Queueing Systems 33 (1999) 73–89

Scheduling strategies and long-range dependence

Venkat Anantharam *

Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, USA

E-mail: ananth@vyasa.eecs.berkeley.edu

Received March 1998; revised February 1999

Consider a single server queue with unit service rate fed by an arrival process of the following form: sessions arrive at the times of a Poisson process of rate λ, with each session lasting for an independent integer time $\tau \geqslant 1$, where $P(\tau = k) = p_k$ with $p_k \sim \alpha k^{-(1+\alpha)} L(k)$, where $1 < \alpha < 2$ and $L(\cdot)$ is a slowly varying function. Each session brings in work at unit rate while it is active. Thus the work brought in by each arrival is regularly varying, and, because $1 < \alpha < 2$, the arrival process of work is long-range dependent. Assume that the stability condition $\lambda E[\tau] < 1$ holds. By simple arguments we show that for *any* stationary nonpreemptive service policy at the queue, the stationary sojourn time of a typical session must stochastically dominate a regularly varying random variable having infinite mean; this is true even if the duration of a session is known at the time it arrives. On the other hand, we show that there exist causal stationary preemptive policies, which do not need knowledge of the session durations at the time of arrival, for which the stationary sojourn time of a typical session is stochastically dominated by a regularly varying random variable having finite mean. These results indicate that scheduling policies can have a significant influence on the extent to which long-range dependence in the arrivals influences the performance of communication networks.

Keywords: single server queue, scheduling policies, long-range dependence

1. Introduction

The past few years have witnessed an increasing interest in the phenomenon of *long-range dependence* in communication network traffic. This interest was sparked by a number of papers that presented statistical analyses of measurements of communication traffic from a wide variety of sources, see, for instance, [5,14,19]. In particular, a number of papers have studied analytical models of queues driven by long-range dependent arrivals, with the goal of deriving insights about the performance of communication networks that need to handle such traffic. Examples of such papers include [1,2,12,13,15–18,20,21]. A feature of most of these works is the observation that the tail behavior of queues with long-range dependent inputs decays much slower than exponentially, either as a Weibull law or according to a power law, depending on the

* Research supported in part by NSF grant NCR 94-22513.

model. This means that many performance related quantities of interest can be much larger than predicted by conventional arrival models. For instance, a consequence of the results of [1] is that the stationary mean sojourn time of a session at a single server first come first served (FCFS) queue, with the arrival process considered in that paper, is infinite; see also [8] for related results. In addition to such inferences from analyses, the presence of a qualitative difference in queueing behavior in queues driven by empirical traffic traces such as those measured in [14], as compared to that predicted by conventional models, is also observed – see, for instance, [10]. The experimental analysis of [10], by modifying the measured traffic traces of [14] in order to break up the correlations at either short or long time scales, makes a strong case for the hypothesis that this difference in queueing behaviour is due to long-range dependence. The persistence of long-range dependence in packet traffic even after it is regulated by flow control schemes closely related to those that have been standardized for use in practice is demonstrated in [20]. This further strengthens the case for the importance of studying the effects of long-range dependence on queueing systems. A bibliography of work in this area as of the middle of 1996 is available in [23].

This paper is another contribution to understanding the effect that long-range dependent traffic can have on the performance of communication networks. Our goal here is to investigate the role of scheduling policies in controlling such effects. We carry out this investigation in the framework of a single server queuing model, described below.

We consider a single server queue with unit service rate fed by an arrival process of the following form: sessions arrive at the times of a Poisson process of rate λ, with each session lasting for an independent integer time $\tau \geqslant 1$, where $P(\tau = k) = p_k$ with $p_k \sim \alpha k^{-(1+\alpha)} L(k)$, where $1 < \alpha < 2$ and $L(\cdot)$ is a slowly varying function. Each session brings in work at unit rate while it is active. In general the duration of a session would not be known at the time it arrives, although, in the course of the analysis, we also consider the situation where this duration is known at the time of arrival. The arrival process of workload is long-range dependent. Assume that the stability condition $\lambda E[\tau] < 1$ holds.

By simple arguments we show that for *any* stationary nonpreemptive service policy at the queue, the stationary sojourn time of a typical session must stochastically dominate a regularly varying random variable having infinite mean; this is true even if the duration of a session is known at the time it arrives. Further, this is true whether the service policy is work conserving or not, and whether it is causal or not. This result is proved in theorem 1 of section 4. On the other hand, we show that there exist causal stationary preemptive policies, which do not need knowledge of the session durations at the time of arrival, for which the stationary sojourn time of a session is stochastically dominated by a regularly varying random variable having finite mean. This result is proved in theorem 3 of section 6.

The contrast between these two results indicates that scheduling policies can have a very significant influence on the extent to which long-range dependence influences the performance of communication networks. Our choice of the arrival model was made for technical convenience, and because there are strong arguments in favor of the relevance of arrival models of the Cox type [9] we consider, see [13,24]. It is natural to expect that similar results will hold for more general arrival processes, e.g., for a Cox type arrival process where the session durations are of the subexponential class, see [3,12] for more on such arrival processes. Studies on the role of scheduling policies at routers and switches in controlling the influence of long range dependence on network performance may require simulation or more complex analysis than what we have done for the single server queue, and are also likely to be of interest, for the reasons mentioned in the first paragraph of the introduction.

2. Formal description of the arrival model

Consider a Poisson process of rate λ on the real line $\mathbb{R}$. Let

$$\cdots < T_{-1} < T_0 \leqslant 0 < T_1 < T_2 < \cdots$$

denote the points of this process.

Let $\{\tau_n,\ n \in \mathbb{Z}\}$ be independent and identically distributed (i.i.d.) positive integer valued random variables. Denoting $P(\tau = k)$ by p_k, we assume that $p_k \sim \alpha k^{-(1+\alpha)} L(k)$, where $1 < \alpha < 2$. Denoting $P(\tau \geqslant k)$ by q_k, $k \geqslant 1$, so that $q_k = \sum_{j=k}^{\infty} p_j$, by using [6, proposition 1.5.11(ii)], we then have $q_k \sim k^{-\alpha} L(k)$, i.e., the session duration is regularly varying. For the basic facts about regularly varying probability distributions, a good source is the book by Bingham et al. [6].

At each time T_n of the Poisson process, a session of duration τ_n starts. This session generates work at rate 1 throughout its duration. Thus, if $A(t)$ denotes the rate of generation of work at time $t \in \mathbb{R}$, we may write

$$A(t) = \sum_{\{n:\ T_n \leqslant t\}} \mathbf{1}(T_n + \tau_n > t),$$

where $\mathbf{1}(B)$ denotes the indicator function of the event B, and the convention is chosen so as to make $(A(t),\ t \in \mathbb{R})$ right continuous with left limits. The process $(A(t),\ t \in \mathbb{R})$ is time stationary.

We consider a single server queue whose input is $(A(t),\ t \in \mathbb{R})$. The server has the ability to serve work at rate 1. To ensure stability of the queue, we therefore need $\lambda E[\tau] < 1$, which we assume. Note that $E[A(0)] = \lambda E[\tau] < 1$.

Let

$$r(u) \triangleq E\big[A(0)A(u)\big] - \big(\lambda E[\tau]\big)^2$$

denote the autocovariance function of $(A(t),\ t \in \mathbb{R})$. It is straightforward to verify that

$$\int_{-\infty}^{\infty} \big|r(u)\big| = \infty. \tag{1}$$

This computation is left to the reader. Following Cox [9], equation (1) says, by definition, that $(A(t),\ t \in \mathbb{R})$ is long-range dependent. The arrival model described above is very similar to the ones considered in several earlier papers, e.g., [1,2,12,13, 15,16,18,20], and the practical relevance of considering arrival models of this type has been argued in [13,24].

3. Some preliminaries

It is convenient to gather in one place some notation that will be used in the subsequent discussion. Recall that $p_k \triangleq P(\tau = k)$ and $q_k \triangleq P(\tau \geqslant k) = \sum_{l=k}^{\infty} p_l$. We write

$$a_k = \sum_{l=1}^{k} p_l, \qquad b_k = \sum_{l=1}^{k} l p_l, \qquad c_k = \sum_{l=1}^{k} l^2 p_l, \qquad c_k^{(m)} = \sum_{l=1}^{k} l^{1+m} p_l, \quad m \geqslant 1. \tag{2}$$

Since we assumed that $p_k \sim \alpha k^{-(1+\alpha)} L(k)$, where $1 < \alpha < 2$ and $L(\cdot)$ is slowly varying, we have

$$\begin{aligned} &a_k \uparrow 1, \qquad b_k \uparrow E[\tau], \qquad c_k \sim \frac{\alpha}{2-\alpha} k^{2-\alpha} L(k), \\ &c_k^{(m)} \sim \frac{\alpha}{1+m-\alpha} k^{1+m-\alpha} L(k), \quad m \geqslant 1, \end{aligned} \tag{3}$$

where, for the latter two statements, we use Karamata's theorem, see [6, proposition 1.5.11(i), p. 26].

The stationary residual time distribution associated with the distribution of the typical session duration τ is the distribution $(q_k/E[\tau],\ k \geqslant 1)$. Note that, since

$$\sum_{l \geqslant k} \frac{q_l}{E[\tau]} \sim \frac{1}{(\alpha-1)E[\tau]} k^{1-\alpha} L(k), \tag{4}$$

by using [6, proposition 1.5.11(ii)], this distribution is also regularly varying, but it has infinite mean.

Throughout the paper, the symbol c_0 will be reserved for a finite positive constant whose value may be different in different equations. In each case, a suitable value can be determined explicitly in terms of the parameters, if desired. Similarly, the notation $L_0(\cdot)$ will be reserved for a slowly varying function, which may be different in different equations. An explicit function can be determined in each case, if desired.

4. Nonpreemptive service

Consider the system described in section 1, where the service policy is an *arbitrary* nonpreemptive policy. We will only assume that this policy is stationary, i.e., at any decision instant, it makes the same decision if the information it has is the same, without reference to the current time. Nonpreemptiveness just means that once service is initiated on a session, it has to continue being served until all the work associated with it is complete. In theorem 1 below, we do not make any assumption about whether the policy is causal or not, whether it is work conserving or not, or whether there is workload lookahead or not, i.e., whether or not it is known, at the time a session arrives, what the length of the session is going to be.

Theorem 1. For any nonpreemptive stationary service policy the stationary sojourn time of a typical session stochastically dominates a regularly varying distribution of infinite mean.

Proof. Recall that

$$\cdots < T_{-1} < T_0 \leqslant 0 < T_1 < T_2 < \cdots$$

denotes the times at which sessions arrive into the system and that these are the times of a Poisson process of rate λ. Consider an arbitrary nonpreemptive stationary service policy. Let

$$\cdots < I_{-1} < I_0 \leqslant 0 < I_1 < I_2 < \cdots$$

denote the times at which sessions enter service under this policy. Let N_k be defined such that the session that enters service at time I_k entered the system at time T_{N_k}. Then we have, for all k,

$$T_{N_k} \leqslant I_k \quad \text{and} \quad I_k + \tau_{N_k} \leqslant I_{k+1}.$$

Consider the graph of the residual duration of the session in service, if any. This can be visualized as a collection of nonoverlapping right angled triangles, the kth triangle starts with a jump of height τ_{N_k} at time I_k and decreases at rate 1, vanishing at time $I_k + \tau_{N_k}$. From this, a simple, standard argument based on time averages yields that the stationary residual duration of the session in service is equal to 0 with probability $1 - \lambda E[\tau]$ and, with probability $\lambda E[\tau]$, equals the stationary residual time distribution of the typical session duration. In particular this is a regularly varying distribution with infinite mean.

Let S have the distribution of the stationary sojourn time of a typical session. Because Poisson arrivals see time averages (PASTA), see [4], [25], this distribution is the same as that of a marked session that, at time 0, enters the stationary system. By the nonpreemptive nature of the service discipline, this marked session will certainly have to wait for the completion of the session currently under service before it can even begin to get served. The conclusion is that the stationary sojourn time of a typical session stochastically dominates the stationary residual duration of the session under service, i.e.,

$$P(S > t) \geqslant c_0 t^{1-\alpha} L(t)$$

for all sufficiently large t, see equation (4). This completes the proof of the theorem.

Note that all we used was the stationarity and the nonpreemptiveness of the service discipline. In particular, no assumption was made about whether the policy used is causal or not, whether it is work conserving or not, or whether there is workload lookahead or not. □

Remark. For the FCFS policy, which is stationary and nonpreemptive, the system essentially becomes an M/G/1 queue with arrival rate λ and service distribution the same as the distribution of the session duration τ, the only difference being that the work associated with an arrival is assumed to arrive at rate 1 after its arrival time, instead of all at once. With some abuse of notation we will continue to call such a system an M/G/1 queue in the sequel. Cohen [8] has studied this system and determined that

$$P(S > t) \sim \frac{\lambda}{(\alpha - 1)(1 - \lambda E[\tau])} t^{1-\alpha} L(t)$$

in this case.

5. Preemptive service with workload lookahead

Consider the system described in section 1, but allow the service policy to be preemptive, i.e., service on a session can be interrupted, whenever desired, to begin or continue work on another session. We will again only be interested in stationary policies. With preemptive service, it does not make sense to consider policies that are not work conserving. Therefore, in both this section and the next, we consider only work conserving policies, although a policy that is not a work conserving policy will be used as a thought experiment in the next section. In this section, we assume that the duration of a session is known at the time that it arrives (workload lookahead) – this assumption will be removed in the next section. In theorem 2 below, we will argue that there is a causal, stationary, work conserving, preemptive service policy for which the sojourn time distribution of a typical session is stochastically dominated by a regularly varying distribution of finite mean. This is in marked contrast to the result of theorem 1 of the preceding section. Of course, for now, we are assuming workload lookahead.

We may consider the arrival process as the sum of a countable number of independent arrival processes. The component arrival process of class k, $k \geqslant 1$, is the one corresponding to the sessions that have duration k – in this process sessions arrive at the times of a Poisson process of rate λp_k, and each session has duration k, bringing in work at rate 1 while it is active. The preemptive policy we will analyze in this section is very simple: sessions of lower class will have preemptive priority over sessions of higher class. Within each class, the sessions will be worked on in FCFS order. To be able to causally implement this policy, we need workload lookahead.

We remark also that, by PASTA, the stationary sojourn time of a typical session has the same distribution as the sojourn time of a marked session, arriving at time 0, with the system in stationarity. Let $\underline{K}$ denote the class of this marked arrival, i.e., the duration of the marked session. Then $\underline{K}$ is a random variable, with $P(\underline{K} = k) = p_k$, and it is independent of all the information associated with the stationary queueing process into which it arrives, up to the time just prior to its arrival.

5.1. Residual work just prior to the arrival of the marked session

Just prior to the arrival of the marked session at time 0, the system in general contains several partially completed sessions (at most one of each class), as well as sessions of various classes which have already arrived, but on which work has not yet commenced. We define a sequence of random variables

$$Z^{(1)} \leqslant Z^{(2)} \leqslant \cdots \leqslant Z^{(k)} \leqslant \cdots \leqslant Z \tag{5}$$

as follows: $Z^{(k)}$ is the total amount of work corresponding to all sessions of class k or less, partially worked on or not, that it is known, just prior to the arrival of the marked session at time 0, that the server will have to handle. Put in another way, $Z^{(k)}$ is the first time after time 0 that the server, working at rate 1, with the preemptive policy just described, without the marked session arriving, and with no sessions arriving after time 0, will finally exhaust all sessions of class k or less. In equation (5), Z is defined as $\lim_{k\to\infty} Z^{(k)}$.

In appendix A we investigate in more detail the distribution of the random variables defined in equation (5).

5.2. A sequence of M/G/1 queues

We next consider, for each $k \geqslant 1$, a FCFS M/G/1 queue, started empty at time 0, whose arrival process is a Poisson process of rate λa_k (see equation (2) for the notation) and where each arrival can be thought of as bringing in work at rate 1 for a duration l, $1 \leqslant l \leqslant k$, with respective probability p_l/a_k. We call this queue the M/G/1 queue of type k. Naturally, in so far as workload, one may identify the evolution of this queue with the evolution of the workload of sessions of class k or less in the original system, after time 0, assuming that all work corresponding to the marked session and corresponding to sessions that arrived before time 0 is instantaneously erased at time 0. However, we find it convenient not to worry about this identification, and, in particular, it is not necessary to discuss the relationship between the M/G/1 queues of type k for different k. This is because, in all our calculations, we are going to condition on the class of the marked session.

The state of the server in the M/G/1 queue of type k consists of alternating between idle and busy periods, starting with an idle period. Let $I_1^{(k)}, I_2^{(k)}, \ldots$ denote the successive idle periods, and $B_1^{(k)}, B_2^{(k)}, \ldots$ the successive busy periods. The idle periods are independent exponentially distributed random variables of mean $1/(\lambda a_k)$, while the busy periods are i.i.d., and independent of the idle periods. The distribution of these busy periods is discussed in more detail in appendix B.

For each $t \geqslant 0$, we also define, for the M/G/1 queue of type k,

$$I^{(k)}(t) = \text{total idle time elapsed up to time } t, \tag{6}$$

$$B^{(k)}(t) = \text{total busy time elapsed up to time } t, \tag{7}$$

$$N^{(k)}(t) = \text{total \# of distinct idle periods completely elapsed up to time } t. \tag{8}$$

Note that these are functions of the earlier defined variables $((I_n^{(k)}, B_n^{(k)}),\ n \geqslant 1)$.

5.3. *Stochastically dominating the sojourn time*

The result to be proved in this section is the following:

Theorem 2. Consider the preemptive work conserving policy with workload lookahead that gives preemptive priority to sessions of shorter duration, with sessions of the same duration being served in FCFS order. The stationary sojourn time of a typical session in this system is stochastically dominated by a regularly varying distribution having finite mean.

Proof. As observed in the preliminary discussion, by PASTA, it is equivalent to discuss the sojourn time distribution of a marked customer that arrives at time 0 into the system in stationarity. Recall that $\underline{K}$ denoted the class of this marked session, and let S denote its sojourn time. Then, with notation as in the preliminary discussion, one has, for any $t > 0$,

$$P(S > t \mid \underline{K} = k) = P\big(Z^{(k)} + k > I^{(k-1)}(t)\big),$$

where, by definition, $I^{(0)}(t) = t$.

We may now write, for any $\gamma > 0$,

$$\begin{aligned} P(S > t) &= \sum_{k=1}^{\infty} P(\underline{K} = k) P(S > t \mid \underline{K} = k) \\ &\leqslant \sum_{k=1}^{\gamma t} P(\underline{K} = k) P\big(Z^{(k)} + k > I^{(k-1)}(t)\big) + P(\underline{K} > \gamma t). \end{aligned} \tag{9}$$

The second term on the right hand side of equation (9) is just $q_{\gamma t}$. Thus we have

$$P(\underline{K} > \gamma t) \sim (\gamma t)^{-\alpha} L(\gamma t). \tag{10}$$

We next observe that, for any $0 < A < 1$,

$$P\big(Z^{(k)} + k > I^{(k-1)}(t)\big) \leqslant P\big(Z^{(k)} + k > At\big) + P\big(I^{(k-1)}(t) < At\big). \tag{11}$$

Now, if we choose

$$3\gamma < A, \tag{12}$$

then, because $\nu_k \triangleq E[Z^{(k)}]$ is nondecreasing in k, and by equation (26) which says in particular that ν_k grows at a sublinear rate (see also [6, proposition 1.3.6(v), p. 16]), we have

$$\nu_k + k < 2\gamma t$$

for *all* $k \leqslant \gamma t$, for all t sufficiently large. Thus

$$\begin{aligned}\sum_{k=1}^{\gamma t} P(\underline{K} = k) P\bigl(Z^{(k)} + k > At\bigr) &\leqslant \sum_{k=1}^{\gamma t} P(\underline{K} = k) P\bigl(\widetilde{Z}^{(k)} + k > At\bigr) \\ &\leqslant \sum_{k=1}^{\gamma t} p_k P\bigl(\widetilde{Z}^{(k)} - \nu_k > \gamma t\bigr) \\ &\leqslant \sum_{k=1}^{\gamma t} p_k \frac{\eta_k}{\gamma^2 t^2}. \end{aligned} \tag{13}$$

In the first step $\widetilde{Z}^{(k)}$ is defined as in appendix A, and we recall that it stochastically dominates $Z^{(k)}$. The last step just uses Chebyshev's inequality.

By equation (27) and the running assumption about p_k, we have

$$p_k \eta_k \sim c_0 k^{2-2\alpha} \bigl(L(k)\bigr)^2.$$

We consider three cases:

- If $\frac{3}{2} < \alpha < 2$, we have $\sum_{k=1}^{\infty} p_k \eta_k < \infty$, so that the left hand side of equation (13) is asymptotically upper bounded by $c_0 t^{-2}$.
- If $1 < \alpha < \frac{3}{2}$, we have

  $$\sum_{k=1}^{\gamma t} p_k \eta_k \sim c_0 t^{3-2\alpha} \bigl(L(\gamma t)\bigr)^2,$$

  by appealing to [6, proposition 1.5.11(i), p. 26]. Now the left hand side of equation (13) is asymptotically upper bounded by $t^{1-2\alpha} L_0(t)$, for a slowly varying function $L_0(\cdot)$.
- If $\alpha = \frac{3}{2}$, observe that for any $\varepsilon > 0$ we have $\sum_{k=1}^{\gamma t} p_k \eta_k$ asymptotically upper bounded by $c_0 t^{\varepsilon}$, by appealing to [6, proposition 1.5.11(i), p. 26, and proposition 1.3.6(v), p. 16], so that the left hand side of equation (13) is now asymptotically upper bounded by $c_0 t^{-2+\varepsilon}$.

In the first case observe that $-2 < -\alpha$. In the second case, observe that $1 - 2\alpha < -\alpha$. In the last case, for ε sufficiently small, again $-2 + \varepsilon < -\alpha$. Thus we have, for the entire range $1 < \alpha < 2$,

$$\sum_{k=1}^{\gamma t} P(\underline{K} = k) P\bigl(Z^{(k)} + k > At\bigr) \leqslant t^{-\alpha} \tag{14}$$

for all t sufficiently large.

Turning now to the second term on the right hand side of equation (11), we may write, for any $\delta > 0$,

$$P\big(I^{(k-1)}(t) < At\big) = P\big(I^{(k-1)}(t) < At,\ N^{(k-1)}(t) \geqslant \delta t\big) + P\big(I^{(k-1)}(t) < At,\ N^{(k-1)}(t) < \delta t\big). \tag{15}$$

The first term on the right hand side of equation (15) can be upper bounded as

$$P\big(I^{(k-1)}(t) < At,\ N^{(k-1)}(t) \geqslant \delta t\big) \leqslant P(E_1 + \cdots + E_{\delta t} < At),$$

where $E_1, \ldots, E_{\delta t}$ are independent exponential random variables of mean $1/\lambda$. If we choose

$$\delta > 2\lambda A, \tag{16}$$

then, using standard Chernoff estimates, we can find $\kappa > 0$ such

$$P\big(I^{(k-1)}(t) < At,\ N^{(k-1)}(t) \geqslant \delta t\big) \leqslant \exp(-\kappa t),$$

uniformly over *all* k, for all t sufficiently large, from which it also follows that

$$\sum_{k=1}^{\gamma t} P(\underline{K} = k) P\big(I^{(k-1)}(t) < At,\ N^{(k-1)}(t) \geqslant \delta t\big) \leqslant \exp(-\kappa t) \tag{17}$$

for all t sufficiently large.

We finally write

$$P\big(I^{(k-1)}(t) < At,\ N^{(k-1)}(t) < \delta t\big) \leqslant P\left(\sum_{n=1}^{\delta t} B_n^{(k-1)} > (1-A)t\right).$$

Recall the estimate in equation (29) for $\mu_{k-1} \triangleq E[B_1^{(k-1)}]$. If we now choose

$$\delta < \frac{(1-A)(1-\lambda E[\tau])}{2E[\tau]}, \tag{18}$$

then we may write

$$\begin{aligned}
&\sum_{k=1}^{\gamma t} P(\underline{K} = k) P\left(\sum_{n=1}^{\delta t} B_n^{(k-1)} > (1-A)t\right) \\
&\quad \leqslant \sum_{k=1}^{\gamma t} p_k P\left(\sum_{n=1}^{\delta t} B_n^{(k-1)} - \delta t \mu_{k-1} > \frac{(1-A)}{2} t\right) \\
&\quad \leqslant \sum_{k=1}^{\gamma t} p_k \frac{16}{(1-A)^4 t^4} \big(\delta t \xi_{k-1} + 3\delta t(\delta t - 1)(\psi_{k-1})^2\big)
\end{aligned}$$

$$\leqslant c_0 \left(\frac{1}{t^3} \sum_{k=1}^{\gamma t} p_k \xi_{k-1} + \frac{1}{t^2} \sum_{k=1}^{\gamma t} p_k (\psi_{k-1})^2 \right). \tag{19}$$

The second step here uses Markov's inequality for the centered fourth moment.

By equation (31) and the running assumption about p_k, we have $p_k \xi_{k-1} \sim c_0 k^{3-2\alpha}(L(k))^2$. Appealing to [6, proposition 1.5.11(i), p. 26], we have

$$\sum_{k=1}^{\gamma t} p_k \xi_{k-1} \sim c_0 t^{4-2\alpha} \big(L(\gamma t)\big)^2.$$

Thus, $(1/t^3) \sum_{k=1}^{\gamma t} p_k \xi_{k-1}$ is asymptotically upper bounded by $t^{1-2\alpha} L_0(t)$, where $L_0(\cdot)$ is a slowly varying function. We observe that $1 - 2\alpha < -\alpha$. Hence

$$\frac{1}{t^3} \sum_{k=1}^{\gamma t} p_k \xi_{k-1} \leqslant t^{-\alpha} \tag{20}$$

for all sufficiently large t.

By equation (30) and the running assumption about p_k, we have $p_k(\psi_{k-1})^2 \sim c_0 k^{3-3\alpha}(L(k))^3$. We consider three cases:

- If $\frac{4}{3} < \alpha < 2$, we have $\sum_{k=1}^{\infty} p_k(\psi_{k-1})^2 < \infty$, so that $(1/t^2) \sum_{k=1}^{\gamma t} p_k(\psi_{k-1})^2$ is asymptotically upper bounded by $c_0 t^{-2}$.
- If $1 < \alpha < \frac{4}{3}$, we have

  $$\sum_{k=1}^{\gamma t} p_k (\psi_{k-1})^2 \sim c_0 t^{4-3\alpha} \big(L(\gamma t)\big)^3,$$

  by appealing to [6, proposition 1.5.11(i), p. 26]. Now $(1/t^2) \sum_{k=1}^{\gamma t} p_k(\psi_{k-1})^2$ is asymptotically upper bounded by $t^{2-3\alpha} L_0(t)$, for a slowly varying function $L_0(\cdot)$.
- If $\alpha = \frac{4}{3}$, observe that for any $\varepsilon > 0$ we have $\sum_{k=1}^{\gamma t} p_k(\psi_{k-1})^2$ asymptotically upper bounded by $c_0 t^{\varepsilon}$, by appealing to [6, proposition 1.5.11(i), p. 26, and proposition 1.3.6(v), p. 16], so that $(1/t^2) \sum_{k=1}^{\gamma t} p_k(\psi_{k-1})^2$ is now asymptotically upper bounded by $c_0 t^{-2+\varepsilon}$.

In the first case observe that $-2 < -\alpha$. In the second case, observe that $2 - 3\alpha < -\alpha$. In the last case, for ε sufficiently small, again $-2 + \varepsilon < -\alpha$. Thus we have, for the entire range $1 < \alpha < 2$,

$$\frac{1}{t^2} \sum_{k=1}^{\gamma t} p_k (\psi_{k-1})^2 \leqslant t^{-\alpha} \tag{21}$$

for all sufficiently large t.

Putting together equations (20) and (21), we may continue with equation (19) and conclude that

$$\sum_{k=1}^{\gamma t} P(\underline{K} = k) P\left(\sum_{n=1}^{\delta t} B_n^{(k-1)} > (1-A)t\right) \leqslant t^{-\alpha} \tag{22}$$

for all sufficiently large t.

It is not hard to see that the constant A can be chosen sufficiently small that both γ and δ can be chosen to meet each of the requirements specified in equations (12), (16), and (18). Thus, from equations (9), (10), (14), (17), and (22), we have shown that

$$P(S > t) \leqslant c_0 t^{-\alpha} L(t) \tag{23}$$

for all t sufficiently large. This completes the proof of the theorem. □

Remark. The sojourn time of a typical session must necessarily be at least as large as its duration. Hence for any $\varepsilon > 0$,

$$P(S > t) \geqslant (1-\varepsilon) t^{-\alpha} L(t)$$

for all t sufficiently large.

6. Preemptive service without workload lookahead

In this section, we assume that at the time of arrival of a session its duration is not known. The main result to be proved is the following:

Theorem 3. There are causal preemptive work conserving policies without workload lookahead for which the stationary sojourn time of a typical session in the system under consideration is stochastically dominated by a regularly varying distribution having finite mean.

Proof. The key observation is that if we define a new process whose points are the times at which the sessions complete bringing in work, then this process is stochastically identical to the original arrival process, but can now be considered as having workload lookahead. To see this, recall that the original arrival process can be considered as the sum of a countable number of independent arrival processes: the component arrival process of class k, $k \geqslant 1$, is the one corresponding to the sessions that have duration k. In the kth component process sessions arrive at the times of a Poisson process of rate λp_k, so if we shift the times of this process by the fixed shift k, we once again have a Poisson process of the same rate. Summing these independent shifted processes now gives a process that is stochastically identical to the original arrival process.

We may now consider the server working with the policy of the preceding section, but taking for the arrival process not the original one, but the modified one defined in the preceding paragraph. This policy, which we denote by π, is of course not work conserving, but we note that the sojourn time of a typical session under this policy is just the sum of its duration and the time its spends in the system from the time it stops bringing in work to the time it actually departs (this can be interpreted as a "sojourn time" relative to the points of the modified arrival process). In terms of the random variables used in the proof of the preceding section, the distribution of the sojourn time of a typical session is therefore just that of $\underline{K} + S$. Note that $\underline{K}$ and S are dependent random variables. However, we have

$$P(\underline{K} + S > t) \leqslant P\left(\underline{K} > \frac{t}{2}\right) + P\left(S > \frac{t}{2}\right) \leqslant c_0 t^{-\alpha} L(t), \tag{24}$$

by equations (10) and (23).

Now π is not work conserving. However, we can choose one of any number of work conserving policies all of which release each session no later than π would. Basically, during the times that π would idle while there is still work in the system, the server may choose to work on any of the sessions present in the system. If π would require work on a particular session and that session is not present (because it left earlier) the server may again choose to work on any session it pleases among those present. All that is necessary is that when π requires work on a particular session and that session is present, the server should work on it. If the choices of which session to work on are made in a time stationary way, the resulting policy is causal, stationary, work conserving, and does not require workload lookahead. If $\widehat{S}$ denotes the stationary sojourn time of a typical session under this policy, then equation (24) implies that

$$P\big(\widehat{S} > t\big) \leqslant c_0 t^{-\alpha} L(t) \tag{25}$$

for all sufficiently large t. This completes the proof of the theorem. □

Remark. For the same reasons as in the remark after the proof of theorem 2, we have, for any $\varepsilon > 0$,

$$P\big(\widehat{S} > t\big) > (1 - \varepsilon) t^{-\alpha} L(t)$$

for all t sufficiently large.

7. Concluding remarks

We considered a class of long-range dependent arrival processes of the Cox type [9] at a single server queue, and considered the role of scheduling policies in influencing the performance seen by a typical session. The motivation for our work comes from empirical studies which demonstrate that the traffic that communication networks have to carry exhibits features that suggest the presence of long-range dependence. We chose a Cox type arrival process because of its analytical tractability and because it has been argued to be a good match for the mechanisms that generate long-range dependence in practice, see [13,24]. Within the framework of our model, we have demonstrated that the scheduling policy can have a dramatic influence on the performance seen by typical sessions. Similar analyses and simulation based studies of the role of routing and switching policies in alleviating the problems due to long-range dependent traffic in communication networks are likely to be of considerable interest.

Appendix A

We investigate in more detail the random variables $Z^{(k)}$ defined in equation (5). More precisely, for each $k \geqslant 1$, we will identify a random variable $\widetilde{Z}^{(k)}$ which stochastically dominates $Z^{(k)}$, and estimate the mean, $\nu_k \triangleq E[\widetilde{Z}^{(k)}]$ and the variance $\eta_k^2 \triangleq \mathrm{Var}(\widetilde{Z}^{(k)})$. These estimates will be useful in handling the expression $P(Z^{(k)} + k > At)$, which appears in equation (13).

Because our policy gives preemptive priority to sessions of smaller class and is work conserving, $Z^{(k)}$ has the same distribution as the stationary workload in the M/G/1 queue of type k described in section 5.2. By PASTA, applied now to this M/G/1 queue of type k, this distribution is the same as the distribution of the workload seen on arrival by a typical arrival to this queue. But the workload seen by a typical arrival to this queue can be expressed as the maximum of a random walk with negative drift, using the familiar Lindley equation for the workload in a single server FCFS queue [4,22]. Since we find it convenient to refer to certain formulas derived in [11, pp. 408–412] we will instead express the workload as the minimum of a random walk with positive drift – this just involves a change of sign in Lindley's equation.

Consider a random walk $(S_n,\ n \geqslant 0)$ with $S_0 = 0$, and with step size having distribution $F = A * B$, where $*$ denotes convolution, A is a distribution supported on the nonnegative half line $[0, \infty)$, and B is a distribution supported on the nonpositive half line $(-\infty, 0]$, and assume that the step size has positive mean. We are following the notation in [11, pp. 408–412]. Specializing to the situation of interest to us requires taking A to be the exponential distribution of mean $1/(\lambda a_k)$, and B to be the discrete distribution supported on the points $-j,\ 1 \leqslant j \leqslant k$, equalling $-j$ with probability p_j/a_k. Following the calculations in [11, pp. 408–412], the strict descending ladder height distribution for this random walk can be seen to be given by the defective density

$$g(y) = \begin{cases} \lambda \displaystyle\sum_{l=j}^{k} p_l & \text{if } -j < y \leqslant -j+1,\ 1 \leqslant j \leqslant k, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $g(\cdot)$ integrates to $\lambda b_k < 1$. $1 - \lambda b_k$ is the probability with which the strict descending ladder height random variable is undefined, i.e., the probability with which the random walk $(S_n,\ n \geqslant 0)$ is always nonnegative. The distribution of $Z^{(k)}$ is now immediately described: it is the sum of a geometric number of independent copies of random variables having the density

$$\widehat{g}(y) = \begin{cases} \dfrac{1}{b_k} \displaystyle\sum_{l=j}^{k} p_l & \text{if } j-1 \leqslant y < j,\ 1 \leqslant j \leqslant k, \\ 0 & \text{otherwise}, \end{cases}$$

where the mean number of such variables that one needs to sum is $\lambda b_k/(1-\lambda b_k)$ and the number of variables summed is independent of the variables themselves.

We find it convenient to observe that the distribution with density $\widehat{g}(\cdot)$ is stochastically dominated by the discrete distribution

$$\tilde{g}(y) = \sum_{j=1}^{k} \left(\frac{1}{b_k} \sum_{l=j}^{k} p_l \right) \delta(y-j),$$

and that therefore $Z^{(k)}$ is stochastically dominated by $\widetilde{Z}^{(k)}$, which is the sum of a geometric number of independent copies of random variables having the density $\tilde{g}(\cdot)$, where the mean number of such variables that one needs to sum is $\lambda b_k/(1-\lambda b_k)$ and the number of variables summed is independent of the variables themselves.

It is straightforward to verify that

$$\nu_k \triangleq E\big[\widetilde{Z}^{(k)}\big] = \frac{\lambda}{2(1-\lambda b_k)}(b_k + c_k) \sim c_0 k^{2-\alpha} L(k), \tag{26}$$

where we used equation (3). Further, one can check that

$$\eta_k \triangleq \mathrm{Var}\big(\widetilde{Z}^{(k)}\big) \sim c_0 k^{3-\alpha} L(k), \tag{27}$$

where we again used equation (3).

Appendix B

We study the busy period distribution of the M/G/1 queue of type k introduced in section 5.2, for each $k \geqslant 1$. We present estimates for the mean $\mu_k \triangleq E[B_1^{(k)}]$, the variance $\psi_k \triangleq E[(B_1^{(k)} - \mu_k)^2]$, and the centered fourth moment $\xi_k \triangleq E[(B_1^{(k)} - \mu_k)^4]$, of this distribution. These estimates will be used to handle the expression that shows up on the right hand side of the first line of equation (19).

Let

$$\beta_k(\theta) \triangleq E\big[\exp\big(-\theta B_1^{(k)}\big)\big]$$

denote the moment generating function (mgf) of this distribution. Let

$$F_k(\theta) \triangleq \frac{1}{a_k} \sum_{j=1}^{k} p_j \exp(-\theta j)$$

denote the mgf of the session duration of a typical session in this queue. It is straightforward to see that $\beta_k(\theta)$ and $F_k(\theta)$ are finite in an open interval around $\theta = 0$.

We have

$$\beta_k(\theta) = F_k\big(\theta + \lambda a_k - \lambda a_k \beta_k(\theta)\big). \tag{28}$$

This can be seen by a standard branching process argument: if $\tau^{(k)}$ denotes the duration of the session that starts the busy period, then

$$B_1^{(k)} \overset{\mathrm{d}}{=} \tau^{(k)} + \text{sum of Poisson } (\lambda a_k) \text{ independent copies of } B_1^{(k)}.$$

From equation (28) it is straightforward to verify that

$$\mu_k \triangleq E\big[B_1^{(k)}\big] = \frac{b_k}{a_k(1 - \lambda b_k)} \uparrow \frac{E[\tau]}{1 - \lambda E[\tau]}, \tag{29}$$

where we use equation (3) and the obvious fact that the mean busy period of the M/G/1 queue of type k is nondecreasing in k.

For the variance of $B_1^{(k)}$, we have

$$\psi_k \triangleq E\big[\big(B_1^{(k)} - \mu_k\big)^2\big] \sim c_0 k^{2-\alpha} L(k), \tag{30}$$

where we used equation (3).

We also find that

$$\xi_k \triangleq E\big[\big(B_1^{(k)} - \mu_k\big)^4\big] \sim c_0 k^{4-\alpha} L(k), \tag{31}$$

where we again used equation (3).

References

[1] V. Anantharam, On the sojourn time of sessions at an ATM buffer with long-range dependent input traffic, in: *Proceedings of the 34th IEEE Conference on Decision and Control*, Vol. 1, New Orleans, LA (December 13–15, 1995) pp. 859–864.

[2] V. Anantharam, Networks of queues with long-range dependent traffic streams, in: *Stochastic Networks: Stability and Rare Events*, eds. P. Glasserman, K. Sigman and D. Yao, Lecture Notes in Statistics, Vol. 117 (Springer, Berlin, 1996) pp. 237–256.

[3] S. Asmussen and J. Teugels, Convergence rates for M/G/1 queues and ruin problems with heavy tails, Report R-95-2005, Institute for Electronic Systems, Aalborg University (1995).

[4] F. Baccelli and P. Bremaud, *Palm Probabilities and Stationary Queues*, Lecture Notes in Statistics, Vol. 41 (Springer, Berlin, 1980).

[5] J. Beran, R. Sherman, M.S. Taqqu and W. Willinger, Long-range dependence in variable-bit-rate video traffic, IEEE Trans. Commun. 43(2–4) (February–April 1995) 1566–1579.

[6] N.H. Bingham, C.M. Goldie and J.L. Teugels, *Regular Variation* (Cambridge University Press, New York, 1987).
[7] C.S. Chang, Stability, queue length, and delay of deterministic and stochastic networks, IEEE Trans. Automatic Control 39(5) (1994) 913–931.
[8] J. Cohen, On the tail of the stationary waiting time distribution and limit theorems for the M/G/1 queue, Ann. Inst. H. Poincaré Section B 8(3) (1972) 255–263.
[9] D.R. Cox, Long-range dependence: A review, in: *Statistics: An Appraisal*, eds. H.A. David and H.T. David (Iowa State University Press, 1984) pp. 55–74.
[10] A. Erramilli, O. Narayan and W. Willinger, Experimental queueing analysis with long-range dependent packet traffic, IEEE/ACM Trans. Networking (April 1996) 4(2) 209–223.
[11] W. Feller, *An Introduction to Probability Theory and Its Applications*, Vol. 2 (Wiley, New York, 1971).
[12] P.R. Jelenkovic and A.A. Lazar, Asymptotic results for multiplexing subexponential on–off sources, Adv. Appl. Probab., submitted.
[13] T. Konstantopoulos and S.-J. Lin, Fractional Brownian motion and Lévy motions as limits of stochastic traffic models, in: *Proceedings of the 34th Annual Allerton Conference on Communications, Control, and Computing*, Urbana, IL (1996) pp. 913–922.
[14] W.E. Leland, M.S. Taqqu, W. Willinger and D.V. Wilson, On the self-similar nature of Ethernet traffic (extended version), IEEE/ACM Trans. Networking 2(1) (1994) 1–15.
[15] N. Likhanov, B. Tsybakov and N.D. Georganas, Analysis of an ATM buffer with self-similar ("fractal") input traffic, in: *Proceedings of the 14th Annual IEEE Infocom* (1995) pp. 985–992.
[16] Z. Liu, P. Nain, D. Towsley and Z.-L. Zhang, Asymptotic behavior of a multiplexer fed by a long-range dependent source, INRIA Report No. 3230 (August 1997). Available via `http://www.inria.fr/mistral/personnel/Philippe.Nain/research.html`.
[17] I. Norros, A storage model with self-similar input, Queueing Systems 16 (1994) 387–396.
[18] M. Parulekar and A.M. Makowski, Tail probabilities for a multiplexer with self-similar traffic, in: *Proceedings of the IEEE INFOCOM* (1996) pp. 1452–1459.
[19] V. Paxson and S. Floyd, Wide area traffic: The failure of Poisson modeling, IEEE/ACM Trans. Networking 3(3) (1995) 226–244.
[20] S. Vamvakos and V. Anantharam, On the departure process of a leaky bucket system with long-range dependent input traffic, Queueing Systems (to appear). Available as UCB/ERL Memorandum, No. M96/51, and UCB/ERL Memorandum, No. M96/70, Electronics Research Laboratory, University of California, Berkeley.
[21] B. Venkateshwara Rao, K.R. Krishnan and D.P. Heyman, Performance of finite-buffer queues under traffic with long-range dependence, in: *Proceedings of the IEEE GLOBECOM 1996*, Vol. 1, London, UK (1996) pp. 607–611.
[22] J. Walrand, *An Introduction to Queueing Networks* (Prentice-Hall, Englewood Cliffs, NJ, 1988).
[23] W. Willinger, M.S. Taqqu and A. Erramilli, A bibliographical guide to self-similar traffic and performance modeling for modern high-speed networks, in: *Stochastic Networks Theory and Applications*, eds. F. Kelly, S. Zachary and I. Ziedins (Oxford University Press, Oxford, 1996) 339–366.
[24] W. Willinger, M.S. Taqqu, R. Sherman and D.V. Wilson, Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level, IEEE/ACM Trans. Networking 5(1) (February 1997) 71–86.
[25] R.W. Wolff, Poisson arrivals see time averages, Oper. Res. 30 (1982) 223–231.