x + e_i)$.

The service rate $r_i(t)$ received by the tagged customer at time t is

$$r_i(t) = \frac{\phi_i(X^i(t) + e_i)}{X_i^i(t) + 1}.$$

The process $X^i(t)$ remains stable, and is also regenerative. Denoting the associated attained-service process with $R_i(x) = \int_{t=0}^x r_i(t) dt$, we conclude that $R_i(x)/x$ converges to a limit γ_i given by

$$\gamma_i = C_\Phi \sum_x \frac{\phi_i(x + e_i)}{x_i + 1} \Phi(x + e_i) \prod_{j=1}^K \rho_j^{x_j}.$$

Example 7.1 Proportional fairness

We now consider a proportional fair network as a particular example of a balanced multi-class PS network. For more examples in the single-class setting, such as the M/G/s PS queue, we refer to Section 4 of [28].

Suppose there are three classes of flows indexed by 0, 1, and 2, and that the bandwidth allocation functions are given by

$$\phi_0(x) = x_0/(x_0 + x_1 + x_2),$$

$$\phi_i(x) = (x_1 + x_2)/(x_0 + x_1 + x_2), i = 1, 2.$$

For this particular network, this allocation is called *proportional fair*, see e.g. Bonald & Massoulié [8]. Taking $c_i = 1$ for $i = 0, 1, 2$ and $k = 0$, Theorem 7.1 yields that, if $\rho_0 + \rho_1 + \rho_2 < 1$, then

$$\mathbb{P}\{V_i > x\} \sim \mathbb{P}\{B > \gamma_i x\},$$

for $i = 0, 1, 2$. In [8], the constants γ_i , also known as the flow throughput, have been computed. They are given by $\gamma_i = 1 - \rho_0 - \rho_i$ for $i = 1, 2$, and

$$\gamma_0 = \frac{1 - \rho_0}{1 + \rho_1/\gamma_1 + \rho_2/\gamma_2}.$$

Note that the condition $\rho_0 + \rho_1 + \rho_2 < 1$ is quite strong, even more so for larger networks: In a linear network with $L > 2$ nodes, the corresponding condition is $\rho_0 + \rho_1 + \dots + \rho_L < 1$ which may be highly restrictive for large values of L . Therefore, an interesting direction for further research is to relax this condition while preserving the validity of (9).

8 Light-tailed service requirements

Except for the two-class system discussed in Subection 4.2, we have concentrated on PS models where the service requirements are heavy-tailed. Since the analysis of queueing models with heavy tails has a relatively short history, one would expect that the theory for sojourn time asymptotics for light-tailed service requirements would be at least as developed. This is far from true however, and sojourn time asymptotics in the light-tailed case are in fact far more difficult to obtain than in the heavy-tailed case. The results known to the authors are reviewed in this section.

8.1 The M/M/1 PS queue

The asymptotic tail behavior of the sojourn time in the M/M/1 PS queue (with arrival rate λ and service rate μ) is rather complicated:

$$\mathbb{P}\{V > x\} \sim \alpha_m x^{-5/6} e^{-\beta_m x^{1/3}} e^{-\gamma_m x},$$

with

$$\alpha_m = \lambda^{-5/6} 2^{2/3} 3^{-1/2} \pi^{5/6} \rho^{5/12} \frac{1 + \sqrt{\rho}}{(1 - \sqrt{\rho})^3} \exp\left(\frac{1 + \sqrt{\rho}}{1 - \sqrt{\rho}}\right),$$

$$\beta_m = 3\lambda^{1/3} \left(\frac{\pi}{2}\right)^{2/3} \rho^{-1/6},$$

$$\gamma_m = \lambda(1/\sqrt{\rho} - 1)^2.$$

Borst *et al.* [12] recently obtained the result by showing that

$$\mathbb{P}\{V > x\} = \frac{1}{\rho} \mathbb{P}\{W_{ROS} > x\},$$

with W_{ROS} the waiting time in the M/M/1 queue with random order of service. The asymptotic behavior of $\mathbb{P}\{W_{ROS} > x\}$ was established by Flatto [26]. Flatto's proof builds upon an integral representation for $\mathbb{P}\{W_{ROS} > x\}$, and an application of the Laplace method, which is made possible after a number of ingenious transformations. It is rather striking that Flatto's results have already been published (without proof) in 1946 by Pollaczek [50].

8.2 The M/D/1 PS queue

A further case where the tail behavior of $\mathbb{P}\{V > x\}$ is known, is the M/D/1 PS queue. Egorova *et al.* [23] recently showed that

$$\mathbb{P}\{V > x\} \sim \alpha_d e^{-\gamma_d x}, \quad (19)$$

with γ_d the unique positive real solution of

$$\frac{\lambda D(\lambda - s) + s - s e^{(\lambda - s)D}}{D(\lambda - s)(\lambda - s e^{(\lambda - s)D})} = \frac{1}{\rho},$$

and

$$\alpha_d = \frac{(1 - \rho)(\lambda - \gamma_d)}{2\lambda(1 - \rho) - \gamma_d \rho(2 - \rho)}.$$

The proof in [23] exploits a geometric random sum representation for V . This representation is also valid for the conditional sojourn time $V(\tau)$ for customers with service requirement τ in the M/G/1 PS queue; the extension of (19) in this direction is forthcoming. It is remarkable that the asymptotic behavior of $\mathbb{P}\{V > x\}$ is completely different from that in the M/M/1 case. In the next subsection we elaborate further on this.

8.3 General distributions and discussion

The above results suggest that it is harder to derive exact asymptotics when the inter-arrival times and service requirements have a general distribution with moment generating functions $\Phi_A(s)$ and $\Phi_B(s)$, respectively. Mandjes & Zwart [36] considered the problem of finding logarithmic asymptotics for $\mathbb{P}\{V > x\}$. Under the assumption that $\mathbb{E}\{e^{\epsilon \exp\{B\}}\} = \infty$ for all $\epsilon > 0$, it is shown that

$$\log \mathbb{P}\{V > x\} \sim -\gamma_g x, \quad (20)$$

with $\gamma_g = \sup_{s \geq 0} [s - \Psi(s)]$, and $\Psi(s) = -\Phi_A^{\leftarrow} \left(\frac{1}{\Phi_B(s)} \right)$. A further assumption made in [36] is that $\Psi'(s^*) = 1$ for some $s^* > 0$, and that $\Phi_B(s) < \infty$ in a neighborhood of s^* .

If both A and B are exponentially distributed, then γ_g indeed reduces to γ_m as defined in Subection 8.1. However, when B is deterministic, γ_d is different from the corresponding value of γ_g in (20). This shows that the tail of B must be “heavy enough” for (20) to hold, although the condition $\mathbb{E}\{e^{\epsilon \exp\{B\}}\} = \infty$ for all $\epsilon > 0$, may be too strong.

To explain these different asymptotics, note that in general, three different types of events may contribute to a long sojourn time: (i) The tagged customer has a long service requirement; (ii) The number of customers in the system at time 0 is exceptionally large; and (iii) the input in the system during the sojourn time of the tagged customer is larger than usual. If the service requirements are heavy-tailed, then event (i) is primarily responsible for the occurrence of a long sojourn time. Using large-deviations theory, it is shown in [36] that event (iii) determines the logarithmic asymptotics (20). Specifically, the input rate after time 0 increases from ρ to 1. This behavior is very similar to that of the busy period, for which the same logarithmic asymptotic behavior has been shown to hold, see Palmowski & Rolski [49] and Nuyens & Zwart [46].

If the service requirements are deterministic, then changing the drift from ρ to 1 does not suffice. The analysis in [23] shows that long sojourn times then occur due to a combination

of events (ii) and (iii). In particular, it is shown that the number of customers seen upon arrival by a customer with sojourn time larger than x , is of the order x .

The two-class system considered in [15], which we discussed in Section 4.2, is an example where a long sojourn time is caused by a combination of effects (i) and (ii). In Theorem 4.3 the various factors on the right-hand side of the asymptotic relation have a clear interpretation. The factor $\mathbb{P}\{B_1 > \frac{x}{M}\}$ corresponds to a sufficiently large service requirement of the customer itself. The second factor $\frac{(1-\rho)\rho_2^{M-1}}{1-\rho^{M+1}}$ entails the probability that the system occupancy is nearly saturated upon arrival of the class-1 customer. This probability, in fact, corresponds to there being sufficiently many *class-2* customers in the system. This is because of the effect embodied by the third factor $(\mathbb{P}\{B_2^r > \frac{x}{M}\})^{M-1}$: each of the customers present should stay longer than an amount of time x . For class-2 customers this is quite probable, as their remaining service requirement distribution is regularly varying, but for class-1 customers (with light-tailed service requirements) this makes such a scenario highly unlikely.

9 Concluding remarks and suggestions for further research

We have surveyed several methods for deriving sojourn time asymptotics in PS queues with heavy-tailed service requirements. In particular, we have established general necessary and sufficient conditions for a reduced service rate approximation to hold, and identified a strong relationship between such an asymptotic equivalence and a geometrically bounded queue length distribution. In addition, we have briefly discussed the case of light-tailed service requirements, which turns out to exhibit a drastically different and substantially more complicated large-deviations behavior.

Based on the existing results, several fascinating topics for further research present themselves. First of all, it would be interesting to develop a deeper understanding of the relationship between a reduced service rate approximation and a geometrically bounded queue length distribution. A second promising yet challenging direction would be to further explore the intriguing case of non-egalitarian (Discriminatory) PS and extensions to PS networks. For DPS we have established the asymptotic equivalence between the tails of the service requirement and sojourn time distributions under various assumptions in Theorems 5.1, 6.1 and 7.2. Like for PS networks, however, a full understanding of necessary conditions is yet to be developed. A related subject that seems to deserve study, is a comparison of the sojourn time asymptotics for PS with alternative service disciplines. As a final topic, the study of PS queues with light-tailed service requirements is very much worth pursuing.

References

- [1] Altman, A., Avrachenkov, K., Ayesta, U. (2006). A survey on Discriminatory Processor Sharing. In this special issue.
- [2] Asmussen, S. (2003). *Applied Probability and Queues*. Springer-Verlag, New York.
- [3] Asmussen, S., Klüppelberg, C., Sigman, K. (1999). Sampling at subexponential times, with queueing applications. *Stoch. Proc. Appl.* **79**, 265–286.

- [4] Bekker, R., Borst, S.C., Núñez-Queija, R. (2005). Performance of TCP-friendly streaming sessions in the presence of heavy-tailed elastic flows. *Perf. Evaluation* **61**, 143–162.
- [5] Van den Berg, J.L., Boxma, O.J. (1991). The M/G/1 queue with processor sharing and its relation to a feedback queue. *Queueing Systems* **33**, 43–71.
- [6] Bingham, N.H., Doney, R.A. (1974). Asymptotic properties of supercritical branching processes I: the Galton-Watson process. *Adv. Appl. Prob.* **6**, 711–731.
- [7] Bingham, N.H., Goldie, C.M., Teugels, J.L. (1987). *Regular Variation*. Cambridge University Press.
- [8] Bonald, T., Massoulié, L. (2001). Impact of fairness on Internet performance. In: *Proc. ACM Sigmetrics/Performance 2001*.
- [9] Bonald, T., Proutière, A. (2002). Insensitivity in processor sharing networks. *Perf. Evaluation* **49**, 193–209.
- [10] Bonald, T., Proutière, A. (2003). Insensitive bandwidth sharing in data networks. *Queueing Systems* **44**, 69–100.
- [11] Bonald, T., Proutière, A. (2004). On stochastic bounds for monotonic processor sharing networks. *Queueing Systems* **47**, 81–106.
- [12] Borst, S.C., Boxma, O.J., Morrison, J.A., Núñez-Queija, R. (2003). The equivalence between processor sharing and service in random order. *Oper. Res. Lett.* **31**, 254–262.
- [13] Borst, S.C., Boxma, O.J., Nunez-Queija, R., Zwart, A.P. (2003). The impact of the service discipline on delay asymptotics. *Perf. Eval.* **54**, 177–206.
- [14] Borst, S.C., Núñez-Queija, R., Van Uitert, M.J.G. (2002). User-level performance of elastic traffic in a differentiated-services environment. *Proc. Performance 2002*.
- [15] Borst, S.C., Núñez-Queija, R., Zwart, A.P. (2004). Bandwidth sharing with heterogeneous service requirements. *Ann. Telecommun.* **59**, 1297–1311.
- [16] Borst, S.C., Van Ooteghem, D.T.M.B., Zwart, A.P. (2005). Tail asymptotics for discriminatory processor-sharing queues with heavy-tailed service requirements. *Perf. Evaluation* **61**, 281–298.
- [17] Boxma, O.J., Dumas, V. (1998). The busy period in the fluid queue. *Perf. Eval. Rev.* **26**, 100–110.
- [18] Breiman, L. (1965). On some limit theorems similar to the arc-sin law. *Theory of Prob. Appl.* **10**, 351–360.
- [19] Coffman, E.G., Jr., Muntz, R.R., Trotter, H. (1970). Waiting time distribution for processor-sharing systems. *J. ACM* **17**, 123–130.
- [20] Cohen, J.W. (1979). The multiple phase service network with generalised processor sharing. *Acta Informatica* **12**, 245–284.

- [21] Crovella, M., Bestavros, A. (1996). Self-similarity in World Wide Web traffic: evidence and possible causes. In: *Proc. ACM Sigmetrics '96*, 160–169.
- [22] Delcoigne, F., Proutière, A., Régnié (2002). Modelling integration of streaming and data traffic. In: *Proc. ITC Specialist Seminar*, Würzburg, Germany.
- [23] Egorova, R., Zwart, A.P., Boxma, O.J. (2005). Sojourn time tails in the M/D/1 processor sharing queue. SPOR-Report 2005-05, Eindhoven University of Technology. *Prob. Eng. Inf. Sciences*, to appear
- [24] Embrechts, P., Klüppelberg, C., Mikosch, T. (1997). *Modeling Extremal Events*. Springer.
- [25] Fayolle, G., Mitrani, I., Iasnogorodski, R. (1980). Sharing a processor among many job classes. *J. ACM* **27**, 519–532.
- [26] Flatto, L. (1997). The waiting time distribution for the random order service M/M/1 queue. *Ann. Appl. Prob.* **7**, 382–409.
- [27] Foss, S., Korshunov, D.A. (2000). Sampling at random time with a heavy-tailed distribution. *Markov Proc. Rel. Fields* **6**, 543–568.
- [28] Guillemin, F., Robert, Ph., Zwart, A.P. (2004). Tail asymptotics for processor sharing queues. *Adv. Appl. Prob.* **36**, 525–543.
- [29] Heyman, D.P., Lakshman, T.V., Neidhardt, A.L. (1997). A new method for analysing feedback-based protocols with applications to engineering Web traffic over the Internet. In: *Proc. ACM Sigmetrics '97*, 24–38.
- [30] Jagerman, D.L., Sengupta, B. (1991). The GI/M/1 processor-sharing queue and its heavy traffic analysis. *Comm. Statist. – Stoch. Models* **7**, 379–395.
- [31] Jelenković, P.R. (1999). Network multiplexer with truncated heavy-tailed arrival streams. In: *Proc. IEEE Infocom '99*, New York, NY, USA, 625–633.
- [32] Jelenković, P.R., Momčilović, P. (2003). Large deviation analysis of subexponential waiting times in a processor sharing queue. *Math. Oper. Res.* **28**, 587–608.
- [33] Jelenković, P., Momčilović, P., Zwart, A.P. (2004). Reduced load equivalence under subexponentiality. *Queueing Systems* **46**, 97–112.
- [34] Kleinrock, L. (1964). Analysis of a time-shared processor. *Nav. Res. Log. Quarterly* **11**, 59–73.
- [35] Kleinrock, L. (1967). Time-shared systems: a theoretical treatment. *J. ACM* **14**, 242–261.
- [36] Mandjes, M.R.H., Zwart, A.P. (2004). Large deviations for sojourn times in processor sharing queues. *Queueing Systems*, to appear.
- [37] Massoulié, L., Roberts, J.W. (1999). Bandwidth sharing: Objectives and algorithms. In: *Proc. IEEE Infocom '99*, 1395–1403.

- [38] Mitra, D. (1981). Waiting time distributions from closed queueing network models of shared-processor systems. In: *Proc. Performance 1981*, ed. F.J. Kylstra, North-Holland, Amsterdam, 113–131.
- [39] Mitra, D., Morrison, J.A. (1983). Asymptotic expansions of moments of the waiting time in closed and open processor-sharing systems with multiple job classes. *Adv. Appl. Prob.* **15**, 813–839.
- [40] Morrison, J.A. (1985). Response-time distribution for a processor-sharing system. *SIAM J. Appl. Math.* **45**, 152–167.
- [41] Morrison, J.A., Mitra, D. (1985). Heavy-usage asymptotic expansions for the waiting time in closed processor-sharing systems with multiple classes. *Adv. Appl. Prob.* **17**, 163–185.
- [42] Nabe, M., Murata, M., Miyahara, H. (1998). Analysis and modelling of World Wide Web traffic for capacity dimensioning of Internet access lines. *Perf. Eval.* **34**, 249–271.
- [43] Núñez-Queija, R. (2000). *Processor-Sharing Models for Integrated-Services Networks*. Ph.D. thesis, Eindhoven University of Technology, ISBN 90-646-4667-8 (also available from the author upon request).
- [44] Núñez-Queija, R. (2002). Queues with equally heavy sojourn time and service requirement distributions. *Ann. Oper. Res.* **113**, 101–117.
- [45] Nuyens, M., Wierman, A., Zwart, A.P. (2005). Preventing large sojourn times with SMART scheduling. SPOR Report 2005-13, Eindhoven University of Technology. Submitted for publication.
- [46] Nuyens, M., Zwart, A.P. (2005). A large deviations analysis of the GI/GI/1 SRPT queue. SPOR Report 2005-06, Eindhoven University of Technology. Submitted for publication.
- [47] Van Ooteghem, D.T.M.B., Zwart, A.P., Borst, S.C. (2004). A stochastic mean-value method for the derivation of delay asymptotics in heavy-tailed processor-sharing systems. SPOR Report 2004-11, Eindhoven University of Technology.
- [48] Ott, T.J. (1984). The sojourn time distributions in the M/G/1 queue with processor sharing. *J. Appl. Prob.* **21**, 360–378.
- [49] Palmowski, Z., Rolski, T. (2004). On busy period asymptotics in the GI/GI/1 queue. Available at <http://www.math.uni.wroc.pl/~zpalma/publication.html> Submitted for publication.
- [50] Pollaczek, F. (1946). La loi d'attente des appels téléphoniques. *Comptes Rendus Acad. Sci. Paris* **222**, 352–355.
- [51] Ramaswami, V. (1984). The sojourn time in the GI/M/1 queue with processor sharing. *J. Appl. Prob.* **21**, 437–442.
- [52] Sakata, M., Noguchi, S., Oizumi, J. (1969). Analysis of a processor-shared queueing model for time-sharing systems. In: *Proc. 2nd Hawaii Int. Conf. System Sciences*, 625–628.

- [53] Sakata, M., Noguchi, S., Oizumi, J. (1971). An analysis of the M/G/1 queue under round-robin scheduling. *Oper. Res.* **19**, 371–385.
- [54] Sengupta, B., Jagerman, D.L. (1985). A conditional response time of the M/M/1 processor-sharing queue. *AT&T Techn. J.* **64**, 409–421.
- [55] Yashkov, S.F. (1987). Processor-sharing queues: Some progress in analysis. *Queueing Systems* **2**, 1–17.
- [56] Yashkov, S.F. (1992). Mathematical problems in the theory of processor-sharing queueing systems. *J. Soviet Math.* **58**, 101–147.
- [57] Zwart, A.P. (1999). Sojourn times in a multi-class processor sharing queue. In: *Teletraffic Engineering in a Competitive World, Proc. ITC-16*, Edinburgh, UK, eds. P. Key, D. Smith (North-Holland, Amsterdam), 335–344.
- [58] Zwart, A.P. (2001). Tail asymptotics for the busy period in the GI/G/1 queue. *Math. Oper. Res.* **26**, 475–483.
- [59] Zwart, A.P., Boxma, O.J. (2000). Sojourn time asymptotics in the M/G/1 processor sharing queue. *Queueing Systems* **35**, 141–166.