Network Multiplexer with Generalized Processor Sharing and Heavy-tailed On-Off Flows

Predrag Jelenković and Petar Momčilović

Department of Electrical Engineering, Columbia University,
500 West 120th Street, New York, NY 10027, United States

We consider a set of fluid On-Off flows that share a common server of capacity c and a finite buffer B. The server capacity is divided using Generalized Processor Sharing scheduling discipline. Each flow has a minimum service rate guarantee that exceeds its long term average demand ρ_i. The buffer sharing is unrestricted as long as there is available space; if the buffer is full, the necessary amount of fluid from the most demanding flows is discarded. When the On periods are heavy-tailed, we show that the loss rate of a particular flow i is asymptotically equal to the loss rate in a reduced system with capacity $c - \Sigma_{j \neq i} \rho_j$ and buffer B, where this flow is served in isolation. In particular, each flow perceives to have the whole buffer B to itself. This insight provides a new guideline for efficiently engineering differentiated quality of service in integrated multimedia networks.

1. INTRODUCTION

Modern communication networks are engineered to carry a diverse spectrum of multimedia services, ranging from real-time traffic, such as voice and video, to various data and Web related applications. These services have different Quality of Service (QoS) requirements, e.g., real-time services have stringent delay requirements, but can tolerate relatively high losses. On the other hand, data related services typically could tolerate larger delays, but need minimal or no losses. Providing the QoS differentiation in integrated multimedia networks is usually achieved through priority scheduling mechanisms. The most popular scheduling schemes, e.g. Weighted Fair Queueing, are based on Generalized Processor Sharing (GPS) algorithm. These algorithms offer the flexibility for providing high degree of service differentiation, extracting statistical multiplexing gains as well as protecting individual flows from the ones with high service demands.

Rigorous investigation of stochastic systems with GPS dates back to [9]; see also [9] for some earlier references. This work was centered around the problem of time-shared computer systems. Renewed interest in GPS stems from the previously mentioned networking application. Recent results for traffic models with exponential characteristics can be found in [1,18,20,21]. However, comprehensive statistical measurements in currently

---

This work is supported by the National Science Foundation Presidential Early Career Accomplishment Award No. 9875156.

To appear in: Proceedings of the Seventeenth International Teletraffic Congress (ITC-17), Salvador da Bahia, Brazil, September 2001

deployed networks have found repeated evidence of the presence of high degree of statistical dependency and variability, often referred to as self-similarity, in network traffic streams, e.g., see [16,10,15,7,19]. The most common approaches to modelling these intricate statistical phenomena is through long-range dependent Gaussian processes, such as Fractional Browninan Motion, and heavy-tailed semi-Markov type processes, e.g., On-Off processes. In this paper we focus on the latter.

Motivated by these empirical findings, a series of research studies that develop new analytical techniques for evaluating stochastic heavy-tailed queueing systems have appeared. In particular, very recent investigations of the behavior of GPS in the presence of heavy-tailed arrival streams can be found in [2–6]; the reader may consult the same papers for additional references on processor sharing and heavy tails. These papers consider a system with finite number of heavy-tailed sessions each of which is queued into an infinite buffer queue; the content of the queues is served by a single server that is scheduled using GPS. When the GPS weights are appropriately chosen, these studies show that each session experiences the same queueing behavior as if it were served in isolation with an appropriate constant capacity. However, if the weights are not properly engineered, the flows may experience induced burstiness [3,6]; these results stress the importance of properly selecting the GPS weights.

In this paper we investigate a queueing system in which the heavy-tailed On-Off sessions share a single finite buffer B. The buffer sharing is unrestricted as long as there is available space; if the buffer is full, the necessary amount of fluid from the most demanding flows is discarded. The flows share a common server according to GPS scheduling discipline. A formal description of our model and some preliminary results are presented in Sections 2 and 3, respectively. When each flow receives the minimum service rate guarantee that exceeds its long-term average demand ρ_i, our main result, stated in Theorem 4.1 of Section 4, shows that the loss rate of a particular flow i is asymptotically equal to the loss rate in a reduced system with capacity $c - \sum_{j \neq i} \rho_j$ and buffer B, where this flow is served in isolation. In other words, each flow perceives to have the entire buffer to itself. This finding complements the result derived in [2]. The new qualitative insight may prove useful in deciding on whether to engineer buffers at the periphery or in the core of the network. Numerical validation of our main result is exemplified in Section 5. It is worth observing that the derived asymptotic approximation yields accurate results even for relatively large probabilities in the range of $10^{-2} - 10^{-3}$. Concluding remarks are presented in Section 6.

2. MODEL DESCRIPTION

Consider N independent On-Off flows $A_i(t)$, $1 \leq i \leq N$. An On-Off flow consists of an alternating sequence of independent activity and silence periods. Activity and silence periods of the ith flow are equal in distribution to τ_i^{on} and τ_i^{off}, respectively. Assume that τ_i^{on} and τ_i^{off} are almost surely (a.s.) positive and that their expectations $\mathbb{E}\tau_i^{on}$ and $\mathbb{E}\tau_i^{off}$ are finite. During an activity period flow i generates fluid at constant rate $A_i(t) = r_i$ and for a period of silence $A_i(t) = 0$. Suppose that $A_i(t)$ is right-continuous and stationary; for a precise construction see [11]. Then, the stationary probability that flow i is active and its long term average rate are equal to $p_{on,i} = \mathbb{P}[A_i(t) = r_i] = \mathbb{E}\tau_i^{on}/(\mathbb{E}\tau_i^{on} + \mathbb{E}\tau_i^{off})$ and

$\rho_i = \mathbb{E}A_i(t) = p_{on,i}r_i$, respectively.

These On-Off flows share a common server of capacity c and a buffer space of size B. Let $W_i(t)$ be the unfinished work of flow i at time t, $\sum W_i(t) \leq B$. The server capacity is split among flows according to GPS scheme. Each flow i is assigned a weight $\phi_i > 0$ such that $\sum_{i=1}^N \phi_i = 1$. Weight ϕ_i represents the guaranteed share of the server capacity for flow i. Available excess capacity is redistributed among flows according to the GPS weights ϕ_i. Service rates $c_i(t)$ for each flow i at time t can be computed by using a recursive algorithm described in [8, pp. 4-5]. Let $E(t)$ be the set of flows with $W_i(t) = 0$, which are receiving service at rate $c_i(t) = A_i(t)$, i.e. $E(t) \triangleq \{i : W_i(t)=0,\ c_i(t)=A_i(t)\}$. We would like to observe that $E(t) = \{i : W_i(t)=0\}$ almost everywhere (a.e.) Lebesgue. Therefore, using this property and the characteristics of GPS for any flow $i \notin E(t)$, rate $c_i(t)$ satisfies

$$c_i(t) = \frac{\phi_i \left(c - \sum_{j:W_j(t)=0} A_j(t)\right)}{\sum_{j:W_j(t)>0} \phi_j} \geq c\phi_i \qquad \text{a.e.} \tag{2.1}$$

Buffer sharing is unrestricted as long as there is available space, i.e., the workloads evolve as if they were in the infinite buffer system. When the buffer fills up, the flows with the maximum amount of fluid $W_i(t)$ in the buffer are subject to penalty. They will experience the minimum necessary loss of fluid such that the flows with smaller workloads can be accommodated. Possible extensions to more general buffer sharing policies will be briefly discussed at the end of Section 4.

More formally, following the approach from [8], the evolution of W_i-s can be described with a set of differential equations. In order to account for the finiteness of the buffer space we define

$$D(t) = \left\{ i \ : \ W_i(t) = \max_{1\leq j\leq N} W_j(t),\ \sum_{j=1}^N W_j(t) = B \right\}$$

with $|D(t)|$ denoting the cardinality of $D(t)$; note that $D(t)$ is nonempty only if the buffer is full. The elements of $D(t)$ are flows that could potentially experience losses at time t. Let $M(t)$ be the largest subset of $D(t)$ such that for all $i \in M(t)$ the following inequality holds

$$A_i(t) - c_i(t) > -|M(t)|^{-1} \sum_{j\notin M(t)} (A_j(t) - c_j(t)).$$

The workloads of the flows in $M(t)$ will be reduced according to

$$\dot{W}_i(t) = -|M(t)|^{-1} \sum_{j\notin M(t)} (A_j(t) - c_j(t))\,;$$

this reduction is necessary to ensure $\sum W_i = B$. Next, it is not very difficult to see that $M(t) = D(t)$ a.e. (Lebesgue). Thus, the evolution of the individual workloads is a.e. described by the following set of differential equations

$$\dot{W}_i(t) = \begin{cases} 0 & \text{if } i \in E(t), \\ -|D(t)|^{-1} \sum_{j\notin D(t)} (A_j(t) - c_j(t)) & \text{if } i \in D(t), \\ A_i(t) - c_i(t) & \text{otherwise.} \end{cases} \tag{2.2}$$

The functions $W_i(t)$ are absolutely continuous and therefore sets of Lebesgue measure zero can be ignored for purpose of their characterization [8]. Hence, the system of equations (2.2) completely defines the behavior of the system.

We assume that this system of equations has a unique stationary solution and, unless otherwise specified, $W_i(t)$ will be used to denote this solution. Let A_i, c_i , W_i be random variables equal in distribution to $A_i(t)$, $c_i(t)$, $W_i(t)$ in stationarity, respectively.

The main objective of this paper is the asymptotic computation of the long-term average loss rate for a given flow as the buffer size grows to infinity. We say that flow i is overflowing at time t if $A_i(t)-c_i(t) > \dot{W}_i(t)$. The instantaneous loss rate process for flow i is defined as

$$\gamma_i^B(t) \triangleq A_i(t) - c_i(t) - \dot{W}_i(t), \tag{2.3}$$

and its expected loss rate is equal to $\Gamma_i^B = \mathbb{E}\gamma_i^B(t)$. Note that the buffer management policy implies

$$\left\{A_i(t) - c_i(t) > \dot{W}_i(t)\right\} \subseteq \left\{W_i(t) \geq B/N,\ \sum_{i=1}^{N} W_i(t) = B\right\}. \tag{2.4}$$

This inclusion is a simple consequence of the fact that if the buffer is full and $W_i(t) < B/N$, then there exists a flow j with higher workload than that of flow i, which contradicts the fact that flow i is experiencing losses. It is worth mentioning that the described queueing system is work- and buffer-conserving. Work-conservation follows from the properties of GPS; buffer-conservation means that no flow will experience loss of fluid unless the buffer is full.

3. PRELIMINARY RESULTS

This section contains preliminary sample path bounds and some of the existing results from the literature on fluid queues with heavy-tailed On-Off flows that will be used in deriving our main results.

3.1. Sample path bounds

We start with an introduction of a finite buffer queuing process that will be used to bound the workload processes in the GPS system with buffer sharing that we have described in Section 2. Consider a fluid queue with constant capacity c, finite buffer B and arrival process $A(t)$. Informally, at time t, fluid arrives to this system at rate $A(t)$ and is leaving the system at rate c. When the queue level reaches the buffer limit B fluid arriving in excess of the draining rate c is lost. We use $Q_A^{B,c}(t) \in [0, B]$ to denote the queue content at time t.

In this paper we only consider arrival processes $A(t)$ that are piece-wise constant and right continuous with a.s. increasing jump times $\{T_0 = 0 < T_1 < T_2 < \cdots\}$. In this case, for any initial value $Q_A^{B,c}(0)$ and $t \in (T_n, T_{n+1}], n \geq 0$, the evolution of $Q_A^{B,c}(t)$ is given by

$$Q_A^{B,c}(t) = \min\left((Q_A^{B,c}(T_n) + (t - T_n)(A(T_n) - c))^+, B\right), \tag{3.1}$$

where $(x)^+ = \max(0, x)$. When $A(t)$, i.e. $\{(T_{n+1} - T_n), A(T_n)\}$, is stationary and ergodic, and $\mathbb{E}A(t) < c$, by using Loynes' construction [17], one can show that recursion (3.1) has a unique stationary and ergodic solution. Furthermore, for all initial conditions $Q_A^{B,c}(0)$, the distribution of $Q_A^{B,c}(t)$ converges to that stationary solution as $t \to \infty$. Let $Q_A^{B,c}$ and A be random variables that are equal in distribution to $Q_A^{B,c}(t)$ and $A(t)$, respectively. The loss rate for the described system, $\Lambda_A^{B,c}$, is defined as

$$\Lambda_A^{B,c} \triangleq \mathbb{E}[(A-c)\mathbf{1}\{Q_A^{B,c}=B\}].$$

Proposition 3.1 *If $W_i(t) \le Q_{A_i}^{B,\phi_i c}(t)$ for $t=0$, then the inequality holds for all $t \ge 0$.*

Proof: Follows from inequality (2.1) and the fact that the portion of the buffer available to $W_i(t)$ is not greater than B. $\diamond$

In order to simplify the notation we set $A(t)=\sum_{i=1}^N A_i(t)$, $r=\sum_{i=1}^N r_i$, $A_{-i}(t)=\sum_{j\neq i} A_j(t)$ and $\rho_{-i}=\sum_{j\neq i}\rho_j$. The second sample path bound formalizes an intuitively expected notion that multiplexing reduces the aggregate queueing workload. The proof of the following proposition can be found in [11].

Proposition 3.2 *If $Q_A^{B,\sum_{n=1}^N c_n}(t) \le \sum_{n=1}^N Q_{A_n}^{B,c_n}(t)$ for $t=0$, then the inequality holds for all $t \ge 0$.*

Here, we state another sample path bound from [11]. This bound will be used to limit the amount of free buffer space in a finite buffer queue. Let quantity $Q_c^{\infty,A}(t)$ represent the workload in an infinite buffer fluid queue with constant arrival rate c and variable service rate $A(t)$.

Proposition 3.3 *If $B-Q_A^{B,c}(t) \le Q_c^{\infty,A}(t)$ for $t=0$, then the inequality holds for all $t \ge 0$.*

3.2. Queueing results

This subsection consists of the known results on fluid queues with heavy-tailed On-Off flows that will be used in proving our main result. In the analysis of renewal processes residual (or excess) random variables and distribution functions play an important role. For a non-negative random variable X with distribution F and finite mean $\mathbb{E}X$, the residual distribution F_r is defined by $F_r(x)=(\mathbb{E}X)^{-1}\int_0^x(1-F(u))du$, $x \ge 0$. A random variable X_r with d.f. F_r is called the residual variable of X. Throughout the paper, for any two real functions $f(x)$ and $g(x)$, we use the frequently used notation $f(x) \sim g(x)$ as $x \to \infty$ to denote $\lim_{x\to\infty} f(x)/g(x)=1$ or equivalently $f(x)=g(x)(1+o(1))$ as $x \to \infty$.

The theorem below summarizes the results from [14,13] in the case when the arrival process is a single On-Off flow. Common heuristic explanation of this theorem comes from identifying the most likely buffer overflow scenario, which, in this system, is due to an isolated very long On period. Classes of heavy-tailed distributions $\mathcal{IR}$, $\mathcal{S}$ ($\mathcal{IR} \subset \mathcal{S}$) are defined in Appendix.

Theorem 3.1 *Let $r_i > c > \rho_i$ and $\tau_{i,r}^{on} \in \mathcal{S}$. Then, as $B \to \infty$*

$$\Lambda_{A_i}^{B,c}(r_i-c)^{-1} = \mathbb{P}[Q_{A_i}^{B,c}=B] \sim K\mathbb{P}[Q_{A_i}^{\infty,c}>B] \sim p_{on,i}\mathbb{P}\left[\tau_{i,r}^{on} > \frac{B}{r_i-c}\right],$$

where $K < \infty$. *Furthermore, if* $\tau_{i,r}^{on} \in \mathcal{IR}$ *then*

$$\lim_{\epsilon \uparrow 1} \overline{\lim_{B \to \infty}} \frac{\mathbb{P}[Q_{A_i}^{B,c} \geq \epsilon B]}{\mathbb{P}[Q_{A_i}^{B,c} = B]} = 1.$$

The next theorem provides the asymptotic characterization of the aggregate loss rate in a finite buffer queue with work-conserving scheduling discipline and, therefore, it applies to the total workload $\sum W_i$ in our GPS model. Here, the most likely buffer overflow arises when several very long On periods overlap. Typically, there is more than one way to overflow, which is accounted for with a summation in the theorem. For a detailed proof and experimental investigation of its accuracy see [11] and [12], respectively.

Theorem 3.2 *Let* $r_{\mathcal{Q}} \triangleq \sum_{j \in \mathcal{Q}} r_j + \sum_{j \notin \mathcal{Q}} \rho_j$ *for all* $\mathcal{Q} \subseteq \{1, \ldots, N\}$ *and*

$$\mathcal{O} \triangleq \{\mathcal{Q} : \; r_{\mathcal{Q}} - (r_j - \rho_j) < c \leq r_{\mathcal{Q}}, \forall j \in \mathcal{Q}\}.$$

If $\tau_{i,r}^{on} \in \mathcal{IR}$ *for* $1 \leq i \leq N$ *and* $r_{\mathcal{Q}} \neq c$ *for all* $\mathcal{Q} \in \mathcal{O}$, *then as* $B \to \infty$

$$\Lambda_A^{B,c} \sim \sum_{\mathcal{Q} \in \mathcal{O}} (r_{\mathcal{Q}} - c) \prod_{j \in \mathcal{Q}} p_{on,j} \mathbb{P}\left[\tau_{j,r}^{on} > \frac{B}{r_{\mathcal{Q}} - c}\right].$$

4. MAIN RESULT

Finally, we are ready to state our main result.

Theorem 4.1 *Let* $\phi_j c > \rho_j$ *for all flows. If* $\tau_i^{on} \in \mathcal{IR}$ *and* $r_i > c - \rho_{-i}$ *for some flow* i, *then as* $B \to \infty$

$$\Gamma_i^B \sim \Lambda_{A_i}^{B,c-\rho_{-i}}.$$

Proof: *Upper bound.* Assume that all queueing processes below are equal to zero at time $t = 0$. Based on definition (2.3) of the instantaneous loss rate and its non-negativity, the following expression holds

$$\begin{aligned} \gamma_i^B(t) &= (A_i(t) - c_i(t) - \dot{W}_i(t)) \mathbf{1}\{A_i(t) - c_i(t) > \dot{W}_i(t)\} \\ &\leq (r_i + A_{-i}(t) - c) \mathbf{1}\{A_i(t) - c_i(t) > \dot{W}_i(t)\}, \end{aligned} \tag{4.1}$$

where the last inequality follows from the fact that the buffer is full on event $\{A_i(t) - c_i(t) > \dot{W}_i(t)\}$ and the instantaneous loss rate of a single flow is upper bounded by the total loss rate in the system. Inequality (4.1), inclusion (2.4), Propositions 3.1, and the work-conserving property of the GPS scheduling scheme yield for all $t \geq 0$

$$\begin{aligned} \gamma_i^B(t) &\leq (r_i + A_{-i}(t) - c) \mathbf{1}\left\{W_i(t) \geq B/N, \; \sum_{i=1}^N W_i(t) \geq B\right\} \\ &\leq (r_i + A_{-i}(t) - c) \mathbf{1}\left\{Q_{A_i}^{B,\phi_i c}(t) \geq B/N, \; Q_A^{B,c}(t) \geq B\right\}. \end{aligned}$$

Next, the preceding inequality and Proposition 3.2 result in

$$\gamma_i^B(t) \le (r_i + A_{-i}(t) - c)\mathbf{1}\left\{Q_{A_i}^{B,c-\rho_{-i}-\delta}(t) \ge \epsilon B\right\}$$
$$+ (r - c)\mathbf{1}\left\{Q_{A_i}^{\infty,\phi_i c}(t) \ge B/N\right\}\mathbf{1}\left\{Q_{A_{-i}}^{\infty,\rho_{-i}+\delta}(t) \ge (1-\epsilon)B\right\}$$

where we select $\epsilon \in (0,1)$ and $\delta \in (0, c-\rho_i-\rho_{-i})$. Note that for any choice of ϵ and δ in the given intervals all queueing processes in the last inequality are stable and converge in distribution to proper random variables as $t \to \infty$. Now, by independence of arrival processes, the last inequality renders

$$\Gamma_i^B \le (r_i+\rho_{-i}-c)\mathbb{P}[Q_{A_i}^{B,c-\rho_{-i}-\delta} \ge \epsilon B] + (r-c)\mathbb{P}[Q_{A_i}^{\infty,\phi_i c} \ge B/N]\mathbb{P}[Q_{A_{-i}}^{\infty,\rho_{-i}+\delta} \ge (1-\epsilon)B].$$

Then, Theorem 3.1 and Lemma A.1 imply

$$\overline{\lim_{B\to\infty}} \frac{\Gamma_i^B}{\Lambda_{A_i}^{B,c-\rho_{-i}}} \le \overline{\lim_{B\to\infty}} \frac{\mathbb{P}\left[Q_{A_i}^{B,c-\rho_{-i}-\delta} \ge \epsilon B\right]}{\mathbb{P}\left[Q_{A_i}^{B,c-\rho_{-i}} = B\right]}.$$

To complete the proof of the upper bound pass $\epsilon \uparrow 1$, $\delta \downarrow 0$ and use Theorem 3.1.

Lower bound. Assume that all processes are in their stationary regimes unless otherwise specified. Let $G_i(0,T)$ be the amount of fluid lost from flow i in time interval $(0,T)$. Recall that the loss rate is defined as $\mathbb{E}\gamma_i^B(t)$, where $\gamma_i^B(t)$ is the instantaneous loss rate process of flow i. Clearly, for all $0 < T < \infty$

$$\Gamma_i^B = \mathbb{E}\gamma_i^B(0) = T^{-1}\mathbb{E}\int_0^T \gamma_i^B(t)dt = T^{-1}\mathbb{E}G_i(0,T). \tag{4.2}$$

Next, define quantity $t_i \triangleq (1+\epsilon)B/(r_i+\rho_{-i}-c)$ and two families of events

$$\alpha_j^i \triangleq \{A_j(0) = r_j,\ \inf\{t > 0 : A_j(-t)=0\} > t_i\},$$

$$\xi_i \triangleq \bigcap_{j\neq i}\left\{\alpha_j^{i,c},\ \frac{1}{t_i}\int_{-t_i}^0 A_j(t)dt > \rho_j - \frac{\epsilon(r_i+\rho_{-i}-c)}{(1+\epsilon)N},\ W_j(0) \le \frac{B}{2N}\right\},$$

where $(\cdot)^c$ represents the set complement operation. On event $\{A_j(0) = r_j\}$, the stationarity of arrival process $A_j(t)$ implies $\inf\{t > 0 : A_j(-t) = 0\} \overset{d}{=} \tau_{i,r}^{on}$. Then, from identity (4.2), the loss rate for flow i can be lower bounded by

$$\Gamma_i^B \ge T^{-1}\mathbb{E}\left[G_i(0,T)\mathbf{1}\left\{\alpha_i^i,\ \xi_i,\ \sum_{k=1}^N W_k(0) \ge B-K\right\}\right]. \tag{4.3}$$

Now, by using Proposition 3.3 and recalling that the considered system is work- and buffer-conserving we obtain

$$\left\{\alpha_i^i,\ \xi_i,\ K \ge B - \sum_{k=1}^N W_k(0)\right\} = \left\{\alpha_i^i,\ \xi_i,\ K \ge B - Q_A^{B,c}(0)\right\}$$
$$\supseteq \left\{\alpha_i^i,\ \xi_i,\ K \ge \overline{Q}_{c-r_i}^{\infty,A_{-i}}(0)\right\}, \tag{4.4}$$

where process $\overline{Q}_{c-r_i}^{\infty,A_{-i}}(t)$ denotes the workload of an infinite buffer queue with constant arrival rate $c-r_i$, service rate A_{-i} and the initial condition $\overline{Q}_{c-r_i}^{\infty,A_{-i}}(-t_i)=B$. Using the standard queueing reflection mapping argument quantity $\overline{Q}_{c-r_i}^{\infty,A_{-i}}(0)$ can be represented as

$$\overline{Q}_{c-r_i}^{\infty,A_{-i}}(0)=\max\left(\sup_{-t_i\le s\le 0}\left\{|s|(c-r_i)-\int_s^0 A_{-i}du\right\},\ B+(c-r_i)t_i-\int_{-t_i}^0 A_{-i}du\right).$$

By noting that $B+t_i(c-r_i)-\int_{-t_i}^0 A_{-i}du<0$ on event ξ_i, and $\rho_{-i}>c-r_i$ we conclude that on event ξ_i

$$\overline{Q}_{c-r_i}^{\infty,A_{-i}}(0)\le\sup_{s\le 0}\left\{|s|(c-r_i)-\int_s^0 A_{-i}du\right\}\stackrel{d}{=}Q_{c-r_i}^{\infty,A_{-i}}(0)<\infty\quad\text{a.s.}\tag{4.5}$$

Next, inequality (4.3) and inclusion (4.4) yield for all $B>2NTr$

$$\Gamma_i^B\ge\mathbb{E}\left[\left(-KT^{-1}+T^{-1}\int_0^T A(u)du-c\right)\mathbf{1}\left\{\overline{Q}_{c-r_i}^{\infty,A_{-i}}(0)\le K,\ \alpha_i^i,\ \xi_i\right\}\right],$$

due to the fact that on event ξ_i for large enough B the workloads of all flows other than i are smaller than B/N for all $t\in(0,T)$ and, therefore, those flows do not experience loss of fluid.

Now, let t_i^* be the first time after $t=0$ that flow i is not generating fluid, i.e. $t_i^*\triangleq\inf\{t>0: A_i(t)=0\}$. Then,

$$\begin{aligned}\Gamma_i^B\ge&-(KT^{-1}+c)\mathbb{P}[\overline{Q}_{c-r_i}^{\infty,A_{-i}}(0)\le K,\ \alpha_i^i,\ \xi_i]\\&+T^{-1}\mathbb{E}\left[\int_0^T A_{-i}(u)du\,\mathbf{1}\{\overline{Q}_{c-r_i}^{\infty,A_{-i}}(0)\le K,\ \alpha_i^i,\ \xi_i\}\right]\\&+r_iT^{-1}\mathbb{E}\left[\min(T,t_i^*)\mathbf{1}\{\overline{Q}_{c-r_i}^{\infty,A_{-i}}(0)\le K,\ \alpha_i^i,\ \xi_i\}\right],\end{aligned}\tag{4.6}$$

where in the last term we used $\int_0^T A_i(u)du\ge r_i\min(T,t_i^*)$ on event α_i^i. Next, we explore asymptotic behavior of the three terms in the preceding equation. By the independence of $A_{-i}(t)$ and $A_i(t)$, the fact that $\lim_{B\to\infty}\mathbb{P}[\xi_i]=1$ and expression (4.5)

$$\underline{\lim}_{B\to\infty}\frac{\mathbb{P}[\overline{Q}_{c-r_i}^{\infty,A_{-i}}(0)\le K,\ \alpha_i^i,\ \xi_i]}{\mathbb{P}[\alpha_i^i]}\ge\mathbb{P}[Q_{c-r_i}^{\infty,A_{-i}}(0)\le K]\tag{4.7}$$

and

$$\begin{aligned}\underline{\lim}_{B\to\infty}\frac{\mathbb{E}\left[\int_0^T A_{-i}(u)du\,\mathbf{1}\{\overline{Q}_{c-r_i}^{\infty,A_{-i}}(0)\le K,\ \alpha_i^i,\ \xi_i\}\right]}{\mathbb{P}[\alpha_i^i]}&\ge\mathbb{E}\left[\int_0^T A_{-i}(u)du\,\mathbf{1}\{Q_{c-r_i}^{\infty,A_{-i}}(0)\le K\}\right]\\&\ge\mathbb{E}\left[\int_0^T A_{-i}(u)du\right]-rT\mathbb{P}[Q_{c-r_i}^{\infty,A_{-i}}(0)>K].\end{aligned}\tag{4.8}$$

Next, for all finite T, due to $\tau^{on}_{i,r} \in \mathcal{IR} \subset \mathcal{L}$ and the independence of arrival processes, it follows that

$$\begin{aligned}\varliminf_{B\to\infty} \frac{\mathbb{E}\left[\min(T,t_i^*)\mathbf{1}\{\overline{Q}^{\infty,A_{-i}}_{c-r_i}(0)\le K,\ \alpha_i^i,\ \xi_i\}\right]}{\mathbb{P}[\alpha_i^i]} &\ge T\mathbb{P}[Q^{\infty,A_{-i}}_{c-r_i}(0)\le K]\ \varliminf_{B\to\infty}\mathbb{P}[t_i^* > T|\alpha_i^i] \\ &= T\mathbb{P}[Q^{\infty,A_{-i}}_{c-r_i}(0)\le K]. \end{aligned} \tag{4.9}$$

Inequality (4.6), together with (4.7), (4.8) and (4.9), implies

$$\varliminf_{B\to\infty}\frac{\Gamma_i^B}{\mathbb{P}[\alpha_i^i]} \ge \left(-\frac{K}{T}+r_i-c+\mathbb{E}\left[\frac{1}{T}\int_0^T A_{-i}(u)du\right]\right)\mathbb{P}[Q^{\infty,A_{-i}}_{c-r_i}(0)\le K] - r\mathbb{P}[Q^{\infty,A_{-i}}_{c-r_i}(0)> K].$$

The preceding inequality holds for all $T>0$ and, therefore, passing $T\to\infty$ renders

$$\varliminf_{B\to\infty}\frac{\Gamma_i^B}{\Lambda^{B,c-\rho_{-i}}_{A_i}} \ge \left(\mathbb{P}[Q^{\infty,A_{-i}}_{c-r_i}(0)\le K] - \frac{r}{r_i+\rho_{-i}-c}\mathbb{P}[Q^{\infty,A_{-i}}_{c-r_i}(0)> K]\right)\varliminf_{B\to\infty}\frac{\mathbb{P}[\tau^{on}_{i,r} > \frac{(1+\epsilon)B}{r_i+\rho_{-i}-c}]}{\mathbb{P}[\tau^{on}_{i,r} > \frac{B}{r_i+\rho_{-i}-c}]}.$$

Finally, by setting $\epsilon\downarrow 0$, recalling from (4.5) that $Q^{\infty,A_{-i}}_{c-r_i}(0)$ is almost surely finite and letting $K\to\infty$ the proof of the lower bound follows. This completes the proof of the theorem. $\diamond$

Possible generalizations At this point we would like to discuss possible generalizations of the preceding theorem. The strict stability condition $c\phi_i > \rho_i$ for all i represents a natural engineering condition. However, from a theoretical perspective the behavior of the system remains unclear if one or more flows have higher average demands than their minimum guaranteed rates. The following corollary represents an easy extension of this type. It states that a flow with guaranteed service rate lower than its expected rate will not be asymptotically affected by other flows if the tail of that flow is sufficiently heavy. In general, the problems of this kind are difficult and remain open.

Corollary 4.1 *Let $\tau^{on}_j \in \mathcal{IR}$ for $1\le j\le N$ and $r_i+\rho_{-i} > c > \rho$ for some flow i. If $\phi_j c > \rho_j$ for all $j\neq i$ and $\prod_{j\in\mathcal{Q}}\mathbb{P}[\tau^{on}_{j,r} > x] = o\left(\mathbb{P}[\tau^{on}_{i,r}>x]\right)$ for all sets $\mathcal{Q}\subseteq\{1,\ldots,N\}$ that satisfy $\sum_{j\in\mathcal{Q}} r_j + \sum_{j\notin\mathcal{Q}}\rho_j \ge c$, then as $B\to\infty$*

$$\Gamma_i^B \sim \Lambda^{B,c-\rho_{-i}}_{A_i}.$$

Proof: The upper bound is a direct consequence of inclusion (2.4) and Theorem 3.2. The proof of the lower bound is the same as in Theorem 4.1. $\diamond$

Careful examining of the proof of Theorem 4.1 shows that it holds for much more general buffer management policies. In particular, this includes the GPS-like rules for buffer management, which can be utilized to improve the behavior of the system for small queue sizes.

5. NUMERICAL EXAMPLES

In this section we present numerical validations of the obtained asymptotic results. Consider a fluid queue with capacity $c = 2.5$ and finite buffer B shared by five On-Off flows. Since the asymptotic results are insensitive to the distribution of Off periods, we choose their distribution to be exponential, i.e. $\mathbb{P}[\tau^{off} > x] = e^{-\mu x}$, $x \geq 0$. On periods are selected from the Pareto family, $\mathbb{P}[\tau^{on} > x] = x^{-\alpha}$, $x \geq 1$, $\alpha > 1$. The peak rates of On-Off flows are defined by vector $(4, 2, 2, 3, 1)$ and the On probabilities are chosen to be $p_{on,i} = 0.1$ for all five flows. The length of the simulated sample path in both examples is set to 10^9.

Example 1 Let the distributions of On periods be defined by $\alpha = (1.6, 1.5, 1.6, 1.5, 1.6)$. The work of the GPS mechanism is fully determined by the weight vector ϕ =(0.3, 0.1, 0.3, 0.2, 0.1). Clearly, the conditions of Theorem 4.1 are satisfied for the first flow. For this flow we simulated the loss rates for buffer sizes $B = 100, 200, \ldots, 1000$. The results of the simulation are presented in Figure 1 with "o" symbols. The approximation of the loss rate is plotted in the same figure with a solid line.

Example 2 Here, consider the system from the previous example with α =(1.6, 3.0, 3.0, 3.4, 1.9) and $\phi = (0.1, 0.2, 0.3, 0.3, 0.1)$. In this example it is easy to verify that the conditions of Corollary 4.1 are satisfied for the first flow. The simulation results ("o" symbols) and the approximation (solid line) for the loss rate are plotted on Figure 2.

It is apparent from both figures that the derived approximations are in excellent agreement with the simulated results. Furthermore, we would like to point out that the asymptotic formulas are very precise even though the estimated probabilities are relatively high, in range of $10^{-2} - 10^{-3}$.

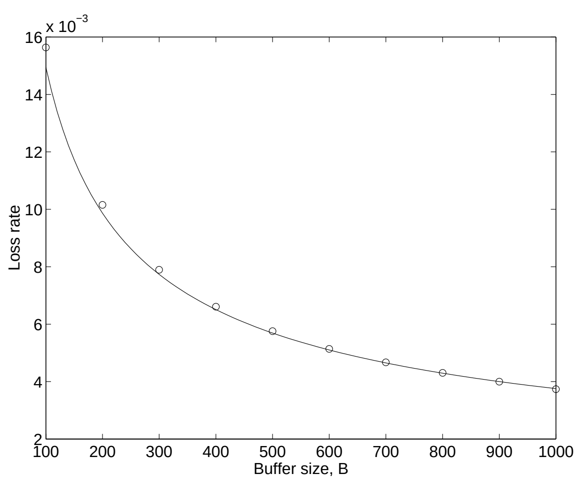

Figure 1. Illustration for Example 1.

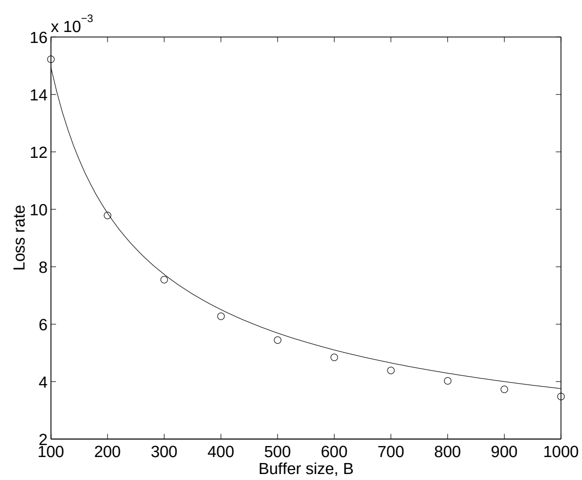

Figure 2. Illustration for Example 2.

6. CONCLUDING REMARKS

In this paper we analyzed a system of N independent heavy-tailed On-Off flows that share a server and a common buffer of size B. The server capacity is divided according to the GPS scheduling discipline and the buffer sharing is unrestricted, unless the buffer is full. When the buffer is full, we introduce a low complexity workload management policy that discards the necessary amount of fluid from the most demanding flows. Each flow is assumed to have its minimum rate guarantee larger than its long-time average rate. Our main result shows that asymptotically each flow experiences the same loss rate as if it were served in isolation with constant capacity and the whole buffer space B exclusively dedicated to it.

The main novelty of our result is that the analyzed system behaves as if it had N times larger buffer than it actually did. This new insight provides additional guideline for deciding on the buffer distribution between the access and core network switching elements. Finally, we observe that the result raises an interesting problem of socially fair buffer pricing. Everybody gets a full share of the resource, but it is not quite clear who and by how much one should pay for it.

APPENDIX: HEAVY-TAILED DISTRIBUTIONS

Definition A.1 *A non-negative random variable X (or its d.f. F) is called long-tailed $X \in \mathcal{L}$ ($F \in \mathcal{L}$) if*

$$\lim_{x\to\infty} \frac{1-F(x-y)}{1-F(x)} = 1, \quad \forall y \in \mathbb{R}.$$

Definition A.2 *A non-negative random variable X (or its d.f. F) is called subexponential $X \in \mathcal{S} \subset \mathcal{L}$ ($F \in \mathcal{S} \subset \mathcal{L}$) if*

$$\lim_{x\to\infty} \frac{1-F^{2*}(x)}{1-F(x)} = 2,$$

where F^{2*} denotes the 2-fold convolution of F with itself, i.e., $F^{2*}(x) = \int_{[0,\infty)} F(x-y)F(dy)$.

Definition A.3 *A non-negative random variable X (or its d.f. F) is called intermediately regularly varying $X \in \mathcal{IR} \subset \mathcal{S} \subset \mathcal{L}$ ($F \in \mathcal{IR} \subset \mathcal{S} \subset \mathcal{L}$) if*

$$\lim_{\eta\uparrow 1} \overline{\lim_{x\to\infty}} \frac{1-F(\eta x)}{1-F(x)} = 1.$$

Lemma A.1 *Let $F \in \mathcal{IR}$, $\eta \in (0,1)$, then*

$$\sup_{x\in[0,\infty)} \frac{1-F(\eta x)}{1-F(x)} < \infty.$$

Proof: Follows immediately from the definition. $\diamond$

REFERENCES

1. D. Bertsimas, I.C. Paschalidis, and J.N. Tsitsiklis. Large deviations analysis of the general processor sharing policy. *Queueing Systems*, 32:319–349, 1999.
2. S. Borst, O. Boxma, and P. Jelenković. Generalized processor sharing with long-tailed traffic sources. In P. Key and D. Smith, editors, *Teletraffic Engineering in a Competitive World, Proc. ITC-16*, pages 345–354, Edinburgh, UK, 1999. North-Holland, Amsterdam.
3. S. Borst, O. Boxma, and P. Jelenković. Induced burstiness in generalized processor sharing queues with long-tailed traffic flows. In *37th Annual Allerton Conference on Communication, Control, and Computing*, 1999.
4. S. Borst, O. Boxma, and P. Jelenković. Asymptotic behavior of generalized processor sharing with long-tailed traffic sources. In *INFOCOM 2000*, Tel-Aviv, Israel, 2000.
5. S. Borst, O. Boxma, and P. Jelenković. Coupled processors with regularly varying service times. In *INFOCOM 2000*, Tel-Aviv, Israel, 2000.
6. S. Borst, O. Boxma, and P. Jelenković. Reduced-load equivalence and induced burstiness in GPS queues with long-tailed traffic flows. preprint, 2000.
7. M. Crovella and A. Bestavros. Self-similarity in world wide web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, 1997.
8. P. Dupuis and K. Ramanan. A Skorohod problem formulation and large deviation analysis of a processor sharing model. *Queueing Systems*, 28(1-3):109–124, 1998.
9. G. Fayolle, I. Mitrani, and R. Iasnogorodski. Sharing a processor among many job classes. *Journal of ACM*, 27(3):519–532, July 1980.
10. D. P. Heyman and T. V. Lakshman. Source models for VBR broadcast-video traffic. *IEEE/ACM Transactions on Networking*, 4:40–48, February 1996.
11. P. Jelenković and P. Momčilović. Asymptotic loss probability in a finite buffer fluid queue with hetrogeneous heavy-tailed on-off processes. submitted for publication, available at `http://www.comet.columbia.edu/~predrag/mypub/fbMux.ps`, 2000.
12. P. Jelenković and P. Momčilović. Capacity regions for network multiplexers with heavy-tailed fluid on-off sources. In *INFOCOM 2001*, Anchorage, Alaska, April 2001.
13. P. R. Jelenković. Subexponential loss rates in a GI/GI/1 queue with applications. *Queueing Systems*, 33:91–123, 1999.
14. P. R. Jelenković and A. A. Lazar. Asymptotic results for multiplexing subexponential on-off processes. *Advances in Applied Probability*, 31(2), June 1999.
15. P. R. Jelenković, A. A. Lazar, and N. Semret. The effect of multiple time scales and subexponentiality of MPEG video streams on queueing behavior. *IEEE JSAC*, 15(6):1052–1071, August 1997.
16. W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic. In *SIGCOMM'93*, pages 183–193, 1993.
17. R. M. Loynes. The stability of a queue with non-independent inter-arrival and service times. *Proc. Cambridge Philos Soc.*, 58:497–520, 1962.
18. L. Massoulié. Large deviations estimates for polling and weighted fair queueing service systems. *Advances in Performance Analysis*, to appear.
19. K. Park and W. Willinger, editors. *Self-similar Network Traffic and Performance Evaluation*. John Wiley and Sons, 2000.
20. Z.-L. Zhang. Large deviations and the generalized processor sharing scheduling for a multiple-queue system. *Queueing Systems*, 28:349–376, 1998.
21. Z.-L. Zhang, D. Towsley, and J. Kurose. Statistical analysis of the generalized processor sharing scheduling discipline. *IEEE JSAC*, 13(6), 1995.