GPS Queues with Heterogeneous Traffic Classes

Sem Borst[†,*,‡], Michel Mandjes[†,⋆], Miranda van Uitert[†]
[†]CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
[*]Bell Laboratories, Lucent Technologies, P.O. Box 636, Murray Hill, NJ 07974, USA
[‡]Department of Mathematics & Computer Science, Eindhoven University of Technology, The Netherlands
[⋆]Faculty of Mathematical Sciences, University of Twente, The Netherlands

Abstract—**We consider a queue fed by a mixture of light-tailed and heavy-tailed traffic. The two traffic classes are served in accordance with the Generalized Processor Sharing (GPS) discipline. GPS-based scheduling algorithms, such as Weighted Fair Queueing (WFQ), have emerged as an important mechanism for achieving service differentiation in integrated networks. We derive the asymptotic workload behavior of the light-tailed class for the situation where its GPS weight is larger than its traffic intensity. The GPS mechanism ensures that the workload is bounded above by that in an isolated system with the light-tailed class served in isolation at a constant rate equal to its GPS weight. We show that the workload distribution is in fact asymptotically equivalent to that in the isolated system, multiplied with a certain pre-factor, which accounts for the interaction with the heavy-tailed class. Specifically, the pre-factor represents the probability that the heavy-tailed class is backlogged long enough for the light-tailed class to reach overflow. The results provide crucial qualitative insight in the typical overflow scenario.**

Keywords— **Generalized Processor Sharing (GPS), heavy-tailed traffic, large deviations, light-tailed traffic, Weighted Fair Queueing (WFQ), workload asymptotics.**

I. Introduction

The next-generation Internet is expected to support a wide variety of services, such as voice, video, and data applications. Voice and video communications induce far more stringent Quality-of-Service (QoS) requirements than the typical sort of data applications which currently account for the bulk of the Internet traffic. The integration of heterogeneous services thus raises the need for differentiated QoS, catering to the specific requirements of the various traffic flows.

One potential approach to achieve service differentiation is through the use of discriminatory scheduling algorithms, which distinguish between packets of various traffic streams. Because of scalability issues, it is practically infeasible though to manipulate packets at the granularity level of individual traffic flows in the core of any large-scale high-speed network. To avoid these complexity problems, traffic flows may instead be aggregated into a small number of classes with roughly similar features, with scheduling mechanisms acting at the coarser level of aggregate streams. With a little simplification, the majority of applications may for example be broadly categorized into just two classes, one containing *streaming* traffic (e.g. audio and video communications), the other one comprising *elastic* traffic (e.g. file transfers). This is a crucial element of the DiffServ proposal [5], which defines the EF class (Expedited Forwarding) for delay-sensitive traffic, and the AF class (Assured Forwarding) for traffic with some degree of delay tolerance.

In view of the delay requirements, it is desirable that streaming applications receive some sort of priority over elastic traffic, at least over short time scales. Strict priority scheduling may however not be ideal, since it may lead to starvation of the best-effort traffic. Even temporary starvation effects may cause end-to-end flow control mechanisms such as TCP to suffer a severe degradation in throughput performance. The Generalized Processor Sharing (GPS) discipline provides a potential mechanism for implementing priority scheduling in a tunable way, with strict priority scheduling as an extreme option [26]. In GPS-based scheduling algorithms, such as Weighted Fair Queueing, the link capacity is shared in proportion to certain class-defined weight factors. By setting the weight factor for the best-effort class relatively low, one can still provide some degree of priority to the streaming applications, while avoiding starvation of the elastic traffic.

Besides achieving service differentiation, scheduling mechanisms also play a role in controlling the performance impact of bursty traffic. Extensive measurements have shown that bursty traffic behavior may extend over a wide range of time scales, and may manifest itself in long-range dependence and self-similarity [21], [28]. The occurrence of these phenomena is commonly attributed to extreme variability and heavy-tailed characteristics in the traffic patterns [3], [12]. These observations have triggered a strong interest in queueing models with heavy-tailed traffic processes, see for instance [27], [31].

Although the presence of heavy-tailed traffic characteristics is widely acknowledged, the practical implications for network performance and traffic engineering remain controversial. For small buffer sizes, the effect of heavy-tailed traffic characteristics is not as dramatic as indicated by theoretical studies for infinite buffer sizes, especially at high levels of multiplexing [11], [16], [22], [30]. For large buffer sizes, flow control mechanisms such as TCP prevent heavy-tailed activity patterns from overwhelming the buffers [2].

In the present paper, we specifically examine the potential role of GPS-based scheduling mechanisms in protecting light-tailed traffic flows from the impact of heavy-tailed traffic processes. Large-deviations results for GPS models with light-tailed traffic may be found in [24], [32]. Workload asymptotics for GPS queues with heavy-tailed traffic flows were obtained in [6], [19]. The latter results show a sharp dichotomy in qualitative behavior, depending on the traffic intensities and the relative values of the weight parameters. For certain weight combinations, an individual flow with heavy-tailed traffic characteristics is effectively served at a *constant* rate, which is only influenced by the average rates of the other flows. In particular, the flow is essentially immune from excessive activity of flows with 'heavier'-tailed traffic characteristics. For other weight combinations however, a flow may be strongly affected by the activity of 'heavier'-tailed flows, and may inherit their traffic characteristics. The latter result in fact also applies for light-tailed flows when their traffic intensity exceeds their GPS weight. In the present paper, we derive the asymptotic workload behavior of the light-tailed class for the more plausible situation where its GPS weight is larger than its traffic intensity.

The remainder of the paper is organized as follows. In Section II, we present a detailed model description and state some important preliminary results. In Section III, we provide an overview of the main results of the paper, which characterize the exact asymptotic behavior of the workload distribution of the light-tailed class. The subsequent sections give a sketch of the proofs. We start in Section IV with deriving lower and upper bounds for the workload distribution of the light-tailed class. In Section V, we proceed to prove some auxiliary results for the light-tailed class in isolation. Although the bounds seem quite crude by themselves, we show in Section VI that they asymptotically coincide, yielding the exact asymptotic behavior. One of the asymptotic terms involves the probability that the heavy-tailed class is backlogged long enough for overflow to occur, which is computed in Sections VII and VIII.

II. Model description

We now present a detailed model description. We consider two traffic flows sharing a link of unit rate. Traffic from the flows is served in accordance with the Generalized Processor Sharing (GPS) discipline, which operates as follows. Flow i is assigned a weight ϕ_i, $i = 1, 2$, with $\phi_1 + \phi_2 = 1$. As long as both flows are backlogged, flow i is served at rate ϕ_i, $i = 1, 2$. If one of the flows is not backlogged, however, then the capacity is reallocated to the other flow, which is then served at the full link rate (if backlogged). (It may occur that one of the flows is not backlogged, while generating traffic at some rate $r_i < \phi_i$. In that case, only the *excess* capacity, i.e., $\phi_i - r_i$, is reallocated to the other flow.) Denote by $A_i(s, t)$ the amount of traffic generated by flow i during the time interval $(s, t]$. We assume that the process $A_i(s, t)$ is reversible and has stationary increments. Denote by $V_i(t)$ the backlog (workload) of flow i at time t. Let $\mathbf{V}_i$ be a random variable with as distribution the limiting distribution of $V_i(t)$ for $t \to \infty$ (assuming it exists). Define $B_i(s, t)$ as the amount of service received by flow i during $(s, t]$. Then the following identity relation holds, for all $s \leq t$,

$$V_i(t) = V_i(s) + A_i(s, t) - B_i(s, t). \quad (1)$$

For any $c \geq 0$, denote by $V_i^c(t) := \sup_{s \leq t}\{A_i(s, t) - c(t - s)\}$ the workload at time t in a queue of capacity c fed by flow i. Denote by ρ_i the traffic intensity of flow i (as will be defined in detail below). For $c > \rho_i$, let $\mathbf{V}_i^c$ be a random variable with as distribution the limiting distribution of $V_i^c(t)$ for $t \to \infty$. Then a similar identity relation as above holds, for all $s \leq t$,

$$V_i^c(t) = V_i^c(s) + A_i(s, t) - B_i^c(s, t). \quad (2)$$

In the next two subsections we describe the traffic model that we consider. We first introduce some additional notation. For any two real functions $g(\cdot)$ and $h(\cdot)$, we use the notational convention $g(x) \sim h(x)$ to denote $\lim_{x\to\infty} g(x)/h(x) = 1$, or equivalently, $g(x) = h(x)(1 + o(1))$ as $x \to \infty$. We use $f(x) \lesssim g(x)$ to denote $\limsup_{x\to\infty} f(x)/g(x) \leq 1$. Also, $f(x) \gtrsim g(x)$ denotes $\liminf_{x\to\infty} f(x)/g(x) \geq 1$. For any two random variables $\mathbf{X}$ and $\mathbf{Y}$, we write $\mathbf{X} \stackrel{d}{=} \mathbf{Y}$ to denote that they have the same distribution function. For any random variable $\mathbf{X}$ with distribution function $F(\cdot)$, $\mathbb{E}\{\mathbf{X}\} < \infty$, denote by $F^r(\cdot)$ the distribution function of the residual lifetime of $\mathbf{X}$, i.e., $F^r(x) = \frac{1}{\mathbb{E}\{\mathbf{X}\}} \int_0^x (1 - F(y))\mathrm{d}y$, and by $\mathbf{X}^r$ a random variable with that distribution. The classes of *subexponential*, *regularly varying*, and *intermediately regularly varying* distributions are denoted with the symbols $\mathcal{S}$, $\mathcal{R}$, and $\mathcal{IR}$, respectively. The definitions of these classes may be found in [4].

A. Traffic model flow 1

We assume that flow 1 is light-tailed. Specifically, we make the assumption that the input process $A_1(s, t)$ is a *Markov-modulated fluid*. Such a process can be described as follows. There is an irreducible Markov chain with a finite state space $\{1, 2, \ldots, d\}$. The corresponding transition rate matrix is denoted by $\Lambda := (\lambda_{ij})_{i,j=1,\ldots,d}$, where we follow the convention that $\lambda_{ii} := -\sum_{j\neq i} \lambda_{ij}$. Since the Markov chain is irreducible, there is a unique stationary distribution, which we denote by the vector π. When the flow is in state i, traffic is generated (as fluid) at constant rate $R_i < \infty$. Let R be the diagonal matrix with the coefficients R_i on the diagonal. Denote the mean rate by $\rho_1 := \sum_{i=1}^d \pi_i R_i$. Denote the peak rate by $R_P := \max_{i=1,\ldots,d} R_i$. It is important to observe that the class of Markov fluid input is closed under superposition, i.e., the superposition of Markov fluid flows can again be modeled as a Markov fluid flow.

Results from Kosten [20], Kesidis *et al.* [18], and Elwalid & Mitra [14] yield the following standard properties.

Property II.1: Take $\rho_1 < c_1 < R_P$. Then

- The moment generating function of traffic generated in an interval of length t is given by, in matrix notation,

$$\mathbb{E}\{\exp(sA_1(0, t))\} = \pi \exp((\Lambda + sR)t)\mathbf{1},$$

with $\mathbf{1}$ the all one vector of dimension d.

- There exists a limiting moment generating function:

$$\frac{1}{t} \log \mathbb{E}\{\exp(s(A_1(0, t) - c_1 t))\} \to M_{c_1}(s).$$

This function is continuous and differentiable. It also holds that there is a finite C such that

$$\mathbb{E}\{\exp(s(A_1(0, t) - c_1 t))\} \leq C\mathrm{e}^{M_{c_1}(s)t}.$$

- The large-buffer asymptotics of a queue with Markov fluid input are given by

$$\liminf_{x\to\infty} \frac{1}{x} \log \mathbb{P}\{\mathbf{V}_1^{c_1} > x\} = -s^*(c_1).$$

Here $s^*(c_1)$ is the unique positive root of $M_{c_1}(s) = 0$. Moreover, $M'_{c_1}(s^*(c_1)) > 0$.

Although we restrict ourselves to Markov fluid input, we believe that our results are valid for a more general class of light-tailed input. We will comment on this issue in Remark 5.1.

B. Traffic model flow 2

We assume that flow 2 is heavy-tailed. We make the assumption that the input process $A_2(s,t)$ is either instantaneous or On-Off, with heavy-tailed burst sizes or On-periods, respectively.

B.1 Instantaneous input

Here, flow 2 generates instantaneous traffic bursts according to a renewal process. The interarrival times between bursts have distribution function $U_2(\cdot)$ with mean $1/\lambda_2$. The burst sizes $\mathbf{B}_2$ have distribution function $B_2(\cdot)$ with mean $\beta_2 < \infty$. Thus, the traffic intensity is $\rho_2 := \lambda_2\beta_2$. We assume that $B_2(\cdot)$ is regularly varying of index $-\nu_2$, i.e., $B_2(\cdot) \in \mathcal{R}_{-\nu_2}$ for some $\nu_2 > 1$. The next result which is due to Pakes [25] then yields the tail behavior of the workload distribution of flow 2 in isolation.

Theorem II.1: If $B_2^r(\cdot) \in \mathcal{S}$, and $\rho_2 < c$, then

$$\mathbb{P}\{\mathbf{V}_2^c > x\} \sim \frac{\rho_2}{c-\rho_2}\mathbb{P}\{\mathbf{B}_2^r > x\}.$$

B.2 Fluid input

Here, flow 2 generates traffic according to an On-Off process, alternating between On- and Off-periods. The Off-periods $\mathbf{U}_2$ have distribution function $U_2(\cdot)$ with mean $1/\lambda_2$. The On-periods $\mathbf{A}_2$ have distribution function $A_2(\cdot)$ with mean $\alpha_2 < \infty$. While On, flow i produces traffic at constant rate r_2, so the mean burst size is $\alpha_2 r_2$. The fraction of time that flow 2 is Off is $p_2 = 1/(1+\lambda_2\alpha_2)$. The traffic intensity is $\rho_2 = (1-p_2)r_2 = (\lambda_2\alpha_2 r_2)/(1+\lambda_2\alpha_2)$.

We assume that $A_2(\cdot)$ is regularly varying of index $-\nu_2$, i.e., $A_2(\cdot) \in \mathcal{R}_{-\nu_2}$ for some $\nu_2 > 1$. The next result which is due to Jelenković & Lazar [17] then yields the tail behavior of the workload distribution of flow 2 in isolation.

Theorem II.2: If $A_2^r(\cdot) \in \mathcal{S}$, and $\rho_2 < c < r_2$, then

$$\mathbb{P}\{\mathbf{V}_2^c > x\} \sim p_2\frac{\rho_2}{c-\rho_2}\mathbb{P}\{\mathbf{A}_2^r > \frac{x}{r_2-c}\}.$$

III. Overview of the results

In this section we provide an overview of the main results of the paper which characterize the exact asymptotic behavior of $\mathbb{P}\{\mathbf{V}_1 > x\}$ as $x \to \infty$. At the end of this section, we present an example. Throughout, we assume $\rho_i < \phi_i$, $i = 1,2$, which ensures stability of both flows. In addition, we make the assumption that $r_2 > \phi_2$ in case of fluid input of flow 2. Otherwise, the workload of flow 2 would be zero, so the workload of flow 1 would be equal to the total workload $\mathbf{V}$. The tail distribution of the latter quantity has been obtained in [9].

To put things in perspective, we first briefly review the case that $\rho_1 > \phi_1$, while $\rho_1 + \rho_2 < 1$. If either (i) $B_2^r(\cdot) \in \mathcal{IR}$ (instantaneous input of flow 2), or (ii) $A_2^r(\cdot) \in \mathcal{IR}$, $r_2 > \phi_2$ (fluid input), then from [6],

$$\mathbb{P}\{\mathbf{V}_1 > x\} \sim \frac{\phi_2-\rho_2}{\phi_2}\frac{\rho_2}{1-\rho_1-\rho_2}\mathbb{P}\{\mathbf{P}_2^r > \frac{x}{\rho_1-\phi_1}\},$$

with $\mathbf{P}_2$ a random variable with as distribution the busy-period distribution in a queue of constant capacity ϕ_2 fed by flow 2.

The above result suggests that the most likely way for flow 1 to build a large queue is that flow 2 generates a large burst, or

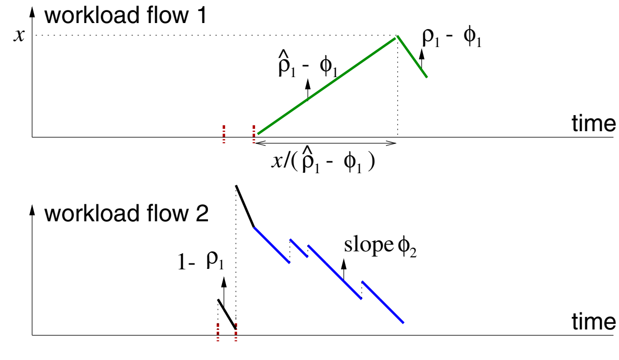

Fig. 1. Overflow scenario - instantaneous input for flow 2.

experiences a long On-period, while flow 1 itself shows roughly average behavior. Note that when flow 2 produces a large amount of traffic, so it becomes backlogged for a long period of time, it receives service at rate ϕ_2. Thus it will experience a busy period as if it were served at constant rate ϕ_2. During that period, flow 1 receives service at rate ϕ_1, while it generates traffic roughly at rate ρ_1, so its queue will grow approximately at rate $\rho_1 - \phi_1$. When flow 2 is not backlogged, its queue will drain approximately at rate $1 - \rho_1$.

Thus, the backlog of flow 2 behaves as that in a queue of constant capacity $1 - \rho_1$ fed by an On-Off source with as On- and Off-periods the busy and idle periods of flow 2 when served at constant rate ϕ_2, respectively. That is reflected in the above result if we use Theorem II.2 to interpret the right-hand side.

We now focus on the case $\rho_1 < \phi_1$. Before presenting the main result, we first provide a heuristic derivation of the asymptotic behavior of $\mathbb{P}\{\mathbf{V}_1 > x\}$ based on large-deviations arguments, see for instance Anantharam [1]. The overflow scenario described above for the case $\rho_1 > \phi_1$ cannot occur, and now flow 1 too must deviate from its 'normal' behavior in order for the queue to grow. Specifically, large-deviations results suggest that flow 1 must behave as if its traffic intensity is temporarily increased from ρ_1 to some larger value $\hat{\rho}_1$ (as will be specified below). During that period, flow 2 is continuously backlogged, consuming capacity ϕ_2, thus leaving capacity ϕ_1 for flow 1. (Notice that if flow 2 were not permanently backlogged, then flow 1 would have to show even greater anomalous activity in order for a given backlog level to occur.) Prior to that period, flow 1 shows 'normal' behavior, leaving an average service rate of $1 - \rho_1$ for flow 2.

To summarize, the intuitive argument is as follows (see Figure 1): a large backlog of level x of flow 1 occurs as a consequence of two rare events: (i) Flow 1 shows similar 'abnormal' behavior as is the typical cause of overflow when served in isolation, thus behaving as if its traffic intensity is increased from ρ_1 to $\hat{\rho}_1$ for a period of time $x/(\hat{\rho}_1 - \phi_1)$. (ii) During that period, flow 2 is constantly backlogged, demanding capacity ϕ_2, with ϕ_1 remaining for flow 1. As we will see later, the persistent backlog is most likely caused by flow 2 generating a large burst or initiating a long On-period prior to that period.

These considerations lead to the following characterization of the asymptotic behavior of $\mathbb{P}\{\mathbf{V}_1 > x\}$:

$$\mathbb{P}\{\mathbf{V}_1 > x\} \sim \mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\}. \quad (3)$$

The second term represents the probability that flow 2 is continuously backlogged during a period of time $x/(\hat{\rho}_1 - \phi_1)$. Here $\mathbf{T}_2^{1-\rho_1}$ is a random variable with as distribution the limiting distribution of $T_2^{1-\rho_1}(t)$ for $t \to \infty$, with

$$T_2^c(t) := \inf\{u \geq 0 : V_2^c(t) + A_2(t, t+u) - \phi_2 u = 0\}$$

representing the drain time in a queue of capacity ϕ_2 fed by flow 2 with initial workload $V_2^c(t)$.

Thus, the workload distribution is asymptotically equivalent to that in an isolated system, multiplied with a certain pre-factor. The isolated system consists of flow 1 served in isolation at constant rate ϕ_1. The pre-factor represents the probability that flow 2 is backlogged long enough for flow 1 to reach overflow. The combination of light-tailed and heavy-tailed large deviations is similar to that in the 'reduced-peak equivalence' result derived in Borst & Zwart [9] as well as that for an M/G/2 queue with heterogeneous servers studied in Boxma *et al.* [10].

Note that the general decompositional form of (3) holds irrespective of the detailed traffic characteristics of the two flows. However, the specific form of the two individual terms in (3) *does* depend on the detailed properties of the traffic processes. In particular, we need to distinguish whether flow 2 generates instantaneous or fluid input. In the latter case, it also depends on whether the peak rate r_2 exceeds $1 - \rho_1$ or not.

We now state the main theorem of the paper.

Theorem III.1: Suppose that the input process $A_1(s,t)$ satisfies Property II.1 and that the input process $A_2(s,t)$ is either instantaneous or On-Off, with regularly varying burst sizes or On-periods, respectively. Assume that $\rho_i < \phi_i$, $i = 1, 2$, and $r_2 > \phi_2$ in case of fluid input of flow 2. Then, with $\hat{\rho}_1 := M'_{\phi_1}(s^*(\phi_1)) + \phi_1$,

$$\mathbb{P}\{\mathbf{V}_1 > x\} \sim \mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1 - \phi_1}\}.$$

Case I: If $B_2^r(\cdot) \in \mathcal{R}_{-\nu_2}$ (instantaneous input), then

$$\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > x\} \sim \frac{\rho_2}{1-\rho_1-\rho_2}\mathbb{P}\{\mathbf{B}_2^r > x(\phi_2 - \rho_2)\}. \quad (4)$$

Case II-A: If $A_2^r(\cdot) \in \mathcal{R}_{-\nu_2}$, $r_2 < 1 - \rho_1$ (fluid input), then

$$\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > x\} \sim (1 - p_2)\mathbb{P}\{\mathbf{A}_2^r > \frac{x(\phi_2 - \rho_2)}{r_2 - \rho_2}\}. \quad (5)$$

Case II-B: If $A_2^r(\cdot) \in \mathcal{R}_{-\nu_2}$, $r_2 > 1 - \rho_1$ (fluid input), then

$$\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > x\} \sim p_2 \frac{\rho_2}{1-\rho_1-\rho_2}\mathbb{P}\{\mathbf{A}_2^r > \frac{x(\phi_2-\rho_2)}{r_2 - \rho_2}\}. \quad (6)$$

Noting that $p_2\rho_2 = (1 - p_2)(r_2 - \rho_2)$, we can observe that in the limiting regime $r_2 \to 1 - \rho_1$, cases II-A and II-B coincide. Also, case I can be seen as the limiting case of II-B if we use $r_2\mathbf{A}_2 = \mathbf{B}_2$ and let $r_2 \to \infty$ so that $p_2 \downarrow 1$. In [7] a qualitatively similar result as in case I is derived for a system with two coupled queues.

Before proceeding to the formal proof of Theorem III.1, we first give an example. Assume flow 1 to behave according to an On-Off process with exponentially distributed On- and Off-periods with means $1/\mu_1$ and $1/\mu_2$, respectively. When the flow is in the On-state, it generates traffic at rate R_1. We assume flow 2 to generate instantaneous input with regularly varying burst sizes of index $-\nu_2$, i.e., $\mathbb{P}\{\mathbf{B}_2 > x\} \sim C_2 x^{-\nu_2} l_2(x)$, with $l_2(\cdot)$ some slowly varying function. First we determine the deviant traffic intensity $\hat{\rho}_1$ using [23],

$$\hat{\rho}_1 = \frac{R_1\phi_1^2}{\mu_2} \Big/ \left(\frac{\phi_1^2}{\mu_2} + \frac{(R_1 - \phi_1)^2}{\mu_1}\right).$$

Using [13], we obtain for flow 1,

$$\mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\} \sim \frac{R_1}{\phi_1}\frac{\mu_2}{\mu_1 + \mu_2}\exp(-(\frac{\mu_1}{R_1 - \phi_1} + \frac{\mu_2}{\phi_1})x).$$

For flow 2, from (4), $\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > x\} \sim$

$$\frac{\rho_2}{1 - \rho_1 - \rho_2}\frac{C_2}{\beta_2(\nu_2 - 1)}(x(\phi_2 - \rho_2))^{1-\nu_2} l_2(x(\phi_2 - \rho_2)).$$

This provides all the ingredients for $\mathbb{P}\{\mathbf{V}_1 > x\}$ as required in Theorem III.1.

The next sections are devoted to the formal proof of Theorem III.1. We start in Section IV by deriving lower and upper bounds for the workload distribution of flow 1. We then proceed in Section V to prove some auxiliary results for flow 1 in isolation. Although the bounds derived in Section IV seem quite crude by themselves, we show in Section VI that they asymptotically coincide, yielding the exact asymptotic behavior of $\mathbb{P}\{\mathbf{V}_1 > x\}$.

In order to determine the drain time distribution of flow 2 as specified in Theorem III.1, we first establish in Section VII some preliminary results for flow 2 in isolation. Note that the specific form of the drain time distribution depends on whether flow 2 generates instantaneous or fluid input. In the latter case, we also need to distinguish whether the peak rate r_2 exceeds $1 - \rho_1$ or not. We calculate the drain time distribution for the case of an instantaneous input process in Section VIII. Due to space constraints, we omit the corresponding analysis for fluid input processes, see [8] for details.

IV. Bounds

In this section we derive lower and upper bounds for the workload distribution of flow 1. The bounds will be instrumental in obtaining the asymptotic behavior of $\mathbb{P}\{\mathbf{V}_1 > x\}$ as given in Theorem III.1. We refer to [8] for detailed proofs of the lemmas in this section.

A. Lower bound

We start with a lower bound for the workload distribution of flow 1. The main idea (see Figure 2) is that the following scenario is sufficient for the event $V_1(t) > x$ to occur (in fact, is the only plausible one, as we will see later). Flow 1 starts to build up at some time s^*, and hence is constantly backlogged throughout the time interval $[s^*, t]$. Flow 2 is also continuously backlogged during $[s^*, t]$. Thus, during that time period, flows 1 and 2 both receive service at rates ϕ_1 and ϕ_2, respectively. Flow 2 already

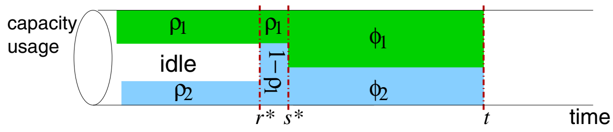

Fig. 2. Intuitive idea lower bound.

becomes backlogged at time $r^* \leq s^*$, and receives service approximately at rate $1-\rho_1$ during $[r^*, s^*]$, while flow 1 then shows roughly average behavior.

Lemma IV.1: Suppose $r^* \leq s^* \leq t$ and y exist such that

$$A_1(s^*, t) - \phi_1(t - s^*) > x,$$

$$A_1(r^*, s^*) - (\rho_1 - \epsilon)(s^* - r^*) \geq -y \text{ and}$$

$$\inf_{s^* \leq u \leq t} \{A_2(r^*, u) - (1 - \rho_1 + \epsilon)(s^* - r^*) - \phi_2(u - s^*)\} \geq y,$$

then $V_1(t) > x$.

Proof By definition, $B_1(s,t) + B_2(s,t) \leq t - s$. Using (1), the GPS discipline implies that $B_2(s,t) \geq$

$$\min\{\phi_2(t-s), V_2(s) + \inf_{s \leq u \leq t}\{A_2(s,u) + \phi_2(t-u)\}\}.$$

Substituting, we have that $V_1(t)$ is bounded below by

$$A_1(s,t) - \phi_1(t-s) +$$

$$\min\{0, V_1(s) + V_2(s) + \inf_{s \leq u \leq t}\{A_2(s,u) - \phi_2(u-s)\}\}.$$

By definition, $B_1(r,s) + B_2(r,s) \leq s - r$. Thus, from (1), for all $r \leq s$,

$$V_1(s) + V_2(s) \geq A_1(r,s) + A_2(r,s) - (s-r).$$

Substituting, we find that $V_1(t)$ is bounded below by

$$A_1(s,t) - \phi_1(t-s) + \min\{0, A_1(r,s) - (\rho_1 - \epsilon)(s-r) +$$

$$\inf_{s \leq u \leq t}\{A_2(r,u) - (1 - \rho_1 + \epsilon)(s-r) - \phi_2(u-s)\}\}$$

for all $r \leq s \leq t$. □

We now translate the above sample-path result into a probabilistic lower bound. We first introduce some additional notation. Define $\mathbf{V}_i^c(w) := \sup_{0 \leq s \leq w}\{A_i(-s,0) - cs\}$ for any $c, w \geq 0$. Note that $\mathbf{V}_i^c(\infty) \stackrel{d}{=} \mathbf{V}_i^c$ for $c > \rho_i$, as defined earlier. For any c, $v \geq 0$, and y, define $\mathbf{T}_2^c(v,y) :=$

$$\inf\{u \geq 0 : \sup_{0 \leq r \leq v}\{A_2(-r,0) - cr\} + A_2(0,u) - \phi_2 u \leq y\}.$$

Thus, $\mathbf{T}_2^c(v,y)$ represents the drain time in a queue of capacity ϕ_2 fed by flow 2 with initial workload $\sup_{0 \leq r \leq v}\{A_2(-r,0) - cr\} - y$. Define, for $c > \rho_2$, $\mathbf{T}_2^c(y) :=$

$$\mathbf{T}_2^c(\infty, y) = \inf\{u \geq 0 : V_2^c(0) + A_2(0,u) - \phi_2 u \leq y\},$$

and note that $\mathbf{T}_2^c(0) \stackrel{d}{=} \mathbf{T}_2^c$ as defined earlier. Also, define

$$\mathbf{T}_2(y) := \mathbf{T}_2^c(0,y) = \inf\{u \geq 0 : A_2(0,u) - \phi_2 u \leq y\}.$$

(note that the latter quantity does not depend on the value of c), and $\mathbf{T}_2 := \mathbf{T}_2(0)$. Denote $P^{\rho_1 - \epsilon}(s^*, v, x, y) :=$

$$\mathbb{P}\{\sup_{s^* - v \leq r \leq s^*}\{(\rho_1 - \epsilon)(s^* - r) - A_1(r, s^*)\} \leq y | A_1(s^*, 0) + \phi_1 s^* > x\}.$$

Corollary IV.1: For any $v \geq 0$ and y,

$$\mathbb{P}\{\mathbf{V}_1 > x\} \geq \mathbb{P}\{\mathbf{V}_1^{\phi_1}(\frac{(1+\alpha)x}{\hat{\rho}_1 - \phi_1}) > x\} P^{\rho_1 - \epsilon}(s^*, v, x, y)$$

$$\mathbb{P}\{\mathbf{T}_2^{1-\rho_1+\epsilon}(v,y) > \frac{(1+\alpha)x}{\hat{\rho}_1 - \phi_1}\}.$$

Proof The proof follows using Lemma IV.1, the independence of $A_1(s,t)$ and $A_2(s,t)$, and the fact that $A_1(s,t)$ and $A_2(s,t)$ have stationary increments. □

B. Upper bound

We proceed to derive an upper bound for the workload distribution of flow 1. The idea is that the lower-bound scenario described above is basically also necessary for the event $V_1(t) > x$ to occur, unless at least one of two other events happen both of which however we will later show are significantly less likely.

Lemma IV.2: Suppose $V_1(t) > x$. Then for all y there exist $r^* \leq s^* \leq t$ such that

$$A_1(s^*, t) - \phi_1(t - s^*) > x, \quad (7)$$

and at least one of the three following events occurs, either

$$A_1(r^*, s^*) - (\rho_1 + \epsilon)(s^* - r^*) > y, \text{ or} \quad (8)$$

$$V_1^{\phi_1}(t) > x + y, \text{ or} \quad (9)$$

$$\inf_{s^* \leq u \leq t}\{A_2(r^*, u) - (1 - \rho_1 - \epsilon)(s^* - r^*) - \phi_2(u - s^*)\} > -2y. \quad (10)$$

Proof Because of the GPS discipline, (7) is implied by $V_1(t) > x$. Hence, there exists an $s \leq t$ such that $A_1(s,t) - \phi_1(t-s) > x$. Define $s^* := \inf\{s : A_1(u,t) - \phi_1(t-u) \leq x \ \forall u > s\}$. Note that flow 1 is continuously backlogged during $[s^*, t]$. It can be shown that $V_1(t) > x$ implies that either (9) holds or $\forall u \in [s^*, t] : B_2(s^*, u) - \phi_2(u - s^*) > -y$. It may be verified that the latter event implies that either (8) or (10) holds. □

We now use the above sample-path relation to obtain a probabilistic upper bound. Denote $Q^{\rho_1 + \epsilon}(s^*, x, y) :=$

$$\mathbb{P}\{\sup_{r \leq s^*}\{A_1(r, s^*) - (\rho_1 + \epsilon)(s^* - r)\} > y \mid A_1(s^*, 0) + \phi_1 s^* > x\}.$$

Corollary IV.2: For any y,

$$\mathbb{P}\{\mathbf{V}_1 > x\} \leq \mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\} Q^{\rho_1 + \epsilon}(s^*, x, y)$$

$$+ \mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\} \mathbb{P}\{\mathbf{T}_2^{1-\rho_1-\epsilon}(-2y) > \frac{(1-\alpha)x}{\hat{\rho}_1 - \phi_1}\}$$

$$+ \mathbb{P}\{\mathbf{V}_1^{\phi_1} > x + y\} + \mathbb{P}\{\mathbf{V}_1^{\phi_1}(\frac{(1-\alpha)x}{\hat{\rho}_1 - \phi_1}) > x\}.$$

Proof The proof follows using Lemma IV.2, the independence of $A_1(s,t)$ and $A_2(s,t)$, and the fact that $A_1(s,t)$ and $A_2(s,t)$ have stationary increments. □

V. PRELIMINARY RESULTS FOR FLOW 1

In this section we prove some auxiliary results for flow 1 in isolation. The results will be crucial in obtaining the asymptotic behavior of $\mathbb{P}\{\mathbf{V}_1 > x\}$ in the GPS model as given in Theorem III.1. We refer to [8] for detailed proofs.

The following result is proven in [9], for a more general class of input processes.

Proposition V.1: If Property II.1 holds with $c_1 = \phi_1$, then, for any $\alpha > 0$, and with $\hat{\rho}_1 := M'_{\phi_1}(s^*(\phi_1)) + \phi_1$,

$$\liminf_{x\to\infty} \frac{\mathbb{P}\{\mathbf{V}_1^{\phi_1}(\frac{(1+\alpha)x}{\hat{\rho}_1-\phi_1}) > x\}}{\mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}} = 1. \quad (11)$$

Lemma V.1: For any $\gamma > 0$, $\epsilon > 0$, $t^* < 0$,

$$\lim_{x\to\infty} \mathbb{P}\{\sup_{r\leq t^*}\{(\rho_1-\epsilon)(t^*-r) - A_1(r,t^*)\} \leq \gamma x \mid A_1(t^*,0) + \phi_1 t^* > x\} = 1.$$

Proof Recall that flow 1 is a Markov fluid source. We condition on the state of the underlying Markov chain at time t^*. Let $E_j(t^*)$ be the event that the state at time t^* is j, $j = 1,\ldots,d$, and $\pi_j(t^*) := \mathbb{P}\{E_j(t^*) \mid A_1(t^*,0) + \phi_1 t^* > x\}$. Then the probability of interest equals

$$\sum_{j=1}^{d} \mathbb{P}\{\sup_{r\leq t^*}\{(\rho_1-\epsilon)(t^*-r) - A_1(r,t^*)\} \leq \gamma x \mid E_j(t^*)\}\pi_j(t^*).$$

The stated then follows by observing that $\forall j = 1,\ldots,d$,

$$\lim_{x\to\infty} \mathbb{P}\{\sup_{r\leq t^*}\{(\rho_1-\epsilon)(t^*-r) - A_1(r,t^*)\} \leq \gamma x \mid E_j(t^*)\} = 1,$$

since $\mathbb{E}\{A_1(-t,0)\} = \rho_1 t$. □

Lemma V.2: For any $\gamma > 0$, $\epsilon > 0$, $\mu > 0$, $t^* < 0$,

$$\lim_{x\to\infty} x^\mu \mathbb{P}\{\sup_{r\leq t^*}\{A_1(r,t^*) - (\rho_1+\epsilon)(t^*-r)\} > \gamma x \mid A_1(t^*,0) + \phi_1 t^* > x\} = 0.$$

Proof Again, condition on the state of the underlying Markov chain at time t^*. Under this condition, the event $\{A_1(t^*,0) + \phi_1 t^* > x\}$ does not provide any extra information. The fact that there exist constants C, a (independent of j) such that [23, Section 4]

$$\mathbb{P}\{\sup_{r\leq t^*}\{A_1(r,t^*) - (\rho_1+\epsilon)(t^*-r)\} > \gamma x \mid E_j(t^*)\} \leq C\mathrm{e}^{-ax}$$

proves the stated. □

Lemma V.3: For any $\gamma > 0$, $\mu > 0$,

$$\limsup_{x\to\infty} \frac{x^\mu \mathbb{P}\{\mathbf{V}_1^{\phi_1} > (1+\gamma)x\}}{\mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}} = 0.$$

Proof The proof follows immediately from the fact that $\mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}$ decays exponentially at rate s^*, where $s^* > 0$ is the solution of $M_{\phi_1}(s) = 0$ [20]. □

Lemma V.4: For any $\alpha > 0$, $\mu > 0$,

$$\limsup_{x\to\infty} \frac{x^\mu \mathbb{P}\{\mathbf{V}_1^{\phi_1}(\frac{(1-\alpha)x}{\hat{\rho}_1-\phi_1}) > x\}}{\mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}} = 0. \quad (12)$$

Proof The proof consists of three steps.

(i) As shown in [8], for $T_x(\alpha) := \lceil(1-\alpha)x/(\hat{\rho}_1 - \phi_1)\rceil$, the probability in the numerator of (12) is bounded by

$$\sum_{t=0}^{T_x(\alpha)} \mathbb{P}\{A_1(0,t) - \phi_1 t > x - (R_P - \phi_1)\}.$$

This immediately leads to

$$\limsup_{x\to\infty} \frac{1}{x}\log \sum_{t=0}^{T_x(\alpha)} \mathbb{P}\{A_1(0,t) - \phi_1 t > x - (R_P - \phi_1)\} \leq$$

$$\limsup_{x\to\infty} \frac{1}{x}\log \sup_{t\in[S_x,T_x(\alpha)]} \mathbb{P}\{A_1(0,t) - \phi_1 t > x - (R_P - \phi_1)\} \quad (13)$$

with $S_x := (x - R_P)/(R_P - \phi_1)$. Notice that we can indeed exclude all t smaller than S_x from the optimization, because in that range no overflow is possible. Clearly, we have proven the stated if we show that the latter decay rate is strictly smaller than s^* for all $\alpha > 0$ since the denominator decays at rate s^*.

(ii) For x large enough, and all t between S_x and $T_x(\alpha)$, due to Chebychev's inequality, and Property II.1,

$$\mathbb{P}\{A_1(0,t) - \phi_1 t > x - (r_P - \phi_1)\} \leq C \inf_{s>0} \frac{\mathrm{e}^{M_{\phi_1}(s)t}}{\mathrm{e}^{s(x-(r_P-\phi_1))}}.$$

Now replace t in (13) by $t_x(\beta) = (1-\beta)x/(\hat{\rho}_1 - \phi_1)$; then the supremum is over $\beta \in [\alpha, 1]$. The infimum over $s > 0$, call the solution $s^*(\beta)$, is calculated by differentiation. Using the observations that $M_{\phi_1}(s^*) = 0$ and $M'_{\phi_1}(s^*) = \hat{\rho}_1 - \phi_1 > 0$ (Property II.1), together with Taylor expansions around s^* of $M'_{\phi_1}(s)$, $M_{\phi_1}(s^*(\beta))$ and $s^*(\beta)$, we obtain

$$\lim_{x\to\infty} \frac{1}{x}\log \inf_{s>0} \frac{\mathrm{e}^{t_x(\beta)M_{\phi_1}(s)}}{\mathrm{e}^{s(x-(r_P-\phi_1))}} = -\frac{\hat{\rho}_1 - \phi_1}{M''_{\phi_1}(s^*)}\beta^2 - s^*.$$

Note that $M''_{\phi_1}(s^*) > 0$ because of convexity.

(iii) Recall that we have to perform the optimization over $\beta \in [\alpha, 1]$. The supremum over β is clearly attained at $\beta = \alpha > 0$. Since the supremum is strictly smaller than s^*, we have proved the stated. □

Remark 5.1 The results of Glynn & Whitt [15] suggest that the derived properties hold for a more general class of arrival processes than just Markov fluid. Upon inspection of the proofs in the present section, we see that only two properties were explicitly exploited: (1) the sources have bounded peak rates, and (2) the 'mild dependence' between $A_1(r,t^*)$ and $A_1(t^*,0)$.

VI. ASYMPTOTIC ANALYSIS

We now use the results from the previous section to show that the lower and upper bounds for $\mathbb{P}\{\mathbf{V}_1 > x\}$ of Section IV asymptotically coincide, resulting in the decompositional form of (3). For the proof, we need to make certain assumptions on the behavior of the drain time distribution $\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\}$. In later sections, we will determine the specific form of the drain time distribution, and find that flow 2 indeed satisfies these assumptions. For notational convenience, we frequently switch to a variable $\hat{x}$, which should be thought of as playing the role of $\frac{x}{\hat{\rho}_1-\phi_1}$.

Lemma VI.1: Suppose that the input process $A_1(s,t)$ satisfies Property II.1 with $c_1 = \phi_1$ and that flow 2 satisfies Assumptions VI.1-VI.3 listed below with $c = 1 - \rho_1$. Assume that $\rho_i < \phi_i$, $i = 1, 2$, and $r_2 > \phi_2$ in case of fluid input of flow 2. Then

$$\mathbb{P}\{\mathbf{V}_1 > x\} \sim \mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1 - \phi_1}\}.$$

Assumption VI.1: For any $\alpha, \gamma, \epsilon > 0$, either (a)

$$\liminf_{\hat{x}\to\infty} \frac{\mathbb{P}\{\mathbf{T}_2^{c+\epsilon}(\gamma\hat{x}) > (1+\alpha)\hat{x}\}}{\mathbb{P}\{\mathbf{T}_2^c > \hat{x}\}} = F^c(\alpha,\gamma,\epsilon),$$

with $\lim_{\alpha,\gamma,\epsilon\downarrow 0} F^c(\alpha,\gamma,\epsilon) = 1$, or (b)

$$\liminf_{\hat{x}\to\infty} \frac{\mathbb{P}\{\mathbf{T}_2 > (1+\alpha)\hat{x}\}}{\mathbb{P}\{\mathbf{T}_2^c > \hat{x}\}} = F(\alpha), \text{ with } \lim_{\alpha\downarrow 0} F(\alpha) = 1.$$

Assumption VI.2: For any $\alpha > 0$, $\gamma > 0$, $\epsilon > 0$,

$$\limsup_{\hat{x}\to\infty} \frac{\mathbb{P}\{\mathbf{T}_2^{c-\epsilon}(-\gamma\hat{x}) > (1-\alpha)\hat{x}\}}{\mathbb{P}\{\mathbf{T}_2^c > \hat{x}\}} = G^c(\alpha,\gamma,\epsilon),$$

with $\lim_{\alpha,\gamma,\epsilon\downarrow 0} G^c(\alpha,\gamma,\epsilon) = 1$.

Assumption VI.3: For some $\mu > 0$,

$$\liminf_{x\to\infty} \hat{x}^{\mu}\mathbb{P}\{\mathbf{T}_2^c > \hat{x}\} \geq 1.$$

Proof of Lemma VI.1 The proof consists of a lower bound and an upper bound which asymptotically coincide.

We start with the lower bound. We distinguish between two cases: Assumption VI.1 (a); Assumption VI.1 (b).

(a) Using Corollary IV.1 with $v = \infty$, $y = \frac{\gamma x}{\hat{\rho}_1-\phi_1}$, Proposition V.1, and Lemma V.1,

$$\liminf_{x\to\infty} \frac{\mathbb{P}\{\mathbf{V}_1 > x\}}{\mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\}} \geq F^{1-\rho_1}(\alpha,\gamma,\epsilon).$$

Letting $\alpha, \gamma, \epsilon \downarrow 0$ completes the proof.

(b) Using Corollary IV.1 with $v = 0$, $y = 0$, and Proposition V.1, noting that $P^{\rho_1-\epsilon}(s^*, 0, x, 0) = 1$,

$$\liminf_{x\to\infty} \frac{\mathbb{P}\{\mathbf{V}_1 > x\}}{\mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\}} \geq F(\alpha).$$

Then let $\alpha \downarrow 0$.

We now turn to the upper bound. Using Corollary IV.2 with $v = \infty$, $y = \frac{\gamma x}{2(\hat{\rho}_1-\phi_1)}$, Lemmas V.2-V.4, and Assumptions VI.2, VI.3, for some $\mu > 0$,

$$\limsup_{x\to\infty} \frac{\mathbb{P}\{\mathbf{V}_1 > x\}}{\mathbb{P}\{\mathbf{V}_1^{\phi_1} > x\}\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\}} \leq G^{1-\rho_1}(\alpha,\gamma,\epsilon).$$

Letting $\alpha, \gamma, \epsilon \downarrow 0$ completes the proof. □

In order to complete the proof of Theorem III.1, it remains to be shown that flow 2 satisfies Assumptions VI.1-VI.3 above, with $\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\}$ as in (4)-(6). This is done in the following two sections. Due to space limitations, we focus on the case of instantaneous input processes. We refer to [8] for the corresponding analysis for fluid input processes.

VII. Preliminary results for flow 2

To determine the behavior of $\mathbb{P}\{\mathbf{T}_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}\}$ as $x \to \infty$, we will reduce the space of all relevant sample paths to a single most-likely scenario, which occurs with overwhelming probability. In this section, we establish some preliminary results which we will use to neglect the contribution of all non-dominant scenarios.

Large-deviations arguments for heavy-tailed distributions suggest that a persistent backlog as associated with the event $\mathbf{T}_2^{1-\rho_1} > \frac{x}{\hat{\rho}_1-\phi_1}$, for large x, is most likely due to just a single large burst. To formalize this idea, we first introduce some additional notation. A burst is called large if the size exceeds $\kappa\hat{x}$, with $\kappa > 0$ some small constant, independent of $\hat{x}$. Denote by $\mathcal{N}_{\kappa\hat{x}}[l,r]$ the number of large bursts of flow 2 arriving in the time interval $[l,r]$. Define $N(t) := \{n : \mathbf{U}_{20}^r + \sum_{i=1}^n \mathbf{U}_{2i} \leq t\}$ as the total number of bursts of flow 2 arriving in the time interval $[0,t]$.

We now state a crucial lemma which will allow us to limit the attention to large bursts, and replace all remaining traffic activity by its average rate. The lemma is a minor modification of Lemma 3 in [29].

Lemma VII.1: Let $\mathbf{S}_n = \mathbf{X}_1 + \ldots + \mathbf{X}_n$ be a random walk with i.i.d. step sizes such that $\mathbb{E}\{\mathbf{X}_1\} < 0$ and $\mathbb{E}\{\mathbf{X}_1^p\} < \infty$ for some $p > 1$. Then, for any $\mu < \infty$, there exists a $\kappa^* > 0$ and a function $\phi(\cdot) \in \mathcal{R}_{-\mu}$ such that for all $\kappa \in (0,\kappa^*]$,

$$\mathbb{P}\{\mathbf{S}_n > \hat{x} | \mathbf{X}_i \leq \kappa\hat{x}, i = 1,\ldots,n\} \leq \phi(\hat{x})$$

for all n and $\hat{x}$.

Note that if $\mathbf{X}_i$ can be represented as the difference of two non-negative independent random variables $\mathbf{X}_i^1$ and $\mathbf{X}_i^2$, then the lemma remains valid if the $\mathbf{X}_i$'s are replaced by the $\mathbf{X}_i^1$'s.

We now use the above lemma to show that the workload of flow 2 cannot significantly deviate from the normal drift over intervals of the order $\hat{x}$ when there are no large bursts.

Lemma VII.2: If $B_2(\cdot) \in \mathcal{R}_{-\nu_2}$, then for any $\eta > 0$, $\theta > 0$, there exists a $\kappa^* > 0$ such that for all $\kappa \in (0,\kappa^*]$,

$$\mathbb{P}\{\mathbf{T}_2((\theta - (\phi_2 - \rho_2)\eta)\hat{x}) > \eta\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,\eta\hat{x}] = 0\} =$$

$$\mathrm{o}(\mathbb{P}\{\mathbf{B}_2^r > \hat{x}(\phi_2 - \rho_2)\}) \text{ as } \hat{x} \to \infty.$$

Proof It is not hard to prove that the event

$$\mathbf{T}_2((\theta - (\phi_2 - \rho_2)\eta)\hat{x}) > \eta\hat{x}$$

implies that

$$\sup_{0\leq u\leq\eta\hat{x}} \{A_2(0,u) - (\rho_2 + \theta/2\eta)u\} > \theta\hat{x}/2.$$

Now let $\mathbf{S}_n := \mathbf{X}_1+\ldots+\mathbf{X}_n$ be a random walk with step sizes $\mathbf{X}_i := \mathbf{B}_{2i} - (\rho_2 + \theta/2\eta)\mathbf{U}_{2i}$, with $\mathbf{U}_{2i}$ and $\mathbf{B}_{2i}$ i.i.d. random variables representing the interarrival times and burst sizes of flow 2, respectively. Note that $\mathbf{X}_i$ represents the net increase in the workload in a queue of capacity $\rho_2 + \theta/2\eta$ between two consecutive bursts, and that $\mathbb{E}\{\mathbf{X}_i\} < 0$. Because of the saw-tooth nature of the process $\{A_2(0,u) - (\rho_2 + \theta/2\eta)u\}$, we have

$$\sup_{0\le u\le t}\{A_2(0,u) - (\rho_2+\theta/2\eta)u\} \le \mathbf{B}_{20} + \sup_{1\le n\le N(t)} \mathbf{S}_n.$$

Thus,

$$\begin{aligned}
&\mathbb{P}\{\mathbf{T}_2((\theta-(\phi_2-\rho_2)\eta)\hat{x}) > \eta\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,\eta\hat{x}] = 0\} \\
\le\ & \mathbb{P}\{\mathbf{B}_{20} + \sup_{1\le n\le N(\eta\hat{x})} \mathbf{S}_n \ge \theta\hat{x}/2, \mathcal{N}_{\kappa\hat{x}}[0,\eta\hat{x}] = 0\} \\
\le\ & \mathbb{P}\{\mathbf{B}_{20} + \sup_{1\le n\le N(\eta\hat{x})} \mathbf{S}_n \ge \theta\hat{x}/2 \mid \mathcal{N}_{\kappa\hat{x}}[0,\eta\hat{x}] = 0\} \\
\le\ & \mathbb{P}\{\mathbf{B}_{20} + \sup_{1\le n\le N(\eta\hat{x})} \mathbf{S}_n \ge \theta\hat{x}/2 \mid \mathbf{B}_{2i} \le \kappa\hat{x}, i\ge 0\} \\
\le\ & \mathbb{P}\{\sup_{1\le n\le N(\eta\hat{x})} \mathbf{S}_n \ge (\theta/2-\kappa)\hat{x} \mid \mathbf{B}_{2i} \le \kappa\hat{x}, i\ge 1\} \\
\le\ & \mathbb{P}\{\sup_{1\le n\le (\lambda_2+\epsilon)\eta\hat{x}} \mathbf{S}_n \ge (\theta/2-\kappa)\hat{x} \mid \mathbf{B}_{2i} \le \kappa\hat{x}, i\ge 1\} \\
& + \mathbb{P}\{N(\eta\hat{x}) > (\lambda_2+\epsilon)\eta\hat{x}\} \\
\le\ & \sum_{i=1}^{(\lambda_2+\epsilon)\eta\hat{x}} \mathbb{P}\{\mathbf{S}_n \ge (\theta/2-\kappa)\hat{x} \mid \mathbf{B}_{2i} \le \kappa\hat{x}, i = 1,\ldots,n\} \\
& + \mathbb{P}\{N(\eta\hat{x}) > (\lambda_2+\epsilon)\eta\hat{x}\}.
\end{aligned}$$

The second term decays exponentially fast as $\hat{x} \to \infty$. According to Lemma VII.1, there exists a $\kappa^* > 0$ and a function $\phi(\cdot) \in \mathcal{R}_{-\mu}$, $\mu > \nu_2$, such that for all $\kappa \in (0, \kappa^*]$, each of the probabilities in the first term is upper bounded by $\phi(\hat{x})$. The statement then follows. □

The following two lemmas show that it is relatively unlikely for flow 2 to cause two rare events to happen. Lemma VII.3 states that flow 2 is not likely to generate two large bursts in an interval of order $\hat{x}$. Lemma VII.4 shows that it is not likely for flow 2 to have a workload of at least order $\hat{x}$ at time 0 and to generate at the same time at least one large burst in an interval of order $\hat{x}$. The proofs can be found in [8].

Lemma VII.3: If $B_2(\cdot) \in \mathcal{R}_{-\nu_2}$, then for any $\alpha < 1$, $\kappa > 0$,

$$\mathbb{P}\{\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \ge 2\} = o(\mathbb{P}\{\mathbf{B}_2^r > \hat{x}(\phi_2-\rho_2)\})$$

as $\hat{x} \to \infty$.

Lemma VII.4: If $B_2(\cdot) \in \mathcal{R}_{-\nu_2}$, then for any $0 < \xi < 1-\alpha$, $\zeta > 0$, $\kappa > 0$,

$$\mathbb{P}\{\mathcal{N}_{\kappa\hat{x}}[\xi\hat{x},(1-\alpha)\hat{x}] \ge 1, V_2^c(0) > \zeta\hat{x}\} =$$

$$o(\mathbb{P}\{\mathbf{B}_2^r > \hat{x}(\phi_2-\rho_2)\}) \text{ as } \hat{x}\to\infty.$$

VIII. Backlog period for instantaneous input

In this section we consider the case where flow 2 generates instantaneous traffic bursts of regularly varying size. The next theorem shows that flow 2 then satisfies Assumptions VI.1-VI.3 and that (4) holds.

Theorem VIII.1: If $B_2(\cdot) \in \mathcal{R}_{-\nu_2}$, then for any $c > \rho_2$, $\alpha > 0$, $\gamma > 0$,

$$\begin{aligned}
&\mathbb{P}\{\mathbf{T}_2^c(\gamma\hat{x}) > (1+\alpha)\hat{x}\} \\
&\gtrsim \frac{\rho_2}{c-\rho_2}\mathbb{P}\{\mathbf{B}_2^r > ((\phi_2-\rho_2)(1+\alpha)+\gamma)\hat{x}\},
\end{aligned} \quad (14)$$

$$\mathbb{P}\{\mathbf{T}_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}\} \lesssim$$

$$\frac{\rho_2}{c-\rho_2}\mathbb{P}\{\mathbf{B}_2^r > ((\phi_2-\rho_2)(1-\alpha) - \gamma\frac{c+\phi_2-2\rho_2}{\phi_2-\rho_2})\hat{x}\}, \quad (15)$$

and

$$\mathbb{P}\{\mathbf{T}_2^c > \hat{x}\} \sim \frac{\rho_2}{c-\rho_2}\mathbb{P}\{\mathbf{B}_2^r > \hat{x}(\phi_2-\rho_2)\}. \quad (16)$$

Before giving the formal proof of the above theorem, we first provide an intuitive argument. Consider a queue of capacity ϕ_2 fed by the arrival process of flow 2. In order for the event $\mathbf{T}_2^c > \hat{x}$ to occur, the workload must remain positive throughout the interval $[0,\hat{x}]$, given that the initial workload is $V_2^c(0)$. Note that the normal drift in the workload is $\rho_2 - \phi_2 < 0$. Thus, there is a 'deficit' $(\phi_2-\rho_2)\hat{x}$, which must be compensated for by the initial workload $V_2^c(0)$ plus possibly flow 2 showing above-average activity during the interval $[0,\hat{x}]$.

We claim that the most likely way for the gap to be filled is by a large initial workload only, i.e., $V_2^c(0) > (\phi_2-\rho_2)\hat{x}$. This in turn is most probably due to an extremely large burst of flow 2 somewhere before time 0, which is consistent with the usual situation for heavy-tailed distributions that a large deviation is caused by just a single exceptional event. Using Theorem II.1, we see that the probability of this event is indeed exactly the right-hand side of (16).

Note that it is unlikely for the gap to be filled by flow 2 producing extra traffic during the interval $[0,\hat{x}]$, because this would require a large burst arriving almost immediately after time 0. The probability of this event is negligibly small compared to that of $V_2^c(0) > (\phi_2-\rho_2)\hat{x}$. A combination of both is even less likely, since this would amount to two rare events occurring simultaneously.

The above arguments will be formalized in the proof below. We first prove that the event $V_2^c(0) > (\phi_2-\rho_2)\hat{x}$ indeed implies that $\mathbf{T}_2^c > \hat{x}$ for large $\hat{x}$, thus obtaining a lower bound for the probability of the latter event. Next we show that for large $\hat{x}$ the event $V_2^c(0) > (\phi_2-\rho_2)\hat{x}$ is also necessary for $\mathbf{T}_2^c > \hat{x}$ to occur, by proving that the probability of all other possible scenarios is negligibly small.

Proof of Theorem VIII.1 We start with the proof of (14). It is not difficult to show that for any $\alpha > 0$, $\gamma > 0$, $\delta > 0$, $\theta > 0$, the event $\mathbf{T}_2^c(\gamma\hat{x}) > (1+\alpha)\hat{x}$ is implied by the events $V_2^c(0) > ((\phi_2-\rho_2+\delta)(1+\alpha)+\gamma+\theta)\hat{x}$ and $\sup_{0\le u\le(1+\alpha)\hat{x}}\{(\rho_2-\delta)u - A_2(0,u)\} \le \theta\hat{x}$.

Hence, using independence of $V_2^c(0)$ and $A_2(0,u)$,

$$\mathbb{P}\{\mathbf{T}_2^c(\gamma\hat{x}) > (1+\alpha)\hat{x}\} \ge \mathbb{P}\{\sup_{0\le u\le(1+\alpha)\hat{x}}\{(\rho_2-\delta)u - A_2(0,u)\} \le \theta\hat{x}\}$$

$$\mathbb{P}\{V_2^c(0) > ((\phi_2-\rho_2+\delta)(1+\alpha)+\gamma+\theta)\hat{x}\}.$$

Using Theorem II.1,

$$\mathbb{P}\{V_2^c(0) > ((\phi_2-\rho_2+\delta)(1+\alpha)+\gamma+\theta)\hat{x}\} \sim$$

$$\frac{\rho_2}{c-\rho_2}\mathbb{P}\{\mathbf{B}_2^r > ((\phi_2 - \rho_2 + \delta)(1+\alpha) + \gamma + \theta)\hat{x})\}.$$

Also, for all $\alpha > 0$, $\delta > 0$, $\theta > 0$,

$$\mathbb{P}\{\sup_{0\le u\le(1+\alpha)\hat{x}} \{(\rho_2 - \delta)u - A_2(0,u)\} \le \theta\hat{x}\} \ge$$

$$\mathbb{P}\{\sup_{u\ge 0}\{(\rho_2 - \delta)u - A_2(0,u)\} \le \theta\hat{x}\} \to 1,$$

as $\hat{x} \to \infty$, since $\mathbb{E}\{A_2(0,u)\} = \rho_2 u$.

Thus, for all $\alpha, \gamma, \delta, \theta > 0$, $\mathbb{P}\{\mathbf{T}_2^c(\gamma\hat{x}) > (1+\alpha)\hat{x}\} \gtrsim$

$$\frac{\rho_2}{c-\rho_2}\mathbb{P}\{\mathbf{B}_2^r > ((\phi_2 - \rho_2 + \delta)(1+\alpha) + \gamma + \theta)\hat{x}\}.$$

Letting $\delta, \theta \downarrow 0$ and using $B_2^r(\cdot) \in \mathcal{R}_{-\nu_2}$, (14) follows.

We now turn to the proof of (15). By partitioning, we obtain for any $\alpha, \gamma, \zeta, \theta, \kappa > 0$, $w \ge 0$,

$$\begin{aligned}
& \mathbb{P}\{\mathbf{T}_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}\} \\
\le\ & \mathbb{P}\{V_2^c(w) > ((\phi_2 - \rho_2)(1-\alpha) - \gamma - \theta)\hat{x} - cw\} \\
+\ & \mathbb{P}\{\mathbf{T}_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[w,(1-\alpha)\hat{x}] = 0, \\
& \quad V_2^c(w) \le ((\phi_2 - \rho_2)(1-\alpha) - \gamma - \theta)\hat{x} - cw\} \\
+\ & \mathbb{P}\{\mathbf{T}_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,w] = 0, V_2^c(0) \le \zeta\hat{x}\} \\
+\ & \mathbb{P}\{\mathcal{N}_{\kappa\hat{x}}[w,(1-\alpha)\hat{x}] \le 1, V_2^c(0) > \zeta\hat{x}\} \\
+\ & \mathbb{P}\{\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha)\hat{x}] \ge 2\} \\
=\ & (\mathrm{A}) + (\mathrm{B}) + (\mathrm{C}) + (\mathrm{D}) + (\mathrm{E}).
\end{aligned}$$

Take $w = \xi\hat{x}$, with $\xi := \frac{\gamma+\zeta+\theta}{\phi_2-\rho_2} < 1-\alpha$.

Now consider term (A). Using Theorem II.1, (A) $\sim$

$$\frac{\rho_2}{c-\rho_2}\mathbb{P}\{\mathbf{B}_2^r > ((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta-\frac{(c-\rho_2)(\gamma+\zeta+\theta)}{\phi_2-\rho_2})\hat{x}\}.$$

Next, it may be shown that term (B) is bounded above by

$$\mathbb{P}\{\mathbf{T}_2(\theta - (\phi_2 - \rho_2)(1-\alpha-\xi))\hat{x}) > (1-\alpha-\xi)\hat{x},$$

$$\mathcal{N}_{\kappa\hat{x}}[0,(1-\alpha-\xi)\hat{x}] = 0\}.$$

Finally, term (C) is bounded above by

$$\mathbb{P}\{\mathbf{T}_2((\theta - (\phi_2 - \rho_2)\xi)\hat{x}) > \xi\hat{x}, \mathcal{N}_{\kappa\hat{x}}[0,\xi\hat{x}] = 0\}.$$

Thus, taking $\eta = \xi$ and $\eta = 1 - \alpha - \xi$ in Lemma VII.2, and using Lemma VII.3, we obtain

$$\mathbb{P}\{\mathbf{T}_2^c(-\gamma\hat{x}) > (1-\alpha)\hat{x}\} \lesssim$$

$$\frac{\rho_2}{c-\rho_2}\mathbb{P}\{\mathbf{B}_2^r > ((\phi_2-\rho_2)(1-\alpha)-\gamma-\theta-\frac{(c-\rho_2)(\gamma+\zeta+\theta)}{\phi_2-\rho_2})\hat{x}\}$$

Letting $\zeta, \theta \downarrow 0$ and using $B_2^r(\cdot) \in \mathcal{R}_{-\nu_2}$, (15) follows.

Finally, note that (16) follows from (14) and (15) by letting $\alpha \downarrow 0$, $\gamma \downarrow 0$, and using the fact that $B_2^r(\cdot) \in \mathcal{R}_{-\nu_2}$. □

IX. Conclusion

We analyzed a GPS queue with two flows, one having light-tailed characteristics, the other one exhibiting heavy-tailed properties. We showed that the workload distribution of the light-tailed flow is asymptotically equivalent to that when served in isolation at its minimum guaranteed rate, multiplied with a certain pre-factor. The pre-factor may be interpreted as the probability that the heavy-tailed flow is backlogged long enough for the light-tailed flow to reach overflow. We did not consider the case where the traffic intensity of the heavy-tailed flow exceeds its minimum guaranteed rate. In this case, the pre-factor – representing again the probability that the heavy-tailed flow is continuously backlogged during the period to overflow of the light-tailed flow – is likely to be some constant. Determining the exact value of the constant seems however a rather challenging task.

In the present paper we have focused on a scenario with two flows. Observe however that the light-tailed flow may be thought of as an aggregate flow, given that the class of Markov-modulated fluid input is closed under superposition of independent processes. In case of instantaneous input, the heavy-tailed flow too may actually represent an aggregate flow, since the superposition of independent Poisson streams with regularly varying bursts produces again a Poisson stream with regularly varying bursts. Unfortunately, the class of On-Off sources is clearly not closed under superposition. In fact, the superposition exhibits a fundamentally more complex structure than a single On-Off-source, which drastically complicates the analysis of the queueing behavior.

References

[1] Anantharam, V. (1988). How large delays build up in a GI/G/1 queue. *Queueing Systems* **5**, 345–368.

[2] Arvidsson, A. Karlsson, P. (1999). On traffic models for TCP/IP. In: *Teletraffic Engineering in a Competitive World, Proc. ITC-16*, Edinburgh, UK, eds. P. Key, D. Smith (North-Holland, Amsterdam), 457–466.

[3] Beran, J., Sherman, R., Taqqu, M.S., Willinger, W. (1995). Long-range dependence in variable-bit-rate video traffic. IEEE *Trans. Commun.* **43**, 1566–1579.

[4] Bingham, N.H., Goldie, C.M., Teugels, J.L. (1987). *Regular Variation* (Cambridge University Press, Cambridge UK).

[5] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., Weiss, W. (1998). An architecture for differentiated services. IETF RFC 2475.

[6] Borst, S.C., Boxma, O.J., Jelenković, P.R. (2000). Reduced-load equivalence and induced burstiness in GPS queues with long-tailed traffic flows. CWI Report PNA-R0016, http://www.cwi.nl/static/publications/reports/abs/PNA-R0016.html, *Queueing Systems*, to appear.

[7] Borst, S.C., Boxma, O.J., Van Uitert, M.J.G. (2001). Two coupled queues with heterogeneous traffic. In: *Teletraffic Engineering in the Internet Era. Proc. ITC-17*, Salvador da Bahia, Brazil (Eds. J. Moreira de Souza, N.L.S. Fonseca, E.A. de Souza e Silva), 1003–1014.

[8] Borst, S.C., Mandjes, M., Van Uitert, M.J.G. (2001). Generalized Processor Sharing queues with heterogeneous traffic classes. CWI Report PNA-R0106, http://www.cwi.nl/static/publications/reports/abs/PNA-R0106.html.

[9] Borst, S.C., Zwart, A.P. (2000). A reduced-peak equivalence for queues with a mixture of light-tailed and heavy-tailed input flows. Technical Report SPOR 2000-04, Eindhoven University of Technology. *J. Appl. Prob.*, to appear.

[10] Boxma, O.J., Deng, Q., Zwart, A.P. (1999). Waiting-time asymptotics for the M/G/2 queue with heterogeneous servers. *Queueing Systems*, to appear.

[11] Cao, J., Ramanan, K. (2001). A Poisson limit for the unfinished work of superposed point processes. Technical Memorandum. Bell Laboratories, Lucent Technologies, Murray Hill, NJ.

[12] Crovella, M., Bestavros, A. (1996). Self-similarity in World Wide Web Traffic: evidence and possible causes. In: *Proc. ACM Sigmetrics '96*, 160–169.

[13] Elwalid, A.I., Mitra, D. (1991). Analysis and design of rate-based congestion control of high speed networks, I: stochastic fluid models, access regulation. *Queueing Systems* **9**, 29–64.
[14] Elwalid, A.I., Mitra, D. (1993). Effective bandwidth of general Markovian traffic sources and admission control of high speed networks. IEEE/ACM *Trans. Netw.* **1**, 329–343.
[15] Glynn, P.W., Whitt, W. (1994). Logarithmic asymptotics for steady-state tail probabilities in a single-server queue. *J. Appl. Prob.* **31A**, 131–156.
[16] Grossglauser, M., Bolot, J.-C. (1999). On the relevance of long-range dependence in network traffic. IEEE/ACM *Trans. Netw.* **7**, 629–640.
[17] Jelenković, P.R., Lazar, A.A. (1999). Asymptotic results for multiplexing subexponential on-off processes. *Adv. Appl. Prob.* **31**, 394–421.
[18] Kesidis, G., Walrand, J., Chang, C.-S. (1993). Effective bandwidths for multi-class Markov fluids and other ATM sources. IEEE/ACM *Trans. Netw.* **1**, 424–428.
[19] Kotopoulos, C., Likhanov, N., Mazumdar, R.R. (2001). Asymptotic analysis of the GPS system fed by heterogeneous long-tailed sources. In: *Proc. Infocom 2001 Conference*, Anchorage, Alaska, USA, 299–308.
[20] Kosten, L. (1986). Liquid models for a type of information buffer problem. *Delft Progress Report* **11**, 71–86.
[21] Leland, W.E., Taqqu, M.S., Willinger, W., Wilson, D.V. (1994). On the self-similar nature of Ethernet traffic (extended version). IEEE/ACM *Trans. Netw.* **2**, 1–15.
[22] Mandjes, M., Kim, J.-H. (2001). Large deviations for small buffers: an insensitivity result. *Queueing Systems* **37**, 349–362.
[23] Mandjes, M., Ridder, A. (1995). Finding the conjugate of Markov fluid processes. *Prob. Eng. Inf. Sci.* **9**, 297–315.
[24] Massoulié, L. (1999). Large deviations estimates for polling and weighted fair queueing service systems. *Adv. Perf. Anal.* **2**, 103–128.
[25] Pakes, A.G. (1975). On the tails of waiting-time distributions. *J. Appl. Prob.* **12**, 555–564.
[26] Parekh, A., Gallager, R. (1993). A generalized processor sharing approach to flow control in integrated services networks: The single node case. IEEE/ACM *Trans. Netw.* **1**, 344–357.
[27] Park, K., Willinger, W. (eds.) (2000). *Self-Similar Network Traffic and Performance Evaluation* (Wiley, New York).
[28] Paxson, A., Floyd, S. (1995). Wide area traffic: the failure of Poisson modeling. IEEE/ACM *Trans. Netw.* **3**, 226–244.
[29] Resnick, S., Samorodnitsky, G. (1999). Activity periods of an infinite server queue and performance of certain heavy tailed fluid queues. *Queueing Systems* **33**, 43–71.
[30] Ryu, B., Elwalid, A. (1996). The importance of long-range dependence of VBR video traffic in ATM traffic engineering: myths and realities. *Comp. Commun. Rev.* **13**, 1017–1027.
[31] Sigman, K. (ed.) (1999). *Queueing Systems* **33**. Special Issue on Queues with Heavy-Tailed Distributions.
[32] Zhang, Z.-L. (1998). Large deviations and the generalized processor sharing scheduling for a multiple-queue system. *Queueing Systems* **28**, 349–376.