Alcom-FT Technical Report Series
ALCOMFT-TR-03-13

Heavy tailed $\mathbf{M/G/1}$-PS queues with impatience and admission control in packet networks

Jacqueline Boyer Fabrice Guillemin Philippe Robert* Bert Zwart

***Abstract*—In this paper we analyze the $\mathrm{M/G/1}$ processor sharing queue with heavy tailed services and with impatient customers. It is assumed that impatience depends on the value of the service required. We prove that a reduced service rate (RSR) approximation holds for estimating the sojourn time of a customer in the system, when the queue capacity is finite or infinite. This allows us to evaluate the reneging probability of customers with very large service times. We then use these results to investigate the impact of admission control on a link of a packet network. Admission control simply consists of limiting the number of simultaneous connections. It turns out that there is a real benefit for the efficiency of the system to perform admission control: It globally increases the fraction of customers, who complete their service (i.e. without being impatient). Finally, we investigate the fairness of the system and propose a criterion to assess the capacity of the system so as to allow the completion of very large service times.**

***Index Terms*—Admission control. Elastic traffic. Processor sharing. Reduced service rate. Heavy tails.**

I. INTRODUCTION

In spite of the fact that over the past few years quality of service (QoS) issues have become less critical for backbone packet networks thanks to over-dimensioning, QoS is still a major concern for network operators in the design of access networks. Before this dichotomy between access and backbone networks, the need for QoS to emerging multimedia and business Internet applications has led the IETF to introduce a series of general traffic management tools. This has given rise to the IntServ and DiffServ models along with the specification of MPLS. All these tools are aimed at describing, in one way or another, the packet flows generated by applications in order to provision sufficient resources in the network so as to control the QoS level offered by the network; QoS is in general expressed in terms of minimum bandwidth, transfer delays, information loss, etc... For this purpose, packet flows, which may correspond to a single application (micro-flow) or to an aggregate (macro-flow), are described by means of traffic parameters, namely the well-known (σ, ρ, π) triplet, which is the basic object of network calculus. (The quantities π and ρ denote the mean and peak rates, respectively and σ is the bucket size.) Traffic parameters are used in particular by admission control algorithms implemented in network elements, which are in charge of limiting the load so as to meet the QoS requirements.

Recognizing that the declaration of traffic parameters is a difficult task for users in general and that a small inaccuracy when assessing traffic parameters may cause excessive packet discarding by traffic policers and then result in a severe degradation of the QoS perceived by users, some authors have recently advocated the introduction of a "smoother" admission control. For instance, Bonald and Roberts [1] and Bonald *et al.* [2] have proposed to control an Internet bottleneck link by blocking some of the flows when its load is above some threshold. For this purpose, they suggest the introduction of an "appli flow" identifier in order to clearly identify flows on the bottleneck. Filtering flows can also be performed by examining port numbers in TCP segments.

One of the major characteristics of the current Internet is in that no flows are blocked by the network. The counterpart is that the number of flows simultaneously active on a link may be very large. As a consequence, the service rate of each flow may take very small values. This may lead some users to interrupt their flows due to impatience (for instance when retrieving large files). This gives rise to reneging, which corresponds to events that abort service prior completion. A fraction of the link rate is thus used by flows, which are eventually not completed, thereby increasing the reneging probability of other flows.

In this paper the performance of a network link supporting long elastic flows, which may reneg, is investigated. A network link crossed by elastic traffic is con-

Jacqueline Boyer and Fabrice Guillemin are with France Telecom R&D, DAC/CPN, 22300 Lannion, France. E-mail: {Jacqueline.Boyer,Fabrice.Guillemin}@rd.francetelecom.fr

* Partially supported by Contract 00-1B320 with France Telecom and the Future and Emerging Technologies programme of the EU under contract number IST-1999-14186 (ALCOM-FT)

Philippe Robert and Bert Zwart are with INRIA Rocquencourt, RAP project, Domaine de Voluceau, 78153 Le Chesnay, France. E-mail: {Philippe.Robert,Bert.Zwart}@inria.fr

sidered and the transmission capacity is assumed to be shared according to processor sharing discipline, which is the ideal bandwidth sharing achieved by TCP. By assuming that impatience is proportional to service and that service times are heavy tailed, a simple admission control of flows is proposed. The improvement of the global performance of the system is examined as well as the positive impact on the reneging probability of a given flow admitted in the system. Problem of fairness are also discussed: does a very long flow always reneg even with admission control ? To address this point, we compute the reneging probability of very large files and we investigate how reneging could be avoided for such flows.

The basic queueing system considered in this analysis is the $M/G/1$ processor sharing (PS) queue with or without impatience and with finite or infinite capacity. It is generally quite difficult to get explicit results concerning queues with impatience. For this reason, the purpose of this paper is to give *qualitative* results on the behavior of these queues. For the $M/G/1$ PS queue under consideration, a result of independent interest is proved: under some assumptions, a RSR (Reduced Service Rate) approximation (also called Reduced Load Equivalence, see Agrawal *et al.* [3]) is shown to hold when the services are heavy tailed. This RSR result sheds some light on the fairness issues mentioned above and may be used to reduce the impatience of very large flows.

The organization of this paper is as follows: In Section II, the problem formulation is presented. In Section III, a general reduced service rate (RSR) approximation is proved. Since the quantities of interest are rather difficult to compute explicitly, an approximating model is introduced. In Section IV, the reneging probability probability for the global system and for large flows is investigated. Simulation results and comparisons between the real system and the approximating model are also presented in this section. In Section V, we discuss some possible extensions of the model. Concluding remarks are presented in Section VI.

II. Problem formulation

A. Modeling a bottleneck link

Consider a bottleneck link, which may be for instance the link at the output of a LAN, connecting the LAN to the Internet, or the link between an access network (e.g., an ADSL area) and the Internet backbone. On this link, several flows with various characteristics are multiplexed. In a first approach, let us consider elastic flows only, that is, flows which are able to adapt to the state of the network (for instance to adjust their transmission rate to the limited transmission capacity at the bottleneck link). In the Internet, such flows are regulated by TCP. Other flows (e.g., UDP flows) will appear as noise later in the analysis.

Some flows may be very short and correspond to short file transfers (typically, Web page retrieval). Such flows do not leave the slow start phase and are not very sensitive to the rate sharing performed by TCP; those flows are referred to as *mice*. (See Floyd [4].) From a modeling point of view, the aggregate corresponding to the superposition of a large number of short flows is in fact a noise traffic, which reduces the transmission capacity of longer data transfer.

At the other extreme, long file transfers (referred to as *elephants*) are more sensitive to rate sharing achieved by TCP. (Typically, a large flow corresponds to a data transfer of several Mbytes.) In a first approximation it is assumed that TCP achieves a fair sharing of the transmission capacity described by the processor sharing (PS) discipline. See Kelly [5] and Massoulié and Roberts [6] for fairness issues in TCP and Thompson *et al.* [7] for the mice/elephants dichotomy. The performance of the system and the efficiency of admission control algorithms depend on the distribution of the volume (in Mbytes) of flows. Measurements on Internet links tend to show that the volume of long flows are heavy tailed (namely with a Pareto tail distribution). For this reason, an $M/G/1$ PS queue with heavy tailed service distributions shall be considered.

The time needed to transmit a file with the PS discipline is very sensitive to the load of the link. At equilibrium, the mean transmission time of a flow of size x is given by $x/c(1-\rho)$, where ρ is the average load of the flows and c the capacity of the link. When ρ is close to 1, the transmission time of a flow can be very large, which may cause a user to interrupt its session. A key property to account for when analyzing the quality of service perceived by users is their impatience.

In this context, since some flows have to be rejected, the purpose of *admission control* is to reject flows in a minimal way so that very few of the accepted flows shall experience impatience during their transmission. When properly done, admission control has several advantages:

- To reject flows before any processing, thereby reducing the overhead of the impatience.
- To increase the proportion of successful transmissions.

Since we are considering ordinary elastic traffic, contrary with usual situations in telecommunications, admission control cannot be performed by using an explicit characterisation of flows (like the (σ, ρ, π) triplet). In this setting, Bonald *et al.* [2] proposed admission control based on an estimation of the load of the link. The admission control considered in this paper is based on the number of accepted flows.

As mentioned earlier, we focus in a first step on traffic offered by elephants. The perturbation due to streaming flows and shorter flows (typically mice) will be be accounted for later on (see Section V). The primary objective of this paper is to qualitatively study the performance of admission control for long elastic flows.

B. *Definitions and notations*

Throughout this paper, we consider a single-server queue with server working at unit speed. Customers (namely long elastic flows) arrive at rate λ and require heavy-tailed service times with generic service times B and with impatience time C. Customers are served according to the PS discipline and a customer leaves the system if it has completed its service or if its sojourn time in the system exceeds C.

Impatience. Customers, who retrieve large files, are ready to wait more than customers retrieving smaller files. The impatience C of a customer may be any non-decreasing concave function of its service time B, if we consider that customers are more and more impatient as time to complete a file transfer becomes large. For admission control we shall assume that impatience is proportional to service, namely $C = \theta B$ for some real number $\theta > 1$.

Admission control. It simply consists in fixing some value N such that any connection is rejected, when N connections are already accepted. This is equivalent to say that the waiting room of the queue is of size N.

Loss probability. When the number of flows which can be simultaneously active is limited to N (the buffer capacity), P_N denotes the *rejection probability* (i.e., the fraction of customers, who cannot enter the system because of the buffer limitation) and κ_N denotes the fraction of customers, who leave the system because of impatience. The global loss probability π_N in the system is defined by

$$\pi_N = \kappa_N + P_N. \tag{1}$$

$1-\pi_N$ is the fraction of arriving customers, who complete their service (i.e., who can enter the system and who do not leave because of impatience). When the buffer capacity is infinite, the rejection probability is clearly equal to 0 (i.e., $P_\infty = 0$).

Efficiency of admission control. Admission control is said to be efficient when N can be chosen so that $\kappa_N \ll P_N$ while $\pi_N \leq \pi_\infty$, i.e., when impatience has significantly less impact than rejection on the global loss probability, which is itself less than the global loss probability without admission control. Thus, we may expect that the QoS perceived by users, who are admitted in the system, is better that the QoS when there is no control: less customers enter the system but once they are in, they do not leave because of impatience. Finally, one may expect an improvement of the utilization of the link (i.e., the fraction of time the server is transmitting bits, which belong to customers, who eventually complete their service).

III. A reduced service rate approximation

A. *Properties of the amount of service received by a customer*

Assuming that a tagged customer stays permanently in the queue, the function $(L(u); u \geq 0)$ denotes the number of (non-tagged) customers in the system at time u. The amount of service received by the tagged customer up to time x is thus given by

$$R(x) = \int_0^x \frac{1}{1+L(u)}\, du. \tag{2}$$

If B is some random variable independent of $(L(u))$ and

$$V = \inf\{y \geq 0; R(y) \geq B\},$$

then V is the time necessary to complete the service of a customer arrived at time 0. If this customer is not impatient, V is its sojourn time in the queue.

Proposition 1: When there is no impatience and provided that the queue is stable, we have

$$R(x)/x \stackrel{\text{def}}{\to} \gamma_N = 1-(1-b_N)\rho, \tag{3}$$

a.s. as $x \to \infty$, where b_N is the rejection probability in the processor sharing queue with a permanent customer and capacity N.

Proof: Let us consider a time interval $[0,t]$. Over this time interval, n_t customers enter the system and bring a total amount of work equal to $\sum_{k=1}^{n_t} B_k$ where B_k is the service time required by the kth customer. The total amount of work received by these customers is

$$\int_0^t \frac{L(u)}{1+L(u)} du.$$

Let D_t be the backlog of work remaining at time t. We have

$$D_t = \sum_{k=1}^{n_t} B_k - \int_0^t \frac{L(u)}{1+L(u)}\, du.$$

Since the system is stable $D_t/t \to 0$ when $t \to \infty$. It follows that

$$(1-b_N)\rho = \lim_{t\to\infty} \frac{1}{t} \int_0^t \frac{L(u)}{1+L(u)}\, du,$$

since $n_t/t \to (1-b_N)\lambda$ as $t \to \infty$. This completes the proof. ■

In the case of an $M/M/1$ PS queue with a permanent customer, the process describing the number of customers in the system over time is a birth and death process with birth rates equal to the arrival rate λ and death rates equal to $n\mu/(n+1)$, $n \geq 1$, where $1/\mu$ is the mean service time. When the queue size is $N < \infty$, the stationary distribution for the queue length is given by

$$\mathbb{P}(L=n) = \frac{(n+1)\rho^n}{\sum_{k=0}^{N-1}(k+1)\rho^k} \tag{4}$$

for $0 \leq n \leq N-1$, with $\rho = \lambda/\mu$. The rejection probability b_N is given by

$$b_N = \frac{(1-\rho)^2 N\rho^{N-1}}{N\rho^{N+1}-(N+1)\rho^N+1}. \tag{5}$$

Now, by using insensitivity results for the $M/G/1$ PS queue (See Burman [8] for an account on insensitivity results in queueing systems), the above expression for b_N holds for any service distribution. As a consequence, for an $M/G/1/N$ PS queue without impatience, we have

$$\gamma_N = (1-\rho)\frac{1}{1-N\rho^N(1-\rho)/(1-\rho^N)}.$$

When the impatience time of a customer is $C = \theta B$, along the same lines it is not difficult to show that there is some non-negative constant $\gamma_{\theta,N}$ such that almost surely

$$R(x)/x \to \gamma_{\theta,N} \tag{6}$$

as $x \to +\infty$. The quantity $\gamma_{\theta,N}$ is such that

$$\gamma_{\theta,N} = \mathbb{E}\left[\frac{1}{1+L}\right],$$

where L is the number of customers in the $M/G/1/N$ PS queue with a permanent customer. It turns out that this quantity is in general very difficult to compute explicitly. This is why we shall study in the following an approximating model.

B. *Sufficient conditions for RSR*

If B is sufficiently heavy-tailed, then the law of large number will pertain and lead to a result of the form of a RSR approximation (also called Reduced Load Equivalence), namely

$$\mathbb{P}(V > x) \sim \mathbb{P}(B > \gamma x) \tag{7}$$

when $x \to \infty$, where $\gamma = \gamma_{\theta,N}$ is defined by relation (6). So far, this result has been only proven for the $M/G/1$ queue with infinite buffer size and without impatience. (See Jelenković and Momčilović [9].) Our goal in this section is to show that equation (7) holds for other systems as well, in particular when customers may be impatient.

For convenience, we assume that the distribution function of B has a regularly varying tail, i.e.,

$$\mathbb{P}(B > x) = \ell(x)x^{-\nu}, \tag{8}$$

with $\nu > 1$ and ℓ a slowly varying function. The following result is crucial.

Proposition 2: The RSR relation (7) holds when

1) B is regularly varying of index $\nu > 1$;
2) $R(x)/x \to \gamma$ a.s. as $x \to \infty$ with $0 < \gamma < 1$;
3) There exists a positive and finite constant K such that $\mathbb{P}(R(x) \leq x/K) = \mathrm{o}(\mathbb{P}(B > x))$.

Proof: Define $g(x) = -\log \mathbb{P}(B > x)$. Since B is regularly varying, then g (which is non-decreasing) satisfies

$$\lim_{\varepsilon \downarrow 0} \limsup_{x\to\infty} g(x(1+\varepsilon)) - g(x) = 0.$$

Thus, for every $\delta > 0$ there exist an $\varepsilon > 0$ and x_0 such that for all $x \geq x_0$, we have

$$g(x+\varepsilon x) \leq g(x) + \delta.$$

Iterating this inequality we get, for $y \geq x_0$,

$$g(y + k\varepsilon y) \leq g(y) + k\delta$$

if $x = y + k\varepsilon x$; this leads to the inequality

$$g(x) - g(y) \leq \delta \left\lceil \frac{x-y}{\varepsilon y} \right\rceil \leq C_0 \left(\frac{x-y}{y} + 1 \right), \tag{9}$$

for $C_0 = \delta/\varepsilon$ and $x \geq y \geq \max(x_0, x/K)$.

Since $\mathbb{P}(B > R(x)) = I + II + III$, where

$$\begin{aligned} I &= \mathbb{P}(B > R(x); R(x) \geq (\gamma+\varepsilon)x), \\ II &= \mathbb{P}(B > R(x); (\gamma-\varepsilon)x < R(x) < (\gamma+\varepsilon)x), \\ III &= \mathbb{P}(B > R(x); R(x) \leq (\gamma-\varepsilon)x), \end{aligned}$$

the three terms I, II and III are treated separately. Note that term I is less than

$$\mathbb{P}(B > (\gamma+\varepsilon)x)\,\mathbb{P}(R(x) > (\gamma+\varepsilon)x),$$

which is $\mathrm{o}(\mathbb{P}(B > \gamma x))$ by Condition 2). Term II can be lower bounded by

$$\begin{aligned} &\mathbb{P}(\gamma x - \varepsilon x < R(x) < \gamma x + \varepsilon x)\,\mathbb{P}(B > \gamma x + \varepsilon x) \\ &\qquad \sim \mathbb{P}(B > \gamma x + \varepsilon x), \end{aligned}$$

and upper bounded by $\mathbb{P}(B > \gamma x - \varepsilon x)$, hence, after letting ε go to 0 and using the regular variation property of the distribution of B, one gets that the quantity II is equivalent to $\mathbb{P}(B > \gamma x)$.

We now turn to Term III. Write $III = IIIa + IIIb$ distinguishing between the two respective cases $R(x) < x/K$ and $R(x) \geq x/K$. Condition 3) implies that $IIIa$ can be neglected. To deal with IIIb, note that

$$IIIb = \mathbb{E}\left(\mathrm{e}^{-g(R(x))}; x/K < R(x) < (\gamma - \varepsilon)x\right)$$

and then

$$\begin{aligned} IIIb &= \mathbb{P}(B > \gamma x) \\ &\quad \times \mathbb{E}\left(\mathrm{e}^{g(\gamma x) - g(R(x))}; x/K < R(x) < (\gamma - \varepsilon)x\right). \end{aligned}$$

We must show that the second factor converges to 0 as $x \to \infty$. Inequality (9) combined with the relation $x/K < R(x) < \gamma x$ gives

$$\begin{aligned} g(\gamma x) - g(R(x)) &\leq C_0\left((\gamma x - R(x))/R(x) + 1\right) \\ &\leq C_0\left(K(\gamma x - R(x))/x + 1\right) \leq C_0(K\gamma + 1), \end{aligned}$$

hence $IIIb$ is upper bounded by

$$\mathbb{P}(B > \gamma x)\, \mathrm{e}^{C_0(K\gamma+1)} \mathbb{P}(R(x) < (\gamma - \epsilon)x).$$

This last quantity is $\mathrm{o}(\mathbb{P}(B > \gamma x))$ by Condition 2). This completes the proof. ■

The above result is analogous to Theorem 5.1.1 in Núñez-Queija [10]. The main strength of Proposition 2 is the weakness of the third condition. In Núñez-Queija [10] a related condition needs to be satisfied *for every* $K > 1/\gamma$. As indicated in [10], this condition can only be checked if detailed information of the sojourn-time distribution is known. In a different context, similar arguments as in the above proof can be found in Foss and Korshunov [11].

It is easy to see that the conditions of Proposition 2 are satisfied if the buffer capacity is finite (just choose K in the third condition of Proposition 2 as the buffer size). In the next section we show that under some weak condition, the RSR approximation applies in the case of an $M/G/1/\infty$-PS with impatience.

C. Infinite buffer with impatience

In this section, we consider an infinite processor sharing queue with heavy tailed service and Poisson arrivals with rate λ. We assume that the impatience time C of a customer is arbitrary and has a finite pth moment.

Proposition 3: Whatever be the dependence of the impatience C with respect to service time, as long as it has a finite pth moment for some $p > 1$ and the service time B is regularly varying of index $\nu > 1$, the RSR approximation (7) holds.

Proof: It is sufficient to show that Condition 3) of Proposition 2 holds. The variable $L(u)$ is stochastically upper bounded by the number of customers $L_1(u)$ in an $M/G/\infty$ queue with service time I and arrival rate λ. We obtain that

$$R(x) \geq \int_0^x \frac{1}{1 + L_1(u)}\, du.$$

For $\epsilon > 0$, a customer is said to be "small" if its service time is smaller than ϵx, otherwise, the customer is "large".

Let $H_1(\epsilon, x)$ be the number of large customers in the system at time 0, clearly enough $H_1(\epsilon x)$ has a Poisson distribution with rate $\lambda\mathbb{E}(I, I > \epsilon x)$. Now let $H_2(\epsilon, x)$ be the number of large customers arrived between time 0 and x, $H_2(\epsilon, x)$ has a Poisson distribution with rate $\lambda x \mathbb{P}(I > \epsilon x)$.

Finally, set $H(\epsilon, x) = H_1(\epsilon x) + H_2(\epsilon x)$. Note that $H(\epsilon, x)$ has a Poisson distribution with rate

$$\mu_1(\epsilon x) \stackrel{\text{def}}{=} \lambda\mathbb{E}\left(I 1_{\{I > \epsilon x\}}\right) + \lambda x \mathbb{P}(I > \epsilon x),$$

Thus, there exists some $C = C_\epsilon$ such that the inequality $\mu_1(\epsilon x) \leq C x^{1-p}$ holds. This implies that for each k and each $\epsilon > 0$,

$$\mathbb{P}(H(\epsilon, x) > k) = \mathrm{O}\left(x^{-k(p-1)}\right) = \mathrm{o}(\mathbb{P}(B > x)),$$

the latter inequality being valid if k is chosen big enough such that $k(p-1) > \nu$. From now on k is chosen so that this is satisfied.

The right hand side of the above inequality

$$\mathbb{P}\left(R(x) < \frac{x}{K}\right) \leq \mathbb{P}\left(\int_0^x \frac{1}{1 + L_1(u)}\, du < \frac{x}{K}\right)$$

can be split as $I + II$ according to the possibilities $H(\epsilon, x) \leq k$ or $H(\epsilon, x) > k$. Note that II is upper bounded by

$$\mathbb{P}(H(\epsilon, x) > k) = \mathrm{O}\left(x^{-k(p-1)}\right) = \mathrm{o}(\mathbb{P}(B > x)).$$

It has to be shown that I is also $\mathrm{o}(\mathbb{P}(B > x))$. For $k \geq 0$, if $H(\epsilon, x) \leq k$ then

$$L_1(u) \leq k + L_{1,<\epsilon x}(u),$$

with $L_{1,<\epsilon x}(u)$ denoting the number of small customers at time u in the queue, $\tau(\epsilon, x)$ denotes the number of busy periods of $(L_{1,<\epsilon x}(u))$ completed at time x. Note that, during each busy period $1 \leq i \leq \tau(\epsilon, x)$, the permanent customer gets at least $E_i/(k+1)$ units of service, with E_i exponentially distributed with rate λ. (A busy period always starts with one customer for an exponential amount of time.)

If for $n \geq 0$, $S_n \stackrel{\text{def}}{=} E_1+\cdots+E_n$, we can conclude that, if $H(\epsilon, x) \leq k$ then $R(x) \geq S_{\tau(\epsilon,x)}/(k+1)$. Thus it is enough to show that there exists a finite constant K such that $\mathbb{P}\left(S_{\tau(\epsilon,x)} \leq x/K\right) = \mathrm{o}(\mathbb{P}(B > x))$. For $a > 0$,

$$\mathbb{P}\left(S_{\tau(\epsilon,x)} \leq x/K\right) \leq \mathbb{P}(\tau(\epsilon,x) < ax) + \mathbb{P}(S_{ax} < x/K),$$

clearly the second term decreases exponentially fast in x, provided that K is such that $a\lambda K > 1$. Thus, it remains to show that

$$\mathbb{P}(\tau(\epsilon,x) < ax) = \mathrm{o}(\mathbb{P}(B > x)).$$

for some ϵ and a.

For $i \geq 1$, $P_i(\epsilon x)$ denotes the length of the ith busy cycle of the $M/G/\infty$ queue when only small customers are considered, $P_0(\epsilon x)$ the remaining busy cycle at time 0. The density function of $P_0(\epsilon x)$ is given by $\mathbb{P}(P_1(\epsilon x) \geq y)\, dy/\mathbb{E}(P_1(\epsilon x))$ for $y \geq 0$.

Since $\mathbb{E}(P_1(\epsilon x)) \leq \mathbb{E}(P_1(\infty))$, one can choose $a > 0$ and $\delta > 0$ small enough such that $a\mathbb{E}(P_1(\epsilon x)) < (1-\delta)$ for all $\epsilon > 0$ and $x > 0$. If for $n \geq 0$, T_n denotes the sum $P_1(\epsilon x) + \cdots + P_n(\epsilon x)$, then

$$\begin{aligned}\mathbb{P}(\tau(\epsilon,x) < ax) &\leq \mathbb{P}\left(P_0(\epsilon x) + T_{\lfloor ax \rfloor} > x\right)\\ &\leq \mathbb{P}(P_0(\epsilon x) > \delta x) + \mathbb{P}\left(T_{\lfloor ax \rfloor} > (1-\delta)x\right).\end{aligned}$$

This last expression is written as $III + IV$.

A bound on the tail distribution of $P_1(\epsilon x)$. Proposition 1 of Resnick and Samorodnitsky [12] implies that for each $\epsilon > 0$ there exists some $\beta > 0$ such that $\mathbb{P}(P_1(\epsilon x) > x) = \mathrm{o}(x^{-\beta})$. Consequently, this implies that $\mathbb{P}(P_0(\epsilon x) > x) = \mathrm{o}(x^{1-\beta})$.

From this result, it follows that, given δ, one can choose β large enough and ϵ small enough such that

$$III = \mathbb{P}(P_0(\epsilon x) > \delta x) = \mathrm{o}(\mathbb{P}(B > x)).$$

One has to show that $IV = \mathbb{P}\left(T_{\lfloor ax \rfloor} > (1-\delta)x\right)$ is also $\mathrm{o}(\mathbb{P}(B > x))$. Take $q > 0$, IV is certainly smaller than

$$\lfloor ax \rfloor \mathbb{P}(P_1(\epsilon x) > qx) + \mathbb{P}\left(P_1(\epsilon x) \wedge qx + \cdots + P_{\lfloor ax \rfloor}(\epsilon x) \wedge qx > (1-\delta)x\right).$$

For given q, the first term is of $\mathrm{o}(\mathbb{P}(B > x))$ if ϵ is chosen suitably small (w.r.t. q). The second term is smaller than

$$\mathbb{P}\left(P_1(\infty) \wedge qx + \cdots + P_{\lfloor ax \rfloor}(\infty) \wedge qx > (1-\delta)x\right). \quad (10)$$

Since I has finite pth moment, the same holds for $P_i(\infty)$, see e.g. Proposition 1 of Daley [13]. If $p_1 = \mathbb{E}(P_1(\infty))$, for $\eta > 0$ define

$$\widetilde{T}_n(qx) = P_1(\infty) \wedge qx + \cdots + P_n(\infty) \wedge qx - n(p_1 + \eta).$$

Then, the quantity (10) is given by

$$\begin{aligned}&\mathbb{P}\left(\widetilde{T}_{\lfloor ax \rfloor}(qx) > (1-\delta)x - (p_1+\eta)\lfloor ax \rfloor\right)\\ &\leq \mathbb{P}\left(\sup_{n \geq 0} \widetilde{T}_{\lfloor ax \rfloor}(qx) > (1-\delta)x - (p_1+\eta)\lfloor ax \rfloor\right).\end{aligned}$$

If η is chosen small so that $(1-\delta) - (p_1+\eta)a > 0$, then Lemma 3.2 (i) of the extension of Jelenković and Momčilović [9] gives that for any $\beta > 0$ there exists a $q > 0$ such that

$$\mathbb{P}\left(\sup_n \widetilde{T}_n(qx) > (1-\delta)x - (p_1+\eta)\lfloor ax \rfloor\right) = \mathrm{o}\left(x^{-\beta}\right).$$

This completes the proof. ■

To estimate the reneging probability, we are led to compute the quantities $\gamma_{\theta,N}$ when the buffer capacity is finite and the corresponding limiting value γ_θ, when the buffer capacity is infinite. $\gamma_{\theta,N}$ is equal to

$$\mathbb{E}\left[\frac{1}{1+L}\right]$$

where L is the number of customers in the stationary regime in the $M/G/1/N$ PS queue with a permanent customer. γ_θ is the corresponding quantity, when $N = \infty$. As mentioned earlier, these two quantities are very difficult to compute explicitly. This is why we study an approximating model in the next section.

D. *Approximating model*

Let us first consider an infinite capacity system and assume that service times are exponentially distributed. Given that in the system considered so far, impatience is proportional to service, when the number of customers in the system is greater than $\lfloor \theta \rfloor$, it is very likely that customers leave the system due to impatience. On the contrary, when the number of customers is less than $\lfloor \theta \rfloor$, the next departure is certainly due to service completion. We are thus led to consider a model based on the $M/M/1$ PS sharing queue, where customers arrive at the system according to a Poisson process with rate λ and leave the system

- due to service completion when the number of customers is less than or equal to $\lfloor \theta \rfloor$ (with rate μ),
- due to impatience when there are more than $\lfloor \theta \rfloor$ customers in the system (with rate $n\mu/\theta$).

The above system is slightly different from that analyzed by Coffman *et al* [14], where impatience is independent of service.

The process describing the number of customers in the above system is a birth and death process with birth rates $\lambda_n = \lambda$ for all $n \geq 0$ and death rates $\mu_n = \mu \max(n/\theta, 1)$ for $n \geq 1$. The stationary distribution of the queue size $\tilde{L}$ is then

$$\mathbb{P}\left(\tilde{L} = n\right) = \frac{1}{G_\infty} \begin{cases} \rho^n & n \leq \lfloor\theta\rfloor, \\ \rho^n \theta^{n-\lfloor\theta\rfloor} \lfloor\theta\rfloor!/n! & n > \lfloor\theta\rfloor, \end{cases} \tag{11}$$

where G_∞ is the normalizing constant, given by

$$G_\infty = \frac{1-\rho^{\lfloor\theta\rfloor+1}}{1-\rho} + \sum_{n=\lfloor\theta\rfloor+1}^{\infty} \frac{\rho^n \theta^{n-\lfloor\theta\rfloor}\lfloor\theta\rfloor!}{n!}.$$

Now, when the buffer capacity is finite (equal to $N > \lfloor\theta\rfloor$), the stationary distribution of the queue size $\tilde{L}_N$ is given by

$$\mathbb{P}\left(\tilde{L}_N = n\right) = \frac{1}{G_N} \begin{cases} \rho^n & n \leq \lfloor\theta\rfloor, \\ \frac{\rho^n \theta^{n-\lfloor\theta\rfloor}\lfloor\theta\rfloor!}{n!} & \lfloor\theta\rfloor < n \leq N. \end{cases}$$

with the normalizing constant G_N given by

$$G_N = \frac{1-\rho^{\lfloor\theta\rfloor+1}}{1-\rho} + \sum_{n=\lfloor\theta\rfloor+1}^{N} \frac{\rho^n \theta^{n-\lfloor\theta\rfloor}\lfloor\theta\rfloor!}{n!}.$$

The rejection probability $\tilde{P}_N$ is

$$\tilde{P}_N = \mathbb{P}\left(\tilde{L}_N = n\right) = \frac{1}{G_N}\frac{\rho^n \theta^{n-\lfloor\theta\rfloor}\lfloor\theta\rfloor!}{n!},$$

and the fraction of reneging customers (among all arriving customers) is

$$\tilde{\kappa}_N = \frac{1}{\rho\theta G_N} \sum_{n=\lfloor\theta\rfloor+1}^{N} n\frac{\rho^n \theta^{n-\lfloor\theta\rfloor}\lfloor\theta\rfloor!}{n!}. \tag{12}$$

It is worth noting that in this model, the fraction of reneging customers is null when $N \leq \lfloor\theta\rfloor$ and the rejection probability is then equal to the rejection probability in an $M/M/1/N$ queue.

Similar computations can be carried out in the case when there is one permanent customer in the system. The number $\tilde{L}'_N$ of customers in the system (excluding the permanent one) evolves as a birth and death process with birth rates λ_n equal to λ for all $n \geq 0$ and with birth rates μ_n given by

$$\mu_n = \begin{cases} n\mu/(n+1) & n \leq \lfloor\theta\rfloor - 1, \\ n\mu/\theta & n \geq \lfloor\theta\rfloor. \end{cases}$$

The stationary distribution of $\tilde{L}'_N$ is given by

$$\mathbb{P}\left(\tilde{L}'_N = n\right) = \frac{1}{G'_N} \begin{cases} (n+1)\rho^n, & n < \lfloor\theta\rfloor, \\ \lfloor\theta\rfloor \frac{\rho^n \theta^{n-\lfloor\theta\rfloor+1}(\lfloor\theta\rfloor-1)!}{n!}, & \lfloor\theta\rfloor \leq n \leq N-1, \end{cases} \tag{13}$$

where G'_N is the normalizing constant.

As we shall see in Section IV-B from simulation results, $\tilde{L}_N$ (resp. $\tilde{L}'_N$) yields a quite fair approximation of the number of customers in the $M/G/1$ PS queue with impatience and without a permanent customer (resp. with a permanent customer). As a consequence, we shall approximate the constant $\gamma_{\theta,N}$ by the mean value

$$\tilde{\gamma}_{\theta,N} = \mathbb{E}\left[\frac{1}{1+\tilde{L}'_N}\right].$$

IV. Admission control considerations

A. Minimal value for the choice of the buffer capacity

The $M/G/1$ PS model is used to describe the dynamic behavior of flows multiplexed on the bottleneck link. The impatience of a user is given by $C = \theta B$ with $\theta > 1$. If flows are rejected when the number of accepted flows is $N_0 = \lfloor\theta\rfloor$, then users cannot be impatient anymore since

$$\int_0^{\theta B} \frac{1}{1+L(u)}\, du \geq \frac{\theta}{N_0} B \geq B.$$

For this choice of the capacity, as soon as a flow is accepted, it certainly completes its transmission. Note that in this case, there is no overhead due to impatience. The blocking probability b_N defined by (5) determines the loss rate of this system.

In spite of the fact that this choice of the buffer size is very conservative, it turns out to be quite efficient in practice. The quantity b_N has to be compared with the reneging probability, which is the loss probability in the system when the buffer size is infinite.

The exact evaluation of the reneging probability, when the buffer is infinite, is very difficult. Nevertheless, it turns out that the approximating model described in Section III-D proves quite accurate, as shown in the next section, where simulation results are reported.

B. Simulation results

Figures 1-3 show the probability distribution of the number of customers in the $M/G/1$ PS queue with an impatience coefficient $\theta = 10$ for different values of the load

ρ, when service times are exponential or Pareto distributed with $\beta = 1$ and $\alpha = 1.25$. In this latter case, the complementary probability distribution function of service times is given by

$$\mathbb{P}\left(B > x\right) = \frac{1}{x^{\alpha}}.$$

for large x.

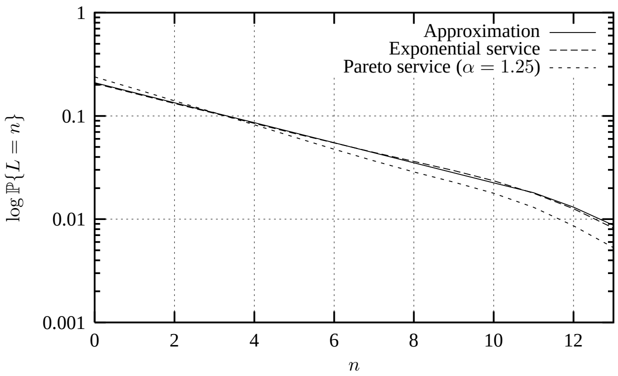

Fig. 1. Approximation vs. simulation results for exponential and Pareto ($\nu = 1.25$) service times for $\rho = 80\%$, $\theta = 10$.

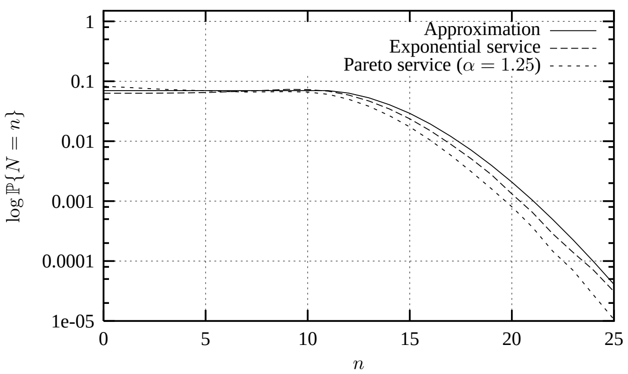

Fig. 2. Approximation vs. simulation results for exponential and Pareto ($\nu = 1.25$) service times for $\rho = 100\%$, $\theta = 10$.

From these figures, it turns out that the approximating model is quite accurate for estimating the number of customers in the queue, when the buffer capacity is infinite. The same simulation experiments have been performed when the impatience coefficient θ is equal to 3. The results also show that $\tilde{L}$ is a fair approximation of the number of customers in the queue with impatience.

It follows that the reneging probability can be well approximated by $\tilde{\kappa}_N$ given by equation (12) when $N = \infty$. Table I gives the reneging probability κ_∞ for exponential and Pareto service times. Once more, we can observe that the approximation $\tilde{\kappa}_\infty$ is quite accurate when the load ρ is not too large and the impatience coefficient θ is sufficiently large. Table I also reports the rejection probability

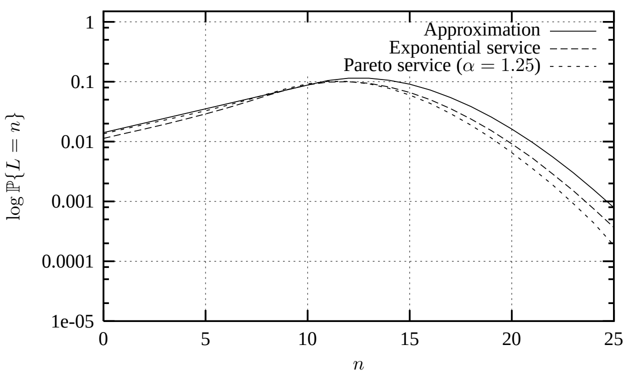

Fig. 3. Approximation vs. simulation results for exponential and Pareto ($\nu = 1.25$) service times for $\rho = 120\%$, $\theta = 10$.

when the buffer capacity is limited to $\lfloor\theta\rfloor$. The salient property of the system is that the global loss probability $P_{\lfloor\theta\rfloor}$ (in the case $N = \lfloor\theta\rfloor$, there is no reneging as observed in the previous section) is less than the loss probability when the buffer capacity is infinite (i.e., when there is no admission control). We hence deduce that there is a real gain in the global performance of the system, when admission control is performed.

To further investigate the benefit of admission control, Figure 4 displays the rejection probability, the reneging probability and the global loss probability as a function of the buffer capacity N for a load $\rho = 80\%$ and an impatience coefficient $\theta = 10$. We observe that for $N \geq \lfloor\theta\rfloor$, there is always a gain to perform admission control. On the contrary, if the buffer capacity is too drastically reduced, the rejection probability may be larger than the reneging probability in the system with no admission control. There is nevertheless a quite large range of the buffer capacity such that the global loss probability can be significantly reduced.

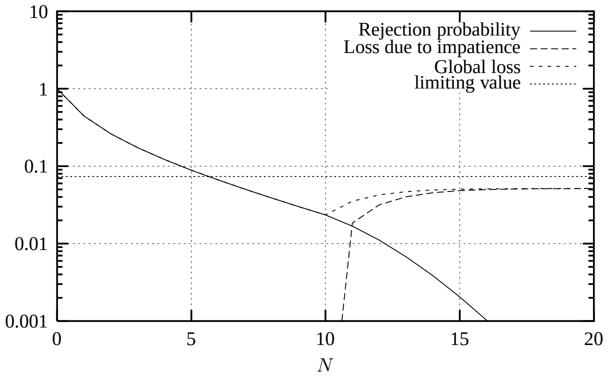

Fig. 4. Loss due impatience and to rejection for $\rho = 80\%$ and $\theta = 10$.

Figure 5 and Figure 6 show the same quantities as Figure 4 for a load $\rho = 100\%$ and a load $\rho = 120\%$; the

TABLE I
LOSS PROBABILITY κ_∞ FOR PARETO ($\nu = 1.25$) AND EXPONENTIAL SERVICE TIMES.

	$\rho = 80\%$		$\rho = 100\%$		$\rho = 120\%$	
	$\theta = 3$	$\theta = 10$	$\theta = 3$	$\theta = 10$	$\theta = 3$	$\theta = 10$
Pareto ($\nu = 1.25$)	0.3642	7.1014e-2	0.5574	.37630	0.7036	0.7158
Exponential	0.3856	7.7067e-2	0.5681	0.3718	0.7182	0.7279
Approximation	0.3016	7.3926e-2	0.641227	0.3720	0.9164	0.8470
$P_{\lfloor\theta\rfloor}$	0.1734	2.3493e-2	0.2500	0.0910	0.3219	0.1925

conclusions remain essentially the same.

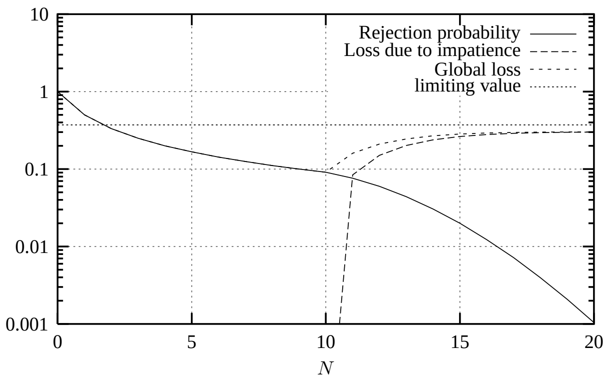

Fig. 5. Loss due impatience and to rejection for $\rho = 100\%$ and $\theta = 10$.

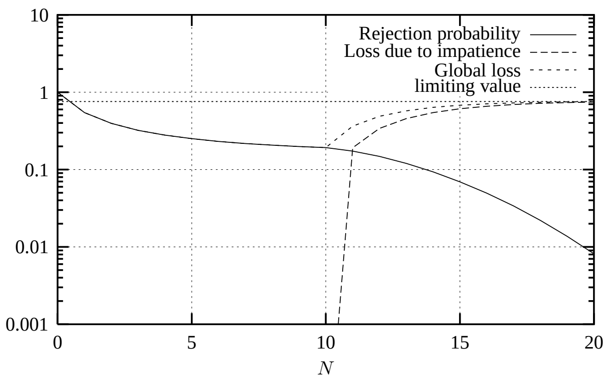

Fig. 6. Loss due impatience and to rejection for $\rho = 120\%$ and $\theta = 10$.

In spite of the fact that admission control improves the global performance of the system, one may wonder whether the buffer capacity can be increased while ensuring that very large elephants have a good chance of completing their transmission. This point, which is more or less related to the fairness of the system (any customer should have a good chance of completing its service) is addressed in the next section.

C. *Transmission of large files: fairness issues*

The value of $N_0 = \lfloor\theta\rfloor$ of Section IV-A is quite conservative. It indeed achieves complete fairness between all customers, since impatience is completely eliminated. But, one may wonder whether the capacity of system could be increased while keeping fairness between all customers. If the global objective is to reduce as much as possible the overhead due to impatience, an option is to secure the success of long transmissions when they are accepted. This option is also meaningful if the impatience of small transfers is supposed to quite large compared to their sizes. In this case admission control has to be fair for the large transfers. Define

$$N_1 = \sup\{N : \theta\gamma_{N,\theta} > 1\},$$

where $\gamma_{N,\theta}$ is defined by (6). Note that $N_1 \geq N_0$. The probability of reneging when the service is x is given by $\kappa_{N,\theta} = \mathbb{P}\left(R(\theta x)/x \leq 1\right)$.

If $\theta\gamma_{N,\theta} < 1$, this implies that $\kappa_{N,\theta}$ converges to 1 as $x \to +\infty$. In this case, if a large elephant is accepted it will not complete its transmission with a probability close to 1. On the other hand if $\theta\gamma_{N,\theta} > 1$, then a large elephant will reneg with a small probability. In this case the reneging probability can be estimated as follows. (See Boyer *et al.* [15] for the proof)

Proposition 4: If B is Pareto with index $\nu > 1$ and if $\theta\gamma_{\theta,N} > 1$, then there exists some $l^* \geq 1$ such that

$$\mathbb{P}\left(R(\theta x) < x\right) \leq C_0 \frac{1}{x^{l^*(\nu-1)}}.$$

This result has the following intuitive explanation, based on large-deviation arguments for heavy tails. For related arguments we refer to Likhanov and Mazumdar [16] and Zwart *et al.* [17].

To let the event $R(\theta x) < x$ happen, the average service rate should become smaller than $1/\theta$ for a long period of time. In Boyer *et al.* [15], it is shown that the most likely way for this to happen is that there are a certain number of other jobs in the system which have size $\mathrm{O}(x)$. Suppose that there are l other jobs in the system. together with the large job under consideration, the system behaves like a processor sharing queue with $l+1$ permanent customers. Let $\gamma_{\theta,N,l}$ be the average service rate for this system (with $\gamma_{\theta,N,0}=\gamma_{\theta,N}$). It is obvious that $\gamma_{\theta,N,l}$ is decreasing in l. The l in the proposition is the smallest possible number of other elephants making the service rate smaller than $1/\theta$, i.e., we have

$$l^* = \arg\min\{l : \gamma_{\theta,N,l} < 1/\theta\}.$$

We refer to Boyer *et al.* [15] for more precise assumptions, a formal proof and some extensions of this result. Note that, in the case $\theta = N = \infty$, we have $\gamma_{\infty,\infty,l} = (1-\rho)/(l+1)$, so in this case it is possible to compute l^* explicitly. In general, one has to resort to approximations.

Proposition 2 suggests the following approximation for the loss probability of an elephant of size x if $\theta\gamma_{\theta,N} > 1$:

$$\kappa_{N,\theta} \approx x^{-l^*(\nu-1)}.$$

This approximation is exponentially decreasing in l^*, while l^* becomes larger if N gets smaller. This suggests that admission control makes sense even if $\theta\gamma_{\theta,N} > 1$ for all N.

In any case, the buffer capacity N should be chosen between N_0 and N_1, a choice close to N_0 shall be conservative and if N is close to N_1, then overhead will significantly increase at the expense of large elephants.

From the above discussion, we see that that a criterion for choosing N is that we should have $\theta.\gamma_{\theta,N} > 1$. From a numerical point of view, it can happen that $\gamma_{\theta,N}$ tends to a limit $\gamma_{\theta,\infty}$, which is such that $\theta.\gamma_{\theta,\infty} > 1$. By taking into the approximating model described in Section III-D, the limiting value $\gamma_{\theta,\infty}$ can be approximated by the limit $\tilde{\gamma}_{\theta,\infty}$ defined by

$$\tilde{\gamma}_{\theta,\infty} = \mathbb{E}\left[\frac{1}{1+\tilde{L}'_\infty}\right].$$

For instance, when $\rho = 80\%$ and $\theta = 10$, the values of $\theta.\gamma_{\theta,N}$ as a function of N are given by Table II for the approximating model. In this case, it turns out that the limiting value $\tilde{\gamma}_{\theta,\infty}$ is such that $\theta.\tilde{\gamma}_{\theta,\infty} = 2.35 > 1$. This means that the capacity of the system can be increased without penalizing large elephants.

For the different set of values $\rho = 120\%$ and $\theta = 10$, the values of $\theta.\gamma_{\theta,N}$ are given by Table III. In this case, we see that the limiting value $\tilde{\gamma}_{\theta,\infty}$ is such that $\theta.\tilde{\gamma}_{\theta,\infty} = 0.926059 < 1$. In fact, the capacity of the system can be increased up to $N = 15$ without excessively penalizing large elephants.

V. Extensions

Some possible generalizations of the above results are briefly presented. (See the paper in preparation by Boyer *et al.* [15] for details.)

TABLE II

Values of $\theta.\gamma_{\theta,N}$ for different values of N when $\rho = 80\%$ and $\theta = 10$.

N	$\theta.\gamma_{\theta,N}$		
	Approx.	exponential	Pareto ($\alpha = 1.25$)
2	6.923077	6.9213	6.9921
3	5.398230	5.4029	5.4450
4	4.494519	4.4947	4.5566
5	3.901578	3.8990	3.9020
6	3.486347	3.4853	3.5601
7	3.182247	3.1799	3.2421
8	2.952249	2.9466	2.9469
9	2.774097	2.7649	2.9097
10	2.633591	2.6265	2.7229
11	2.530838	2.5085	2.5089
12	2.460334	2.4231	2.519
13	2.414770	2.3656	2.5210
14	2.387035	2.3331	2.3778
15	2.371150	2.3057	2.4185
16	2.362592	2.3096	2.3637
17	2.358255	2.3015	2.3514
18	2.356184	2.2952	2.2913
19	2.355250	2.3044	2.3360
20	2.354851	2.2914	2.3711

TABLE III

Values of $\theta.\gamma_{\theta,N}$ for different values of N when $\rho = 120\%$ and $\theta = 10$.

N	$\theta.\gamma_{\theta,N}$		
	Approx.	exponential	Pareto ($\alpha = 1.25$)
2	6.470588	6.4716	6.5017
3	4.715026	4.7155	4.7441
4	3.668671	3.6712	3.7112
5	2.976640	2.9790	2.9812
6	2.486837	2.4888	2.4775
7	2.123216	2.1235	2.1472
8	1.843530	1.8469	1.8467
9	1.622435	1.6214	1.6573
10	1.443806	1.4438	1.4608
11	1.306832	1.2914	1.2982
12	1.203413	1.1745	1.1849
13	1.125581	1.0909	1.0647
14	1.067186	1.0315	0.9976
15	1.023732	0.9885	0.9662
16	0.991889	0.9581	0.9411
17	0.969077	0.9357	0.9160
18	0.953214	0.9208	0.9150
19	0.942569	0.9099	0.9023
20	0.935707	0.9048	0.9067

A. Fluctuating service rate

To take into account the impact of mice or priority traffic, one may assume that the service capacity for long elastic flows fluctuates according to some process $Z(t)$, $t \geq 0$. The quantity $R(x)$ becomes

$$R(x) = \int_0^x \frac{Z(u)}{1 + L(u)}\, du.$$

To check the RSR approximation, there are again two cases: when the buffer size is finite or when there is impatience. In the first case,

$$R(x) \geq C_0 \int_0^x Z(u)\, du.$$

Thus, the validity of Condition 3 of Proposition 2 crucially depends on properties of the process Z. Some examples in which Condition 3 is satisfied are:

- $(Z(t))$ is a Gaussian process with stationary increments and covariance function regularly varying of index $2H < 2$. In particular, this includes fractional Brownian motion ($0 < H < 1$).
- $Z(u)$ is a semi-Markov process as long as periods during which $Z(u)$ is small are sufficiently light-tailed.

The case with $N = \theta = \infty$ has recently been investigated by Borst *et al.* [18], for the special case $Z(u) > \rho$. Using our methods, it is easy to recover their result: lower bound $R(x)$ with a standard PS having capacity $\rho + \epsilon$, and combine Propositions 2 and 4.

B. Discriminatory processor sharing

So far, we considered the case when bandwidth is equally shared among the different elastic flows (max-min fairness). One may, nevertheless, introduce some weighing coefficients, in order to give more bandwidth to some flows. This is precisely the task achieved by the so-called discriminatory processor sharing discipline. Consider the case with K customer classes, each class is assigned a weight coefficient ϕ_i, $i = 1, \ldots, K$.

To check whether the RSR approximation still holds, tag a customer of a given class, say, class 1. In this case $R(x)$ becomes (with obvious notation)

$$R(x) = \int_0^x \frac{\phi_1}{\phi_1 + \sum_{i=1}^K \phi_i N_i(u)} \mathrm{d}u.$$

Of course, Assumption 3 is satisfied if the total amount of customers in the system is bounded. If one considers the case with impatience and infinite buffer, one can prove condition 3 of Proposition 2 following exactly the same approach as in Section III.

C. The $M/G/s$ PS queue

The Processor sharing queue with multiple servers operates as follows. Suppose the number of customers in the system at a given time u is $Q(u)$. When $Q(u) \geq s$, each customer is being served with rate $s/Q(u)$. Otherwise, a customer receives service rate 1. This model formulation could be used to incorporate the fact that maximum transmission rates of files can be substantially smaller than the total system capacity, see e.g. Kherani and Kumar [19].

For this model, if $\theta = \infty$, one can still compute γ explicitly, using the fact the model is insensitive. moreover, it can also be shown that the RSR remains valid for all possible value of N and θ.

VI. Conclusion

In this paper we have used the $M/G/1$ PS processor sharing queue with heavy tailed services and with impatience proportional to service to qualitatively study the problematic of admission control of elastic flows in packet networks. in particular we have established an RSR approximation for computing the sojourn time of a flow in the system. This result has been proved under quite general assumptions. Finally, we have proposed an approximating model, which yields quite accurate results for certain values of the parameters (moderate load and sufficiently large impatience coefficient).

It turns out that under the specific assumption that impatience is proportional to service, there is a real gain in the global performance of the system, when admission control is performed. This function simply consists of limiting the number of flows, which can be simultaneously active on the link. The capacity of the system can be chosen so as to eliminate impatience. Moreover, numerical evidence shows that the capacity can be slightly increased, while ensuring a certain fairness (long flows still have a good chance of completing their service), at the expense of globally increasing impatience and the loss probability.

All the results are encouraging to continue investigations on admission control of elastic flows in packet networks. In particular, it has to be checked, whether the results obtained in this paper for impatience proportional to service remain valid, when impatience take a more general form (for instance an affine function of the service time or more generally a concave function of service). This point will be addressed in a further study.

References

[1] T. Bonald and J. Roberts, “Performance modeling of elastic trafic in overload,” in *SIGMETRICS'2001*, (Cambridge, MA, USA), June 2002.

[2] T. Bonald, S. Oueslaty-Boulahia, and J. Roberts, "QOS is still an issue: we need a new paradigm." France Telecom technical report, 2002.

[3] R. Agrawal, A. M. Makowski, and P. Nain, "On a reduced load equivalence for fluid queues under subexponentiality," *Queueing Systems. Theory and Applications*, vol. 33, no. 1-3, pp. 5–41, 1999.

[4] V. Paxson and S. Floyd, "Wide area traffic: the failure of Poisson modeling," *IEEE/ACM Transactions on Networking*, vol. 3, no. 3, pp. 226–244, 1995.

[5] F. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, 1998.

[6] L. Massoulié and J. Roberts, "Bandwidth sharing: Objectives and algorithms," in *INFOCOM '99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies*, pp. 1395–1403, 1999.

[7] K. Thompson, G. Miller, and R. Wilder, "Wide area internet traffic patterns and characteristics," *IEEE Network Magazine*, December/November 1997.

[8] D. Burman, "Insensitivity in queueing systems," *Advances in Applied Probability*, vol. 13, pp. 846–859, 1981.

[9] P. Jelenković and P. Momčilović, "Resource sharing with subexponential distributions," in *Infocom'2002*, (New York), June 2002.

[10] R. Núñez-Queija, *Processor Sharing Models for Integrated Services Networks*. PhD thesis, Eindhoven University of Technology, 2000.

[11] S. Foss and D. Korshunov, "Sampling at a random time with a heavy-tailed distribution," *Markov Processes and Related Fields*, vol. 6, pp. 643–658, 2000.

[12] S. Resnick and G. Samorodnitsky, "Activity periods of an infinite server queue and performance of certain heavy tailed fluid queues," *Queueing Systems*, vol. 33, pp. 43–71, 1999.

[13] D. J. Daley, "The busy period of the $M/GI/\infty$ queue," *Queueing Systems. Theory and Applications*, vol. 38, no. 2, pp. 195–204, 2001.

[14] E. Coffman, A. Puhalskii, M. Reiman, and P. Wright, "Processor shared queues with reneging," *Performance Evaluation*, vol. 19, no. 1, pp. 25–46, 1994.

[15] J. Boyer, F. Guillemin, P. Robert, and B. Zwart, "Tail asymptotics for processor sharing queues." In preparation, 2002.

[16] N. Likhanov and R. Mazumdar, "Cell loss asymptotics in buffers fed by heterogeneous longtailed," in *Proceedings of Infocom 2000*, (Tel Aviv, Israel), 2000.

[17] A. Zwart, S. Borst, and M. Mandjes, "Exact queueing asymptotics for multiple heavy-tailed on-off flows," in *Proceedings of Infocom 2001*, (Anchorage, USA), 2001.

[18] S. Borst, R. Núñez-Queija, and M. Van Uitert, "User-level performance of elastic traffic in integrated-services networks," in *Proceedings of Performance 2002*, (Rome, Italy), 2002.

[19] A. Kherani and A. Kumar, "Stochastic models for throughput analysis of randomly arriving elastic flows in the internet," in *Proceedings of Infocom 2002*, (New York, NY, USA), 2002.