A Poisson Limit for the Unfinished Work of Superposed Point Processes

Jin Cao and Kavita Ramanan

Abstract

A key criterion in the design of high-speed networks is the probability of the buffer content exceeding a certain threshold. We consider n independent identical traffic sources modelled as point processes, which are fed into a link with speed proportional to n. Under fairly general assumptions on the input processes we show that the steady state unfinished work exceeding a threshold $b > 0$ tends to the corresponding probability assuming Poisson input processes. We verify the assumptions for a broad range of long-range dependent sources commonly used to model data traffic. Simulations show that for realistic values of the exceedance probability and moderate utilisations, convergence takes place at reasonable values of the number of sources superposed. In particular, our results indicate that even in the presence of long-range dependent traffic sources, with superposition the buffer exceedance probability has an exponential tail for even smaller buffers than suggested by previous results, which consider $O(n)$ buffer size. This is particularly relevant for high-speed networks in which the cost of high-speed memory is significant.

Keywords. Long-range dependence, buffer overflow probability, Poisson limit theorem, insensitivity, high-speed data networks, multiplexing gains, smoothing, superposition of sources, exponential decay rate, point process, Weibull, Pareto, on-off sources, network design.

1 Introduction

Empirical studies of high-speed data networks (e.g. Ethernet LAN's [18] and wide-area networks [23]) have shown that network traffic displays long-range dependence or self-similarity, thus bringing into question the use of traditional traffic source models such as on-off processes with exponential on-times, which are short-range dependent. Since on-off processes with heavy-tailed on-time distributions have been shown to exhibit long-range dependent behaviour [29], they have been proposed as more realistic data network source models. The implications of this discovery on network performance has been examined by several authors (see [3] for a survey). In particular, in [17] it is shown that under the appropriate scaling the tail of the queue distribution associated with an on-off heavy-tailed source is sub-exponential, rather than exponential as would be predicted by an exponential source model. An important question that arises in this context is whether this discrepancy in queueing behaviour between short-range and long-range dependent traffic models persists even when a large number of sources are multiplexed together. Some authors contend that even high levels of aggregation would not mitigate the burstiness of long-range dependent traffic [18, 23]. On the other hand, others have shown that in the presence of multiplexing some smoothing does indeed take place [2, 8, 12, 24, 19, 22, 26]. For instance in the context of traffic engineering for ATM multiplexers of VBR video sources it is shown in [24] that long-term correlations do not have a significant impact on the cell loss rate when ATM buffers are of realistic dimensions. Moreover, in [2, 8, 12, 19, 22, 26] a system with n independent identical sources feeding into a link with processing rate $O(n)$ is considered and using a large deviations analysis (under various assumptions on the input processes) it is shown that the steady state probability of the unfinished work exceeding a threshold equal to bn has an exponential tail (rather than a polynomial one) in the asymptotic limit as n tends to infinity. For the class of semi-Markov modulated fluid arrival processes, e.g. on-off sources with general on-time and off-time distributions, [22] also analyses the asymptotic ($n \to \infty$) decay rate of the steady state probability of the unfinished work exceeding nb, in the limit as $b \to 0$, and shows that in this regime the decay rate depends on the on-time and off-time distributions only through its mean. Thus the above mentioned papers provide evidence of statistical multiplexing gains when both the threshold and link speed grow proportionally to the number of sources multiplexed.

With the continual increase in link speeds in modern communication networks, the cost of high-speed memory is likely to become a non-trivial factor in the design of data networks. For example, the current price of 100 ms of buffering in a 40 Gbs system is approximately \$1000, and a typical switch is likely to have many 100 ms buffers supported on different linecards [30]. Therefore it is natural to ask whether the buffer need actually grow proportionally to the link speed in order to realise multiplexing gain, or whether smaller buffers will suffice. In this paper we prove a strong insensitivity result that addresses this question. Our result is stronger than that in [22] and uses a different approach. Firstly, we consider a scaling limit in which although the processing rate is proportional to the number n of sources multiplexed, the buffer size remains $O(1)$. Given fairly general assumptions on the source arrival distribution we show that the exceedance probability tends to the same value as if each source arrival process were Poisson with the same mean. This result supports recent statistical analysis of Internet data [4] and indicates that buffers need not scale with the link speed in order to achieve the required multiplexing gains. Secondly, we consider the actual probability rather than just the logarithmic asymptote considered in [2, 8, 12, 22, 26]. Knowledge of the prefactor usually allows one to more accurately assess actual statistical multiplexing gains, and design admission control policies that are neither too conservative nor too aggressive (see, for example, [6, 13]). Furthermore, we model the packet arrival processes as point processes rather than fluid processes. This seems to better capture the behaviour of real packet traffic, especially on the time scales relevant for buffer overflow (see Section 5 for further discussion of this issue).

The paper is organised as follows. In Section 2 we describe the model and state our main result (Theorem 2.1). In Section 3.1 we state our assumptions on the arrival processes and derive some consequences. The assumptions are shown to be satisfied by a large class of network traffic models in Section 4.2. In Section 3.2 we present the proof of our main result, and in Section 3.3 we contrast the scaling used here with another commonly used scaling. In Section 4 we present some simulation results, and we conclude in Section 5 with a discussion of the implication of our results for network design.

We close this section with some notation used throughout the paper. We let Z and $I\!R$ denote the set of integers and reals respectively, and let Z_+ and $I\!R_+$ be the set of nonnegative integers and reals respectively. For $x \in I\!R$, we use $\lfloor x \rfloor$ to denote the greatest integer less than or equal to x, and let $a \wedge b$ denote the minimum of a and b. Given $E \subseteq I\!R$ we let $\mathcal{D}([0,\infty) : E)$ and $\mathcal{D}([0,T] : E)$ be the space of right continuous functions on $[0,\infty)$ and $[0,T]$, respectively, that take values in E. Given a point process A on $I\!R$, for conciseness we use A_t to denote $A([0,t])$ and note that $A_t \in \mathcal{D}([0,\infty) : Z_+)$. Recall that a function $f(t)$ is said to be $o(t)$ if and only if

$$\lim_{t \to 0} \frac{f(t)}{t} = 0.$$

2 Description of the Model and Main Result

Consider n identical independent sources that send packets to a server that has a deterministic processing rate $n\mu$. We model the arrival of packets from source i as a point process $A^{(i)}$, where $A_t^{(i)}$ represents the total number of packets emitted by source i in the interval $[0,t]$. We assume that each $A^{(i)}$ has the same distribution as a simple stationary point process A that satisfies $\mathrm{E}[A_t] = \lambda t$ for each $t > 0$ for some $\lambda \in (0,\mu)$. Note that $\lambda < \mu$ is the usual stability condition that ensures that the queues do not grow unboundedly. The superposed process $A^n \doteq \sum_{i=1}^n A^{(i)}$ represents the aggregate arrival of packets from all n sources. Packets that cannot be processed immediately are stored in a buffer (which is assumed to be infinite). Moreover, we assume that each packet has the same processing time, which we without loss of generality set to be one by an appropriate choice of units. The unfinished work in the system at time t is defined to be the amount of time required to complete the processing of all packets present in the system at time t. The steady state probability of the unfinished work exceeding a certain level b is used as a surrogate for the steady state buffer overflow probability in a buffer of size b. We let U_{A^n} denote the stationary unfinished work in the system when the number of sources is n, and the arrival process of each source is distributed according to A. It is well known [20] that under the stability assumption $\lambda < \mu$, U_{A^n} exists and is explicitly given by the relation

$$U_{A^n} = \sup_{t>0} \left[A_t^n - n\mu t \right]. \tag{1}$$

Let N^λ denote the stationary Poisson point process with rate λ, and let the unfinished work process corresponding to a stationary Poisson source and a processor with rate μ be denoted by U_{N^λ}. More precisely,

$$U_{N^\lambda} = \sup_{t>0} \left[N_t^\lambda - \mu t \right]. \tag{2}$$

We now state our main result, which holds under fairly general assumptions on the distribution of the arrival process A of each source that are stated precisely in Section 3.1.

Theorem 2.1 *Suppose N^λ is the stationary Poisson point process with intensity $\lambda \in (0, \infty)$ and suppose A satisfies Assumption 3.2. Let U_{A^n} and U_{N^λ} be the steady state unfinished work corresponding to A^n and N^λ, as defined in (1) and (2) respectively. Then for any $b > 0$*

$$\lim_{n\to\infty} \mathrm{P}(U_{A^n} > b) = \mathrm{P}(U_{N^\lambda} > b). \tag{3}$$

It is well known that the unfinished work U_{N^λ} for a Poisson process has a steady state distribution of the form

$$\mathrm{P}(U_{N^\lambda} > b) = 1 - (1-\rho)e^{\rho b} Q_{\lfloor b \rfloor}(b - \lfloor b \rfloor),$$

where $\rho = \lambda/\mu$ is the utilisation of the queue and for $n \in \mathbb{N}$, $Q_n(x)$ are polynomials that that can be calculated using numerically stable recursion relations (see, for example, [16, p. 391]). It is also easy to see from [16, p. 391] that the distribution of U_{N^λ} contains no atoms in $(0, \infty)$. When b is an integer, the expression above takes a simpler form [15, p. 243], given by

$$\mathrm{P}(U_{N^\lambda} > m) = 1 - (1-\rho) \sum_{k=1}^{m} e^{k\rho} (-1)^{m-k} \frac{(k\rho)^{m-k}}{(m-k)!} \qquad \text{for } m \in \mathbb{N}, \tag{4}$$

and

$$\mathrm{P}(U_{N^\lambda} > 0) = \rho.$$

Thus Theorem 2.1 shows that as the number of sources multiplexed increases, the exceedance probability of the unfinished work associated with sources having fairly general distributions approaches the corresponding probability assuming Poisson sources, which can be calculated explicitly.

3 The Poisson Limit Theorem

In Section 3.1 we state the main assumption on the source distribution (Assumption 3.2) and derive some important consequences. The assumption is shown to be satisfied by a large class of processes commonly used to model network traffic in Section 4.2. In Section 3.2 we prove the main result, Theorem 2.1, and in Section 3.3 we contrast the scaling used in this paper with another commonly used large deviations scaling.

3.1 Main Assumptions

We first introduce some definitions required to state the assumptions. Let $\Lambda(x,t)$ be the Legendre transform of the logarithmic moment generating function of A_t. In other words, for $x \in I\!\!R$ and $t \in [0,\infty)$, let

$$\Lambda(x,t) \;=\; \sup_{\theta \in I\!\!R} [\theta x - t^{-1} \log \mathrm{E}[e^{\theta A_t}]]. \tag{5}$$

It is easy to verify (see for example [11, Lemma 2.2.5]) that for $x \geq \mathrm{E}[A_t]$,

$$\Lambda(x,t) \;\doteq\; \sup_{\theta \in [0,\infty)} [\theta x - t^{-1} \log \mathrm{E}[e^{\theta A_t}]]. \tag{6}$$

For $x \in I\!\!R$, we will also find it convenient to define the quantities

$$\Lambda_1(x) \doteq \liminf_{t \to 0} \Lambda(x,t) \quad \text{and} \quad \Lambda_2(x) \doteq \liminf{}_{t \to \infty} \frac{t\Lambda(x,t)}{\log t}. \tag{7}$$

Recall the definition of a simple point process [10, Definition 3.3.II].

Definition 3.1 *The point process A is said to be simple if* $\mathrm{P}(A\{t\} = 0$ *or* 1 *for all* $t) = 1$.

Note that if the point process is defined on $I\!\!R$, according to our notational convention, $A(\{t\}) = A_t - A_{t-}$, where $A_{t-} \doteq \lim_{s \uparrow t} A_s$.

We now state our main assumptions on the process A.

Assumption 3.2 *A is a stationary point process satisfying the following three properties.*

1. *A is simple and there exists $\lambda < \mu$ such that* $\mathrm{E}[A_t] = \lambda t$ *for* $t \in [0,\infty)$.

2. *There exists $\tilde{\theta} > 0$ such that* $\mathrm{E}\left[e^{\tilde{\theta} A_t}\right] < \infty$ *for every* $t \in [0,\infty)$.

3. $\Lambda_2(\mu) > 0$.

The fact that A is simple reflects our modeling assumption that all packets are of the same size, and that multiple packets from a single source do not arrive simultaneously at the server. The latter assumption is quite realistic, while the former is applicable to ATM networks, where packet sizes are all of the same size. (Our framework allows for generalisations to the case of packets with variable size, but we do not consider them in this paper.) Along with the fact that A is stationary and has finite intensity λ, this also implies that the probability of multiple packets arriving within an arbitrary small interval of zero is exceedingly small (see Lemma 3.3). The condition $\lambda < \mu$ is the usual stability condition that ensures existence of a steady state unfinished work (and queue length) process. In Section 4.2 we show that Assumption 3.2 is satisfied by a large class of sources that have been proposed as generic models of network traffic.

In the next two lemmas we establish some consequences of Assumption 3.2 that will be used in the proof of the main result.

Lemma 3.3 *Suppose A satisfies Assumption 3.2(1). Then*

$$\mathrm{P}(A_t = 1) = \lambda t + o(t) \qquad \textit{and} \qquad \mathrm{P}(A_t \geq 2) = o(t). \tag{8}$$

Moreover, if in addition A satisfies Assumption 3.2(2), then there exists $\theta_0 > 0$ such that

$$\lim_{t \to 0} t^{-1} \log \mathrm{E}\left[e^{\theta A_t}\right] = -\lambda + \lambda e^{\theta}, \tag{9}$$

uniformly for $\theta \in [0, \theta_0]$.

Proof. The first property follows directly from the fact that A is a simple stationary process with finite intensity [10, Propositions 3.3.I, 3.3.IV and 3.3.V].

We now prove the second property. For $k = 0, 1, \ldots$, let $p_t(k)$ denote $\mathrm{P}(A_t = k)$. Let $k_{-1} \doteq 1$ and for $l \in Z_+$, define

$$k_l \doteq \inf\{k > k_{l-1} : p_t(k) > 0\}.$$

We assume that the sequence $\{k_l\}$ is infinite. The case when $\{k_l\}$ is finite (i.e. there exists a finite $K < \infty$ such that $p_t(k) = 0$ for all $k \geq K$) follows as a simple consequence. By

Assumption 3.2(2) there exists $\theta_0 > 0$ such that $\mathrm{E}[e^{\theta A_t}] < \infty$ for all $\theta \in [0, 2\theta_0]$. Since $\mathrm{E}[e^{2\theta_0 A_t}] = p_t(0) + p_t(1)e^{2\theta_0} + \sum_{l=0}^{\infty} p_t(k_l)e^{2k_l\theta_0}$, the ratio test for convergent series implies that

$$\limsup_{l \in Z_+} \left[\frac{p_t(k_{l+1})e^{2(k_{l+1}-k_l)\theta_0}}{p_t(k_l)} \right] \leq 1.$$

Thus there exist $Lin(0,\infty)$ and $\alpha \in (0,1)$ such that for every $\theta \in [0,\theta_0]$

$$\sup_{l=L,L+1,\ldots} \left[\frac{p_t(k_{l+1})e^{(k_{l+1}-k_l)\theta}}{p_t(k_l)} \right] < \alpha.$$

Let

$$M \doteq \sum_{l=0}^{L} \frac{p_t(k_l)e^{(k_l-k_0)\theta_0}}{p_t(k_0)},$$

and note that in particular

$$\frac{p_t(k_L)e^{(k_L-k_0)\theta}}{p_t(k_0)} \leq M.$$

Then for $\theta \in [0,\theta_0]$ we obtain

$$\begin{aligned}
\mathrm{E}\left[e^{\theta A_t}\right] &= p_t(0) + p_t(1)e^{\theta} + \sum_{l=0}^{\infty} p_t(k_l)e^{k_l\theta} \\
&= p_t(0) + p_t(1)e^{\theta} + p_t(k_0)e^{k_0\theta}\left[\sum_{l=0}^{L} \frac{p_t(k_l)e^{(k_l-k_0)\theta}}{p_t(k_0)} + \sum_{l=L}^{\infty} \frac{p_t(k_{l+1})}{p_t(k_0)}e^{(k_{l+1}-k_0)\theta}\right] \\
&\leq p_t(0) + p_t(1)e^{\theta} + p_t(k_0)e^{k_0\theta}\left[M + \frac{p_t(k_L)e^{(k_L-k_0)\theta}}{p_t(k_0)} \sum_{l=L}^{\infty} \frac{p_t(k_{l+1})}{p_t(k_L)}e^{(k_{l+1}-k_L)\theta}\right] \\
&\leq p_t(0) + p_t(1)e^{\theta} + Mp_t(k_0)e^{k_0\theta}\left[1 + \sum_{l=L}^{\infty} \alpha^{l+1-L}\right] \\
&\leq p_t(0) + p_t(1)e^{\theta} + Mp_t(k_0)e^{k_0\theta}(1-\alpha)^{-1}.
\end{aligned}$$

On the other hand, clearly

$$\mathrm{E}\left[e^{\theta A_t}\right] \geq p_t(0) + p_t(1)e^{\theta} \geq 0.$$

Observing that $k_0 \geq 2$ and combining the last two displays with (8), one concludes that

$$\lim_{t\to 0} t^{-1}\left[\mathrm{E}\left[e^{\theta A_t}\right] - 1\right] = -\lambda + \lambda e^{\theta}$$

and

$$\lim_{t\to 0} t^{-1}\left[\mathrm{E}\left[e^{\theta A_t}\right] - 1\right]^2 = 0,$$

where the convergence is uniform for $\theta \in [0, \theta_0]$. The last two displays, along with the fact that $x - x^2/2 \leq \log(1+x) \leq x$ for $x \geq 0$, yield

$$\lim_{t \to 0} t^{-1} \log \left(1 + \mathrm{E}\left[e^{\theta A_t}\right] - 1\right) = \lim_{t \to 0} t^{-1} \left[\mathrm{E}\left[e^{\theta A_t}\right] - 1\right] = -\lambda + \lambda e^{\theta},$$

uniformly for $\theta \in [0, \theta_0]$ and thus (9) is established. □

Lemma 3.4 *Suppose A satisfies Assumption 3.2. Then the functions Λ, Λ_1 and Λ_2 defined in (6) and (7) satisfy the following properties.*

1. *The functions $\Lambda(\cdot, t)$, $t \in [0, \infty)$, $\Lambda_1(\cdot)$ and $\Lambda_2(\cdot)$ are non-negative non-decreasing functions on $[\lambda, \infty)$.*

2. *For every $x \geq \mu$, $\Lambda_1(x) > 0$.*

3. *For any $m \in Z_+$, $m > 1$, $t \in [0, \infty)$ and $x \geq \lambda$, $\Lambda(x, t) \geq m^{-1} \Lambda(x, m^{-1} t)$.*

Proof. First note that for all x and t, $\Lambda(x, t) \geq 0$ (since $\theta x - t^{-1} \log \mathrm{E}[e^{\theta A_t}] = 0$ when $\theta = 0$). It is easy to verify that for every $t \in [0, \infty)$, $\Lambda(\cdot, t)$ is convex, $\Lambda(\lambda, t) = 0$ and $\Lambda(\cdot, t)$ is non-decreasing for $x \geq \lambda$ (see, for example, [11, Lemma 2.2.5]). It then automatically follows from the definition that $\Lambda_1(\cdot)$ and $\Lambda_2(\cdot)$ are also non-negative non-decreasing functions on $[\lambda, \infty)$. This establishes the first property.

Since Assumption 3.2 is satisfied, from Lemma 3.3 we know that there exists $\theta_0 > 0$ such that (9) is satisfied uniformly for $\theta \in [0, \theta_0]$. Choose $\theta_* \in (0, \min[\theta_0, \log(\mu/\lambda)])$. Then using (9) for $x \geq \mu$ we observe that

$$\begin{aligned}
\Lambda_1(x) &= \liminf_{t \to 0} \sup_{\theta \in [0, \infty)} \left[\theta x - t^{-1} \log \mathrm{E}\left[e^{\theta A_t}\right]\right] \\
&\geq \liminf_{t \to 0} \sup_{\theta \in [0, \theta_*]} \left[\theta x - t^{-1} \log \mathrm{E}\left[e^{\theta A_t}\right]\right] \\
&\geq \sup_{\theta \in [0, \theta_*]} \liminf_{t \to 0} \left[\theta x - t^{-1} \log \mathrm{E}\left[e^{\theta A_t}\right]\right] \\
&\geq \sup_{\theta \in [0, \theta_*]} \left[\theta x - \lambda(e^{\theta} - 1)\right] \\
&\geq \theta_* x - \lambda(e^{\theta_*} - 1) \\
&> 0,
\end{aligned}$$

where the last inequality holds because $\theta x - \lambda(e^{\theta} - 1)$ is zero for $\theta = 0$ and (for $x \geq \mu$) is a strictly increasing function of θ on $[0, \theta^*]$.

The third property follows because the stationarity of A and Holder's inequality show that for any $m > 1$, $m \in Z_+$, and $t \in [0, \infty)$,

$$\mathrm{E}\left[e^{\theta A_t}\right] = \mathrm{E}\left[e^{\theta \sum_{i=1}^{m}\left(A_{it/m} - A_{(i-1)t/m}\right)}\right] = \mathrm{E}\left[\Pi_{i=1}^{m} e^{\theta A_{t/m}}\right] \leq \mathrm{E}\left[e^{m\theta A_{t/m}}\right].$$

Therefore for $x \in I\!R$, $t \in [0, \infty)$ and $m \in Z_+$, $m > 1$, we obtain

$$\begin{aligned}
\Lambda(x,t) &= \sup_{\theta \in I\!R}\left[\theta x - t^{-1}\log \mathrm{E}\left[e^{\theta A_t}\right]\right] \\
&\geq \sup_{\theta \in I\!R}\left[\theta x - t^{-1}\log \mathrm{E}\left[e^{m\theta A_{t/m}}\right]\right] \\
&= m^{-1}\sup_{\theta' \in I\!R}\left[\theta' x - m t^{-1}\log \mathrm{E}\left[e^{\theta' A_{t/m}}\right]\right] \\
&= m^{-1}\Lambda(x, m^{-1}t),
\end{aligned}$$

which completes the proof. □

We let $\alpha(\cdot) \doteq \Lambda(\mu, \cdot)$ denote the Legendre transform of the logarithmic generating function of A evaluated at the service rate μ. In other words, since $\mu > \lambda$, using (6) we see that

$$\alpha(t) = \sup_{\theta \in [0,\infty)}\left[\mu\theta - t^{-1}\log \mathrm{E}\left[e^{\theta A_t}\right]\right]. \tag{10}$$

Corollary 3.5 *Suppose A satisfies Assumption 3.2 and let $\alpha(\cdot)$ be given by (10). Then there exist τ, δ_1, $\delta_2 > 0$ such that*

1) $\alpha(t) \geq \delta_1$ for $t \in [0, \tau]$.

2) $\dfrac{t\alpha(t)}{\log t} \geq \delta_2$ for $t > \tau$.

Proof. Applying Lemma 3.4(2) with $x = \mu$, it is clear that there exist $t_1 > 0$ and $\delta_1' > 0$ such that $\alpha(t) \geq \delta_1'$ for $t \in [0, t_1]$. In addition Assumption 3.2(3) guarantees the existence of $\tau < \infty$ and $\delta_2 > 0$ such that $t\alpha(t)/\log t \geq \delta_2$ for $t \in [\tau, \infty)$. Without loss of generality we assume $t_1 \leq \tau$. Define $t_0 \doteq t_1^2/(\tau + t_1)$. Then for any $t \in [t_1, \tau]$,

$$\frac{t}{t_1} \leq \frac{\tau}{t_1} \leq \left\lceil \frac{\tau}{t_1} \right\rceil \leq \frac{\tau + t_1}{t_1} = \frac{t_1}{t_0} \leq \frac{t}{t_0},$$

and so there exists $\tilde{t} \in [t_0, t_1]$ such that $m \doteq t/\tilde{t} \in Z_+$. Then from Lemma 3.4(3) for $t \in [t_1, \tau]$ we obtain

$$\alpha(t) \geq \tilde{t}\alpha(\tilde{t})/t \geq t_0\delta_1'/\tau > 0.$$

The corollary then follows by setting $\delta_1 \doteq t_0\delta_1'/\tau$. □

3.2 Proof of the Main Result

In this section we prove Theorem 2.1, which shows that as the number of sources increases (and the processing rate increases proportionally), for a large class of arrival processes the steady state probability of the unfinished work exceeding a *fixed* threshold converges to the corresponding probability assuming the arrival processes were Poisson with the same mean intensity. The main intuition behind the result is that under the given assumptions the most probable time scale of exceedance of the threshold shrinks with increase in the number n of sources multiplexed (see Theorem 3.7). The fact that the finite-dimensional distributions of the scaled arrival processes tend to a Poisson limit as $n \to \infty$ [10, Proposition 9.2.VI] then suggests the result in Theorem 2.1. The issue of relevant time scales is discussed in more detail in Sections 3.3 and 5. We first establish some preliminary results that are required for the proof of the theorem.

We introduce the scaled point process B^n defined by

$$B_t^n \doteq A_{t/n}^n \qquad \text{for } t \in [0, \infty). \tag{11}$$

Note that the steady state unfinished work for the aggregate process A^n with server rate $n\mu$ is the same as that for the scaled process B^n with server rate μ. In other words

$$U_{A^n} = \sup_{t\geq 0}\left[A_t^n - n\mu t\right] = \sup_{t\geq 0}\left[A_{t/n}^n - \mu t\right] = \sup_{t\geq 0}\left[B_t^n - \mu t\right]. \tag{12}$$

Lemma 3.6 *For $c, x \in I\!R$, $n \in Z_+$ and $t \in (0, \infty)$ such that $x + ct^{-1} \geq \lambda$, we have*

$$\mathrm{P}\left(B_t^n > xt + c\right) \leq \exp\left[-t\Lambda(x + ct^{-1}, n^{-1}t)\right].$$

Moreover if $\alpha(\cdot)$ is given by (10) and $c \geq 0$, then

$$\mathrm{P}\left(B_t^n > \mu t + c\right) \leq \exp\left[-t\alpha(n^{-1}t)\right].$$

Proof. For any $\theta \geq 0$,

$$
\begin{aligned}
\mathrm{P}(B_t^n > xt + c) &\leq \exp\left[-xt\theta - c\theta\right] \mathrm{E}\left[\exp[\theta B_t^n]\right] \\
&= \exp\left[-xt\theta - c\theta\right] \left(\mathrm{E}\left[\exp\left[\theta A_{t/n}\right]\right]\right)^n \\
&= \exp\left[-t\left((x + ct^{-1})\theta - nt^{-1} \log \mathrm{E}\left[\exp[\theta A_{t/n}]\right]\right)\right].
\end{aligned}
$$

Taking the infimum of the right hand side over $\theta \geq 0$, using (6) and the fact that $x + ct^{-1} \geq \lambda$ we deduce that

$$\mathrm{P}\left(B_t^n > xt + c\right) \leq \exp[-t\Lambda(x + ct^{-1}, n^{-1}t)].$$

The second statement follows from the first and the fact proved in Lemma 3.4 that $\Lambda(\cdot, t)$ is a non-decreasing function on $[\lambda, \infty)$. $\square$

Theorem 3.7 *Suppose A satisfies Assumption 3.2. Then for any $b_0, \epsilon > 0$, there exist $T = T[b_0, \epsilon] \in \mathbb{R}_+$ and $N = N[b_0, \epsilon] \in Z_+$, such that for any $b \in [b_0, \infty)$ and $n > N$*

$$\mathrm{P}\left(\sup_{t \in [T,\infty)} [B_t^n - \mu t] > b\right) < \epsilon.$$

Proof. Choose $\gamma \in (0, b_0)$, set $s \doteq \gamma\mu^{-1}$, and for $l \in Z_+$ let $t_l \doteq sl$, and $I_l \doteq [t_l, t_{l+1})$. Then Lemma 3.6 and the fact that $b - \gamma > 0$ and B_t^n is non-decreasing in t yields

$$
\begin{aligned}
\mathrm{P}\left(\sup_{t \in I_l} [B_t^n - \mu t] > b\right) &\leq \mathrm{P}(B_{t_{l+1}}^n > b + \mu t_l) \\
&= \mathrm{P}\left(B_{t_{l+1}}^n > b - \gamma + \mu t_{l+1}\right) \\
&\leq \exp[-t_{l+1}\alpha(n^{-1}t_{l+1})].
\end{aligned}
$$

Given $T \in (0, \infty)$ let $L \in Z_+$ be such that $T \in [t_L, t_{L+1})$. From the above display we infer that for some $\tau \in (0, \infty)$ and all $n > L/\tau$,

$$
\begin{aligned}
\mathrm{P}\left(\sup_{t \in [T,\infty)} [B_t^n - \mu t] > b\right) &\leq \mathrm{P}\left(\sup_{t \in [sL,\ s\lfloor n\tau \rfloor)} [B_t^n - \mu t] > b\right) + \mathrm{P}\left(\sup_{t \in [s\lfloor n\tau \rfloor,\ \infty)} [B_t^n - \mu t] > b\right) \\
&\leq \sum_{l=L+1}^{\lfloor n\tau \rfloor} \exp(-t_l\alpha(n^{-1}t_l)) + \sum_{l=\lfloor n\tau \rfloor + 1}^{\infty} \exp(-t_l\alpha(n^{-1}t_l)).
\end{aligned}
$$

By Corollary 3.5 there exist $\tau, \delta_1, \delta_2 > 0$ such that $\alpha(t) \geq \delta_1$ for $t \in [0, \tau]$, and $t\alpha(t)/\log t \geq \delta_2$ for $t \in [\tau, \infty)$. Therefore

$$\sum_{l=L+1}^{\lfloor n\tau \rfloor} \exp(-t_l \alpha(n^{-1} t_l)) \leq \sum_{l=L+1}^{\lfloor n\tau \rfloor} \exp(-t_l \delta_1) < \sum_{l=L+1}^{\infty} \exp(-t_l \delta_1),$$

which tends to zero as $L \to \infty$, and similarly

$$\begin{aligned} \sum_{l=\lfloor n\tau \rfloor+1}^{\infty} \exp(-t_l \alpha(n^{-1} t_l)) &\leq \sum_{l=\lfloor n\tau \rfloor+1}^{\infty} \exp(-n\delta_2 \log(n^{-1} t_l)) \\ &= \sum_{l=\lfloor n\tau \rfloor+1}^{\infty} \left(\frac{ls}{n}\right)^{-\delta_2 n} \\ &\leq s^{-\delta_2 n} \int_{\tau}^{\infty} x^{-\delta_2 n} dx, \end{aligned}$$

which tends to zero as $n \to \infty$. Thus we can choose $L = L[\gamma, b_0]$ and $N = N[\gamma, b_0]$ such that for $T = sL$ and $n > N$ each term in the last two displays is less than $\epsilon/2$, for $n > N$. This establishes the lemma with $T \doteq sL$. $\square$

We let $F : \mathcal{D}([0,\infty) : I\!R_+) \to \mathcal{D}([0,\infty) : I\!R_+)$ be the supremum map defined by

$$F(f)(T) = F_T(f) = \sup_{t \in [0,T]} [f_t - \mu t] \quad \text{for } T \in [0, \infty). \tag{13}$$

We assume that $\mathcal{D}([0,\infty) : I\!R_+)$ is endowed with the Skorokhod J_1 metric $d(\cdot,\cdot)$ defined below. For $T < \infty$ let $\mathcal{S}_T$ represent the class of strictly increasing continuous mappings of $[0, T]$ onto itself, and for $\sigma \in \mathcal{S}_T$ we introduce the norm

$$||\sigma||_T \doteq \sup_{s,t \in [0,T], s \neq t} \left| \log \left(\frac{\sigma(t) - \sigma(s)}{t - s} \right) \right|.$$

The Skorokhod J_1 metric on $\mathcal{D}([0,\infty) : I\!R_+)$ [1, p. 112], [28, p. 73] is given by

$$d(f, g) = \int_0^{\infty} [d_T(r_T(f), r_T(g)) \wedge 1] e^{-T} dT,$$

where $r_T : \mathcal{D}([0,\infty) : I\!R_+) \to \mathcal{D}([0,T] : I\!R_+)$ is simply the restriction of f to $[0,T]$, and for $f, g \in \mathcal{D}([0,T] : I\!R_+)$, $T \in [0,\infty)$, we define

$$d_T(f, g) \doteq \inf \left\{ \varepsilon > 0 : \text{ there exists } \sigma \in \mathcal{S}_T \text{ such that } ||\sigma||_T < \varepsilon \text{ and } \sup_{t \in [0,T]} |f(t) - g(\sigma(t))| < \varepsilon \right\}.$$

We now prove the main theorem.

Proof of Theorem 2.1. From the superposition theorem for point processes we know that B^n converges weakly to N^λ in $\mathcal{D}([0,\infty) : I\!R_+)$ [10, Lemmas 9.1.IV and 9.1.X and Proposition 9.2.VI]. From [28, Theorem 6.2] and the Skorokhod Representation Theorem (see for example [28, p.68]) it follows that $F(B^n) \Rightarrow F(N^\lambda)$ in $\mathcal{D}([0,\infty) : I\!R_+)$. Using the continuous mapping theorem and the P a.s. continuity of the projection operator $f \to f_T$ for T a.e. [1, Theorems 5.1 and 15.1], we conclude that for Lebesgue almost every $T \in [0,\infty)$ $F_T(B^n) \Rightarrow F_T(N^\lambda)$. Then by Portmanteau's theorem [1, Theorem 2.1] for any $b > 0$

$$\lim_{n\to\infty} \mathrm{P}(F_T(B^n) > b) = \mathrm{P}(F_T(N^\lambda) > b). \tag{14}$$

Now recall that for any $b > 0$,

$$\mathrm{P}(U_{A^n} > b) = \mathrm{P}\left(\sup_{t\in[0,\infty)} [A_t^n - n\mu t] > b\right) = \mathrm{P}\left(\sup_{t\in[0,\infty)} [B_t^n - \mu t] > b\right),$$

and

$$\mathrm{P}\left(U_{N^\lambda} > b\right) = P\left(\sup_{t\in[0,\infty)} \left[N_t^\lambda - \mu t\right] > b\right).$$

Thus for any $T \in (0,\infty)$ and $b > 0$ we have

$$\mathrm{P}\left(F_T(B^n) > b\right) \le \mathrm{P}(U_{A^n} > b) \le \mathrm{P}\left(F_T(B^n) > b\right) + \mathrm{P}\left(\sup_{t\in[T,\infty)} [B_t^n - \mu t] > b\right),$$

and likewise

$$\mathrm{P}\left(F_T(N^\lambda) > b\right) \le \mathrm{P}(U_{N^\lambda} > b) \le \mathrm{P}\left(F_T(N^\lambda) > b\right) + \mathrm{P}\left(\sup_{t\in[T,\infty)} [N_t^\lambda - \mu t] > b\right).$$

Theorem 3.7, along with the observation that N_t^λ has the same distribution as $\sum_{i=1}^n N_{t/n}^\lambda$ for any $n \in I\!N$, shows that given any $\varepsilon > 0$ for all large enough T and n the last terms in the above two displays are less than $\varepsilon/2$. Since one can always choose T such that (14) is also satisfied, we conclude that for all $b > 0$

$$\lim_{n\to\infty} \mathrm{P}\left(U_{A^n} > b\right) - \varepsilon \le \mathrm{P}\left(U_{N^\lambda} > b\right) \le \lim_{n\to\infty} \mathrm{P}\left(U_{A^n} > b\right) + \varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, sending $\varepsilon \downarrow 0$ yields the desired result. $\square$

3.3 Comparison of Different Scalings

Recall the definition (11) of the scaled arrival processes B^n and consider the sequence of exceedance probabilities $\mathrm{P}\left(\sup_{t\in[0,\infty)}[B_t^n - \mu t] > b_n\right)$ associated with a sequence of thresholds b_n, $n \in I\!N$. In this paper we consider the case when $b_n = b$, $n \in I\!N$. Another more commonly used scaling (see for instance [2, 19, 22]) is $b_n = nb$ for some $b > 0$. In this section we compare the mathematical analysis of the limiting probabilities in these two cases. The practical implications of these scalings is discussed in Section 5.

Let $p_n(b_n)$ be the steady state probability of the buffer exceeding b_n in the presence of n sources, i.e.

$$p_n(b_n) \doteq \mathrm{P}\left(\sup_{t\in[0,\infty)}[B_t^n - \mu t] > b_n\right).$$

Then $p_n(b_n)$ is clearly bounded by the sum of the following two probabilities:

$$\mathrm{P}\left(\sup_{t\in[0,T]}[B_t^n - \mu t] > b_n\right) \quad \text{and} \quad \mathrm{P}\left(\sup_{t\in[T,\infty)}[B_t^n - \mu t] > b_n\right). \tag{15}$$

Under the scaling $b_n = b$, the first term is the dominant term since, as shown in Theorem 3.7, the second term is arbitrarily small for sufficiently large T and n. On the other hand, we show below that under the scaling $b_n = nb$ the second term is dominant. Assume that there exists $t_0 > 0$ such that $\mathrm{E}[e^{\theta A_t}] < \infty$ for any $\theta > 0$ and $t \in [0, t_0]$. Using Holder's inequality this condition can be shown to imply Assumption 3.2 (2). For n large enough such that $nT^{-1}b > \lambda$, from Lemma 3.6 we obtain

$$\mathrm{P}\left(\sup_{t\in[0,T]}[B_t^n - \mu t] > nb\right) \le \mathrm{P}\left(B_T^n > nb\right) \le \exp[-T\Lambda(nT^{-1}b, n^{-1}T)].$$

Therefore we obtain

$$\begin{aligned} \limsup_{n\to\infty}\left[n^{-1}\log\mathrm{P}\left(\sup_{t\in[0,T]}[B_t^n - \mu t] > nb\right)\right] &\le \limsup_{n\to\infty}[-n^{-1}T\Lambda(nT^{-1}b, n^{-1}T)] \\ &\le -\liminf_{t\to 0}[t\Lambda(bt^{-1}, t)]. \end{aligned} \tag{16}$$

Substituting the definition (5) of $\Lambda(\cdot, t)$ we have

$$\liminf_{t\to 0} t\Lambda(bt^{-1}, t) = \liminf_{t\to 0}\sup_{\theta\in I\!R}[\theta b - \log\mathrm{E}[e^{\theta A_t}]] \ge \sup_{\theta\in I\!R}\liminf_{t\to 0}[\theta b - \log\mathrm{E}[e^{\theta A_t}]]. \tag{17}$$

Given our assumption that $\mathrm{E}[e^{\theta A_t}] < \infty$ for any $\theta > 0$ and $t \in [0, t_0]$, and the fact that A_t is right continuous with $A_0 = 0$ Lebesgue's bounded convergence theorem shows that

$$\limsup_{t\to 0}\log\mathrm{E}[e^{\theta A_t}] = 0.$$

Together with (16) and (17) this implies

$$\limsup_{n\to\infty}\left[n^{-1}\log \mathrm{P}\left(\sup_{t\in[0,T]}[B^n_t-\mu t]>nb\right)\right]\leq -\sup_{\theta\in I\!R}[\theta b]=-\infty. \tag{18}$$

However, previous results (e.g. [12, 8, 26], [22, Theorem 3.2][1]) show that under Assumption 3.2,

$$\lim_{n\to\infty} n^{-1}\log p_n(b_n) = -I(b) < \infty, \tag{19}$$

where

$$I(b) \doteq \inf_{t>0}\,[t\Lambda(\mu+bt^{-1})]$$

is the exponential decay rate of the exceedance probability. (Note that in contrast to the case $b_n = b$, here the asymptotic decay rate is a function of the arrival process distribution.) Since for any two series $\{a_n\}$ and $\{b_n\}$

$$\limsup_{n\to\infty}\log(a_n+b_n) = \max[\limsup_{n\to\infty}\log a_n, \limsup_{n\to\infty}\log b_n],$$

(15) together with (18) and (19) implies that for any $T > 0$

$$\lim_{n\to\infty} n^{-1}\log \mathrm{P}\left(\sup_{t\in[T,\infty)}[B^n_t-\mu t]>nb\right) = -I(b).$$

Thus we have shown that the second term in (15) is the dominant contribution to the exceedance probability under the scaling $b_n = nb$.

Our comparison of different scalings highlights the importance of a *critical time scale* [24] for performance analysis. The critical time scale is loosely defined as the time scale beyond which the source behaviour does not significantly affect the exceedance probability. For the buffer size scaling $b_n = b$ considered in this paper, the critical time scale is inversely proportional to the number n of sources superposed (see Theorem 3.7). Thus as n increases the characteristics of an individual source at the critical time scale become less important, the effect of statistical multiplexing becomes the dominant factor, and the net behaviour tends towards Poisson. However, with the scaling $b_n = nb$, the time scale of interest does not change with superposition. Thus individual source behaviour still plays an important role.

[1] Theorem 3.2 of [22] only considers semi-Markov modulated fluid arrival processes and so does not apply to point processes directly. However, it appears that the result could be generalised to apply to point processes.

4 Study of Convergence Rates via Simulation

In Theorem 2.1 we showed that for a large class of arrival processes, the associated buffer exceedance probability of the unfinished work process converges to a Poisson limit. In this section we simulate the unfinished work process for a variety of representative input source models in order to gauge the qualitative dependence of the rate of convergence on parameters such as the threshold level b and utilisation ρ of the link, as well as on the characteristics of the source model. In Section 4.1 we describe the five classes of models considered and in Section 4.2 we prove that Assumption 3.2 is satisfied by each class of models. We present and examine the simulation results in Section 4.3, and discuss the implications of these results on network design in Section 5.

4.1 Models of the Arrival Processes

We consider five classes of arrival processes that are commonly used as models of network traffic sources. The different classes exhibit key features of observed data network traffic behaviour such as long-range dependence, asymptotic self-similarity [18] and "burstiness" to different extents. Recall that a stochastic process $X(t)$ is said to be long-range dependent if there exists a constant $H \in (1/2, 1)$, referred to as the *Hurst parameter*, such that $\mathrm{Var}(X(t)) \sim O(t^{2H})$, while it is said to be self-similar if the finite-dimensional distributions of the time-changed and rescaled process $\theta^{-H}X(\theta t)$ are the same as the original process, i.e. $\theta^{-H}X(\theta\cdot) \stackrel{fidi}{=} X(\cdot)$ for $\theta > 0$. A commonly used measure for burstiness of traffic is the index of dispersion of counts (IDC) [9, 27]. Of special interest is the limiting value I_∞ of IDC associated with an arrival process A_t, which is defined by

$$I_\infty \doteq \lim_{t\to\infty} \frac{\mathrm{Var}[A_t]}{\mathrm{E}[A_t]}.$$

It is well known that $I_\infty = 1$ for a Poisson process and for renewal processes it is equal to the square of the coefficient of variation (COV) of the inter-arrival distribution (recall that the COV of a random variable is defined to be the ratio of its variance to the square of its mean). A process is considered more (less) bursty than Poisson if I_∞ is greater (smaller) than 1. I_∞ also has the property that it is invariant under the superposition of independent identical sources.

We now describe the five classes in detail. Sources more bursty than Poisson are considered in Classes 1)-4), and for completeness a model that is less bursty than Poisson is considered in Class 5 (also see Table 1). Classes 2) and 4) are long-range dependent and asymptotically self-similar and thus belong to the category of *fractal point processes* [25], which have $I_\infty = \infty$. Classes 1), 2) and 5) are renewal processes, and are therefore completely characterised by the distribution of their inter-arrival times. Classes 3) and 4) belong to the category of what we refer to as Poisson on-off point processes, which are comprised of an underlying on-off process and a conditional mean arrival rate. Packets arrive as a Poisson point process with the given conditional mean arrival rate during an on-period, while no packets arriving during an off-period. Thus Poisson on-off processes are completely defined by three characteristics — the on-period distribution, the off-period distribution and the mean Poisson arrival rate during the on-period.

Class 1 (Weibull): The arrival process is a renewal processes with the inter-arrival times distributed as *Weibull* random variables. Recall that a random variable Z is Weibull with shape parameter $\alpha \in (0, \infty)$ if Z^α is an exponential. The Weibull random variable is heavy-tailed if $\alpha \in (0, 1)$. The inter-arrival distribution of HTTP packet traffic has been empirically observed to be well modelled by a Weibull distribution [5].

Class 2 (InvWeibull): This source model is a renewal process with inter-arrival times distributed according to an *Inverse-Weibull* distribution, where Z is said to have an Inverse-Weibull distribution with shape parameter α if $Z^{-\alpha}$ is an exponential random variable. Recall that a random variable X is a *Pareto* with tail parameter α iff

$$\mathrm{P}(X > x) = Mx^{-\alpha}, \qquad \text{for } x > M^{\frac{1}{\alpha}}, \tag{20}$$

it can be easily shown that the Inverse-Weibull has an asymptotic Pareto tail distribution with tail parameter α, i.e. $\mathrm{P}(Z > z) = O(z^{-\alpha})$ as $z \to \infty$. For $\alpha \in (1, 2)$ the inter-arrival has a finite mean (equal to λ^{-1}) but an infinite variance, so that $I_\infty = \infty$. This process belongs to the class of fractal point process [25] with Hurst parameter $(3 - \alpha)/2$.

Class 3 (ExponExpoff): This is a Poisson on-off process where both the on-periods and off-periods are exponentially distributed, and the mean on-period arrival rate κ is finite.

The model can also be viewed as a renewal process with the inter-arrival time having a *hyper-exponential* distribution, i.e., a mixture of two exponential distributions. This model was used in [6] to study the queueing behaviour of superposed ATM traffic.

Class 4 (ParetoonExpoff): The arrival process is a Poisson on-off process with an exponential off-period distribution and the on-periods are *Pareto* with tail parameter α (20). In addition, it has a finite conditional mean arrival rate κ. If $\alpha \in (1,2)$, this process belongs to the class of fractal point process [25], it has $I_\infty = \infty$ and is long-range dependent with a Hurst parameter $(3-\alpha)/2$.

Class 5 (Gamma): The arrival process is a renewal processes with the inter-arrivals having a *Gamma* (or Erlang) distribution with shape parameter $\alpha \in I\!N$, i.e. it is defined to be the sum of α independent and identically distributed exponentials.

4.2 Verification of Assumption 3.2 for Network Traffic Models

Here we verify Assumption 3.2 for the five classes of models described in the last section.

Lemma 4.1 *Any renewal process A whose inter-arrival distribution X satisfies* $\mathrm{E}[X] = \lambda^{-1} > 0$ *has* $\mathrm{E}[A_t] = \lambda < \infty$ *and satisfies Assumptions 3.2(2) and 3.2(3).*

Proof. The fact that $\mathrm{E}[X] = \lambda^{-1} > 0$ implies automatically that for every $t \in [0,\infty)$, $\mathrm{E}[A_t] = \lambda t$.

Let $\tilde{L}$ be the time to the first arrival, and suppose $L_n, n \in I\!N$, is the time interval between the nth and $(n+1)$th arrivals. Since A_t is a renewal process, $\{\tilde{L}, L_n, n \in I\!N\}$ are independent and each $L_n, n \in I\!N$, has the same distribution as X. Thus for $n \in I\!N$,

$$\mathrm{P}(A_t \geq n) \quad = \quad \mathrm{P}(\tilde{L} + L_1 + L_2 + \ldots + L_{n-1} \leq t),$$

which implies that for any $\phi > 0$,

$$\mathrm{P}(A_t \geq n) \leq \mathrm{P}(L_1 + L_2 + \ldots + L_{n-1} \leq t) \leq e^{\phi t}\left(\mathrm{E}[e^{-\phi X}]\right)^{n-1}. \tag{21}$$

Since $\mathrm{E}[X] = \lambda^{-1} > 0$, X is positive on a set of positive probability, and hence $\mathrm{E}[e^{-\phi X}] \in (0,1)$ for any $\phi > 0$. Thus given any $\phi_0 > 0$, there exists $\tilde{\theta} > 0$ such that

$$c \doteq e^{\tilde{\theta}}\mathrm{E}[e^{-\phi_0 X}] < 1. \tag{22}$$

Combining (21) and (22), for $\theta < \tilde{\theta}$ and $t \in [0, \infty)$ we obtain

$$\mathrm{E}[e^{\theta A_t}] \;=\; \sum_{n=0}^{\infty} \mathrm{P}(A_t = n) e^{\theta n} \;\leq\; \sum_{n=0}^{\infty} \mathrm{P}(A_t \geq n) e^{\tilde{\theta} n} \;\leq\; e^{\phi_0 t} \left(\mathrm{E}[e^{-\phi_0 X}] \right)^{-1} \sum_{n=0}^{\infty} c^n \;<\; \infty,$$

which validates Assumption 3.2(2).

In order to prove Assumption 3.2(3) we use methods similar to those used in [19, Proposition 3.1]. Pick $\beta > 0$ and for $t \in (0, \infty)$ define $\theta_t \doteq \beta \log(t)/t$. Then note that

$$\Lambda_2(\mu) = \frac{t\Lambda(\mu, t)}{\log t} \geq \frac{t\theta_t \mu}{\log t} - \limsup_{t \to \infty} \frac{\log \mathrm{E}[e^{\theta_t A_t}]}{\log t} = \beta\mu - \limsup_{t \to \infty} \frac{\log \mathrm{E}[e^{\theta_t A_t}]}{\log t}. \tag{23}$$

Since $\theta_t \to 0$ as $t \to \infty$, by Assumption 3.2(2) $\mathrm{E}[e^{\theta_t A_t}]$ is finite for all sufficiently large t. Choose $C \in (\lambda, \mu)$, recall that $\lceil Ct \rceil$ is the least integer greater than or equal to Ct and note that

$$\begin{aligned} \mathrm{E}[e^{\theta_t A_t}] &= \textstyle\sum_{n=0}^{\infty} \mathrm{P}(A_t = n) e^{\theta_t n} \\ &= \textstyle\sum_{n=0}^{\lceil Ct \rceil - 1} \mathrm{P}(A_t = n) e^{\theta_t n} + \sum_{n=\lceil Ct \rceil}^{\infty} \mathrm{P}(A_t = n) e^{\theta_t n} \\ &\leq e^{\theta_t C t} + \textstyle\sum_{n=\lceil Ct \rceil}^{\infty} \mathrm{P}(A_t \geq n) e^{\theta_t n} \\ &= t^{C\beta} + \textstyle\sum_{n=\lceil Ct \rceil}^{\infty} \mathrm{P}(A_t \geq n) e^{\theta_t n}. \end{aligned} \tag{24}$$

For $\phi \in [0, \infty)$ define

$$g(\phi) \doteq e^{\phi C^{-1}} \mathrm{E}[e^{-\phi X}],$$

and note that g is differentiable with derivative

$$g'(\phi) = e^{\phi C^{-1}} \left(C^{-1} \mathrm{E}[e^{-\phi X}] - \mathrm{E}[X e^{-\phi X}] \right).$$

Therefore

$$g(0) = 1, \quad g'(0) = C^{-1} - \lambda^{-1} < 0,$$

and so there exists $\phi_0 > 0$ such that

$$g(\phi_0) < 1. \tag{25}$$

To bound the summation on the right side of (24), now let

$$d \doteq \mathrm{E}[e^{-\phi_0 X}] \;\text{ and }\; d_t \doteq e^{\theta_t} \mathrm{E}[e^{-\phi_0 X}],$$

and observe that as $t \to \infty$, $d_t \to d$ and $d < 1$. Using this observation along with (21), (25) and the fact that $\theta_t = \beta \log t / t$, we conclude that for all sufficiently large t,

$$\begin{array}{rclcl} \sum_{n=\lceil Ct \rceil}^{\infty} \mathrm{P}(A_t \geq n) e^{\theta_t n} & \leq & e^{\phi_0 t} \sum_{n=\lceil Ct \rceil}^{\infty} (\mathrm{E}[e^{-\phi_0 X}])^{n-1} e^{\theta_t n} & = & e^{\phi_0 t} d^{-1} (1-d_t)^{-1} d_t^{\lceil Ct \rceil} \\ & \leq & d^{-1}(1-d_t)^{-1} g(\phi_0)^{\lceil Ct \rceil} t^{C\beta} e^{\beta t^{-1} \log t} & \leq & d^{-1}(1-d_t)^{-1} t^{C\beta} e^{\beta t^{-1} \log t}. \end{array}$$

Substituting this into (24) we obtain

$$\limsup_{t \to \infty} \frac{\log \mathrm{E}[e^{\theta_t A_t}]}{\log t} \leq \limsup_{t \to \infty} \frac{C\beta \log t + \log(1 + d^{-1}(1-d_t)^{-1} e^{\beta t^{-1} \log t})}{\log t} = C\beta,$$

which together with (23) establishes the required inequality

$$\Lambda_2(\mu) \geq (\mu - C)\beta > 0.$$

$\square$

Lemma 4.2 *Any Poisson on-off process with a finite conditional mean rate κ satisfies Assumption 3.2(2). Moreover, if the on-time and off-time distributions T_{on} and T_{off} satisfy*

$$\mathrm{E}(T_{on}^{1+\zeta}) < \infty \quad \textit{for some } \zeta > 0, \tag{26}$$

and

$$\mathrm{E}(T_{off}) < \infty, \tag{27}$$

then Assumption 3.2(3) is also satisfied.

Proof. Let Y be the underlying on-off process associated with the Poisson on-off process, so that $Y_t \in \{0,1\}$ for $t \in [0,\infty)$ and $I_t \doteq \int_0^t Y_t \, dt$ represents the cumulative on-time in the interval $[0,t]$. Note that for every $t \in [0,\infty)$ $I_t \leq t$ and $A_t \mid I_t$ is a Poisson random variable with mean κI_t. Thus for any $t \in [0,\infty)$ and $\theta \in I\!R$

$$\mathrm{E}[e^{\theta A_t}] = \mathrm{E}\left[\mathrm{E}[e^{\theta A_t} \mid I_t]\right] = \mathrm{E}[e^{(e^{\theta}-1)\kappa I_t}] \leq e^{(e^{\theta}-1)\kappa t} < \infty, \tag{28}$$

which establishes Assumption 3.2(2).

To prove Assumption 3.2(3), first note that for any $\varepsilon > 0$, there exists γ_ε such that for $\theta \in (0, \gamma_\varepsilon)$,

$$e^{\theta} < 1 + (1+\varepsilon)\theta.$$

Also note that if $\theta_t \doteq \beta \log t/t$ for some $\beta > 0$, $\theta_t \to 0$ as $t \to \infty$. Thus for all t sufficiently large so that $\theta_t \in (0, \gamma_\varepsilon)$, by (28) and the above inequality we have

$$\mathrm{E}[e^{\theta_t A_t}] = \mathrm{E}[e^{(e^{\theta_t}-1)\kappa I_t}] \le \mathrm{E}[e^{(1+\varepsilon)\kappa\theta_t I_t}] \le e^{\varepsilon\kappa\theta_t t}\mathrm{E}[e^{\kappa\theta_t I_t}] = t^{\varepsilon\kappa\beta}\mathrm{E}[e^{\kappa\theta_t I_t}].$$

Therefore we obtain

$$\Lambda_2(\mu) \ge \beta\mu - \limsup_{t\to\infty} \frac{\log \mathrm{E}[e^{\theta_t A_t}]}{\log t} \ge \beta\mu - \limsup_{t\to\infty} \frac{\log \mathrm{E}[e^{\theta_t \kappa I_t}]}{\log t} - \beta\varepsilon\kappa.$$

Sending $\varepsilon \downarrow 0$ we conclude that

$$\Lambda_2(\mu) \ge \beta\mu - \limsup_{t\to\infty} \frac{\log \mathrm{E}[e^{\theta_t \kappa I_t}]}{\log t}. \tag{29}$$

However from the assumptions (26) and (27) on the on-time and off-time distributions, the proof of Proposition 3.3 of Mandjes and Borst [21] shows that for any $C \in (\lambda, \mu)$ there exist positive K and α such that

$$\mathrm{P}(I_t > Ct) \le Kt^{-\alpha}.$$

Then by Proposition 3.1 of [19], it follows that there exists $\beta > 0$ such that the right hand side of (29) is strictly greater than 0. This establishes Assumption 3.2(3). □

Theorem 4.3 *Any simple renewal point process whose inter-arrival distribution has mean* $\lambda^{-1} > 0$ *satisfies Assumption 3.2. Any simple Poisson on-off process with on-time and off-time distribution satisfying (26) and (27) satisfies Assumption 3.2. In particular, the processes in Classes 1)-5) defined in Section 4.1 satisfy Assumption 3.2.*

Proof. The first statement follows directly from Lemma 4.1 and 4.2. This implies that classes 1), 2) and 5) satisfy Assumption 3.2. Since T_{on} satisfies (26) if it is either an exponential distribution or Pareto distribution with tail parameter $\alpha \in (1, 2)$, and T_{off} satisfies (27) if it is an exponential distribution, Classes 3) and 4) satisfy Assumption 3.2 by Lemma 4.2. □

4.3 Simulation Results

We simulate representative examples from each of the five classes described in Section 4.1 with specific parameter values as summarised in Table 1. The parameter values were chosen such that each arrival process has a mean packet arrival rate $\lambda = 5.45$. In order to provide an intuitive feel of the relative burstiness of the the five arrival point processes that were simulated, in Figure 1 we have used Trellis graphics [7] to illustrate a sample realisation of each example in the time period [0, 91.5], which contains approximately 500 packets. The illustration suggests that the Poisson arrival process (top panel) is more regular than Classes 1-4 (bottom 4 panels), which appear to have more clustered arrivals.

Table 1: Description of five source arrival models in the simulation (mean arrival rate = 5.45).

Class	Name	Description and Parameters	I_∞
1	Weibull	inter-arrival time $\sim$ Weibull($\alpha = 0.5$)	5
2	InvWeibull	inter-arrival time $\sim$ Inverse-Weibull($\alpha = 1.5$)	∞
3	ExponExpoff	on period $\sim$ Exp(1), off period $\sim$ Exp(10), rate $\kappa = 60$	100.2
4	ParetoonExpoff	on period $\sim$ Pareto($\alpha = 1.5$), off period $\sim$ Exp(10), rate $\kappa = 60$	∞
5	Gamma	inter-arrival time $\sim$ Gamma($\alpha = 5$)	0.2

Figure 2 plots the buffer exceedance probability (in the $\log_{10}$ scale) against buffer size $b = 0, \ldots, 100$ for processes from the five models (see Table 1), and compares them with the unfinished work distribution U_{N^λ} for the Poisson process, which was obtained using the formula given in (4). We consider three different degrees of superposition, namely $n = 1, 100, 1000$, and two utilisation levels $\rho = 0.3$ and $\rho = 0.7$. Each panel has six curves – the five shaded lines represent the five models, and the thicker black line represents the Poisson arrival process. When $n = 1$, as expected the steady state unfinished work distributions corresponding to the different arrival processes are very different. The two Poisson on-off models show a greater exceedance probability in comparison with the others, and all curves with the exception of the **Gamma** process have an exceedance probability larger than that of Poisson. This is not surprising since Classes 1-4 are burstier than Poisson. However, as n increases, the difference between these curves decreases ($n = 100$) and eventually all the

Figure 1: A Realisation of the Five Arrival Processes in the Simulation.

curves converge to that of Poisson ($n = 1000$). At lower utilisations, convergence seems to be achieved at lower values of n. For example, when $\rho = 0.3$ almost complete convergence is achieved for all classes for $n = 100$, while for $\rho = 0.7$ it takes $n = 1000$ for convergence to occur. Also notice that alhtough the commonly used burstiness measure I_∞ of the arrival processes does not change, the queueing behaviour changes drastically with superposition. In particular, even though I_∞ is infinite for both **InvWeibull** and **ParetoonExpoff**, at $n = 100$ they both show smaller buffer exceedance probabilities (at least for buffer sizes between 0 and 100) than **ExponExpoff**, which has a finite I_∞. In fact, the buffer exceedance probability for the **InvWeibull** is quite close to Poisson even for $n = 1$, particularly for low utilisations. This indicates that the limiting IDC I_∞ is a rather poor metric for measuring the impact of the burstiness of superposed processes on queueing behaviour for small buffers.

For each of the five representative examples, Figure 3 plots the absolute normalised

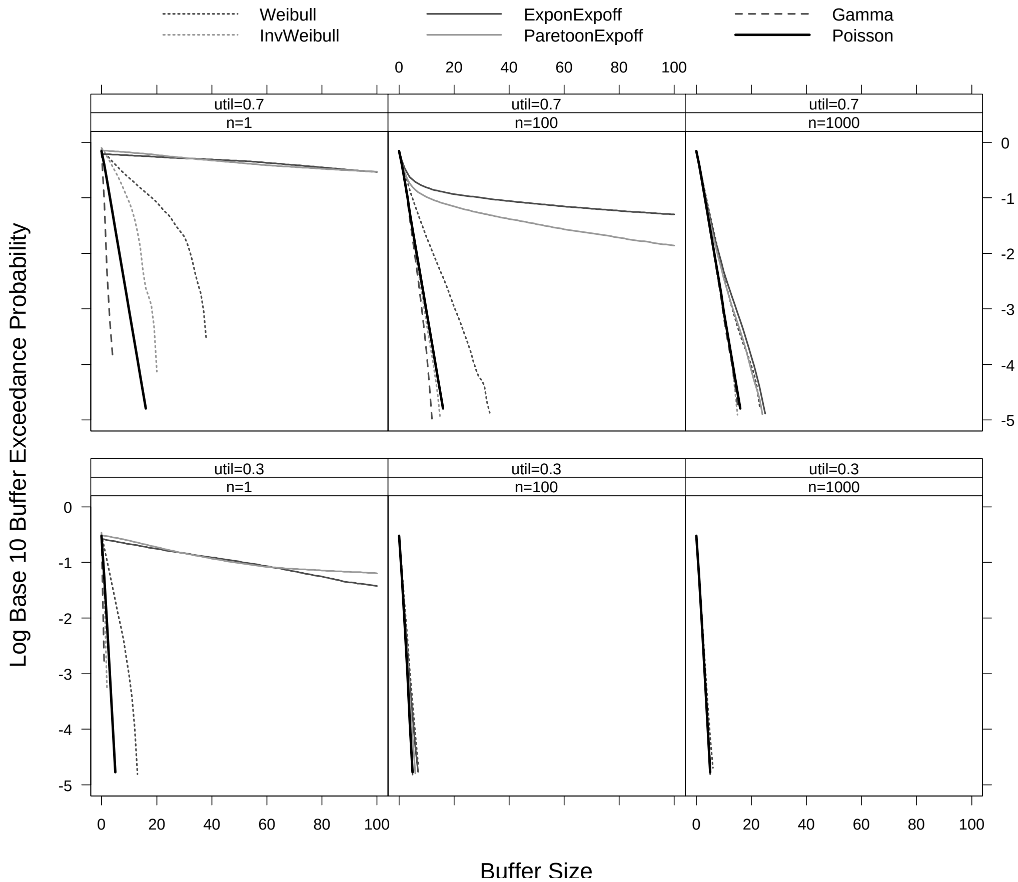

Figure 2: Buffer exceedance probability against buffer sizes, for superposed processes from the five models when $n = 1, 100, 1000$.

difference of the logarithm of the exceedance probability from that of the Poisson process against the square root of the number of superpositions n, once again for two utilisation levels (0.3 and 0.7). Let $p_{A^n}(b)$ be the probability of exceeding a threshold size b for the superposed arrival process A^n, and let $p_{N^\lambda}(b)$ be the corresponding probability for the Poisson process. Then the normalised error is defined as

$$\frac{|\log_{10} p_{A^n}(b) - \log_{10} p_{N^\lambda}(b)|}{\sup_{n\in I\!N} |\log_{10} p_{A^n}(b) - \log_{10} p_{N^\lambda}(b)|}. \tag{30}$$

In the simulation the supremum in the denominator, i.e. the maximum of the absolute error is estimated using the observed values from $1 \leq n \leq 1000$ (Table 2). This maximum is mostly achieved at small values of n. The three panels represent three different buffer sizes, namely $b = 1, 3, 6$, with the five lines representing the five models. Notice we also removed two curves for **InvWeibull** and **Gamma** at utilization 0.3 and buffer size 6. This is because the exceedance probability is very small in these cases (on the order of 10^{-6}), and our

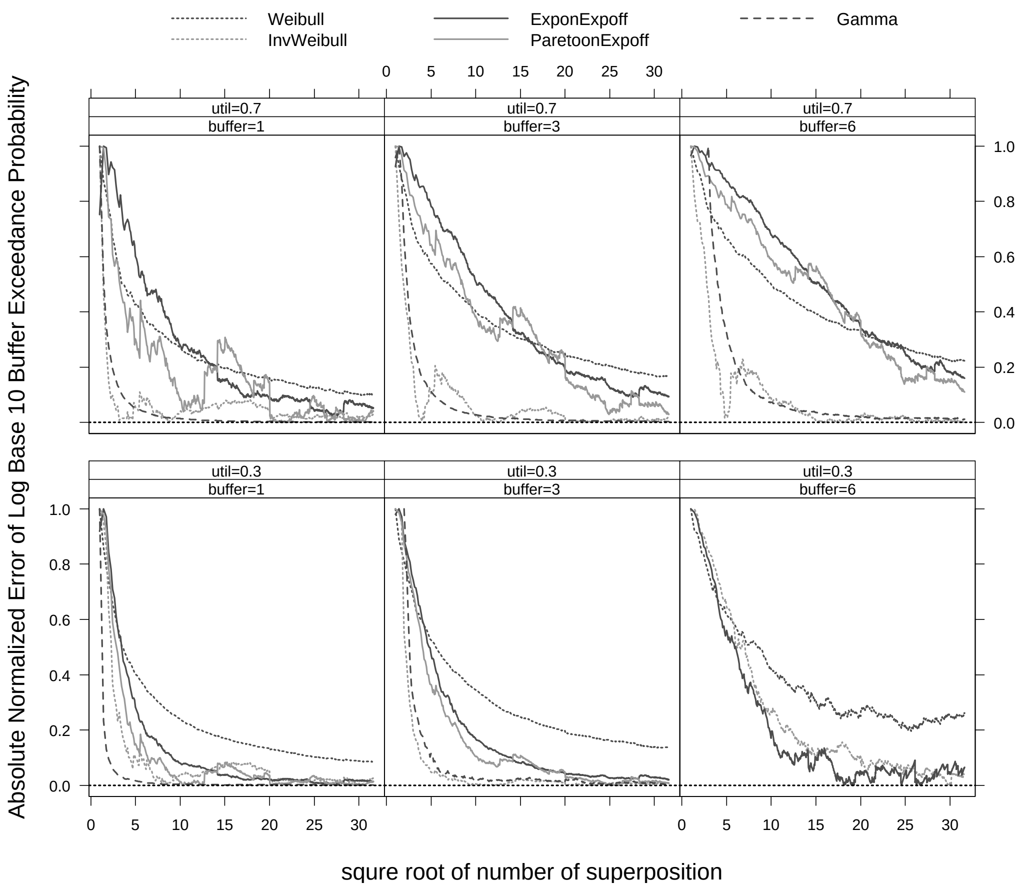

Figure 3: Absolute normalised error of buffer exceedance probability against different superpositions, for buffer size $b = 1, 3, 6$ and for superposed processes from the five models.

simulation does not give enough accuracy for small probabilities. These figures suggest that convergence is slower for larger buffer sizes and higher utilisation levels. Note from Table 2 that the normalising quantity in (30) (i.e. the supremum of the absolute error) is larger for larger buffer sizes, and is smaller for higher utilisations. Overall, convergence to the Poisson limit for large buffers and high utilisation seems to take place at higher values of n (see both Figures 2 and 3). Between different classes of arrival processes the convergence seems fastest for **InvWeibull** and **Gamma** and slowest for **Weibull**. The slow convergence for Weibull may be related to the propensity of the Weibull distribution to create short interarrivals. Indeed it can be shown that as t goes to zero, the **Weibull** arrival process with shape parameter α satisfies

$$\mathrm{P}(A_t \geq 2) = O(t^{1+\alpha}).$$

Table 2: Supremum of the absolute error of the logarithm of the buffer exceedance probability from that of Poisson for the five source models.

Buffer, Utilization	Weibull	InvWeibull	ExponExpoff	ExponExpoff	Gamma
$b = 1,\ \rho = 0.3$	0.52	0.22	0.73	0.74	1.61
$b = 3$	1.78	2.53	2.42	2.44	1.64
$b = 6$	3.74		5.07	5.10	
$b = 1,\ \rho = 0.7$	0.22	0.24	0.25	0.25	0.64
$b = 3$	0.68	0.63	0.82	0.82	2.32
$b = 6$	1.40	1.23	1.69	1.69	3.25

However, for the other models

$$\mathrm{P}(A_t \geq 2) = O(t^2).$$

5 Implications for Network Design

In this paper we have proved a strong insensitivity result for the unfinished work of superposed point processes, in the limit as the number of sources, and proportionally the processing rate, goes to infinity. Specifically, we show that the actual steady state probability (and not just the logarithmic asymptote) of the unfinished work exceeding a *fixed* positive threshold tends to the corresponding probability assuming the arrival processes were Poisson. Simulations suggest that at moderate utilisations, for a large class of traffic models convergence is achieved at reasonable levels of superposition, say when several hundreds of sources are multiplexed.

Our result supports recent statistical analysis of Internet traffic [4, 5]. In [4, 5], Internet traces collected at a corporate research site and several Universities are studied. These sites typically have a few hundred to a few thousand simultaneous streams during a busy period. It is shown empirically that with increase in the number of superpositions the packet arrival process tends toward Poisson and packet sizes tend toward independence. An open-loop queueing study was also carried out in [5] using traces from the corporate research site with different degrees of superposition. It was shown that when the utilisation is kept fixed, the effect of the long-range dependence of network traffic diminishes and the steady state exceedance probabilities (assuming an infinite buffer) tend to the Poisson limit. This is to be contrasted with the study in [14], where a single Ethernet LAN trace is used to demonstrate the impact of long-range dependence on queueing performance, with almost no reference to the degree of superposition. In short, our results suggest that with increase in the amount of superposition, the effect of long-range dependence on queueing reduces even more markedly than suggested by previous results.

As discussed in Section 3.3, our results differ from previous ones with respect to the type of scaling considered. In our analysis the buffer size is $O(1)$, so that the time scale of interest is inversely proportional to the number n of sources superposed (see Theorem 3.7), and hence proportional to the mean packet inter-arrival time. When n increases, this inter-arrival time decreases, while the packet inter-arrivals from an individual source remains the same. Thus on that time scale the characteristics of an individual source become less important, the packet inter-arrivals tend to look more Poisson-like, and the steady state exceedance probability also tends to that of Poisson. Our asymptotic analysis is valid for a whole range of magnitudes of the exceedance probability, and not just for small probabilities, as is often the case in large deviations analyses. Previous analyses [2, 8, 26, 12, 19, 22] consider a different limiting regime, where the buffer size is $O(n)$, in which case the time scale of interest does not change with superposition. Those results show that the logarithm of the exceedance probability decreases exponentially with n, and the decay rate depends on the distribution of the arrival process. Moreover, the nature of the asymptotic analysis implies that the results are valid only in a regime where exceedance probabilities are small. (Note that [19] also provides estimates of the prefactor.)

The different scalings have important practical implications for network design. Both our analysis as well as that of [2, 8, 26, 12, 19, 22] show that statistical multiplexing gains obtained from superposing many sources are significant even when individual sources have long-range dependent behaviour. However the results of [2, 8, 26, 12, 19, 22] only show that multiplexing gain (i.e. an exponential decay rate) can be achieved when the buffer size increases proportionally to the number of sources multiplexed. Our results suggest that for high speed systems (assuming that the characteristics of the individual sources do not change) relatively small buffers may be adequate in order to achieve the same loss rate. In particular this implies that a system design with a linear increase in the buffer size may be overly conservative. As mentioned in the introduction, this is particularly relevant in cases where the cost of high-speed memory is high.

Another important observation is that since the critical time scale becomes with superposition for our scaling, a fluid approximation of the processes may no longer be valid. Thus we model our packet arrival processes as point processes. The point process perspective also seems to reflect more accurately the discrete nature of packet arrivals in routers and switches. Some authors have considered models in which time is discretised and the arrival process is represented in terms of the accumulated workload over each discrete interval. This approach appears to be too crude, especially if the size of the discretised interval is kept fixed with increase in the degree of superposition since in that case the accumulated workload tends to infinity.

In future work, one would like to gain more insight into the dependence of the convergence rates on various parameters of the system, in order to develop a useful heuristic for sizing buffers in a network to smooth burstiness effects. It would also be useful to generalise the analysis here to the case of marked point processes, to model the more realistic case when packets are of different sizes. In particular, it would be of interest to analytically determine if indeed the packet size process tends towards independence with increase in the number of sources multiplexed. Here we have used the probability of exceeding a level b as a proxy for the overflow probability in a finite buffer of size b. It would be worthwhile to see how accurate this approximation is in practice. Finally, here we have considered only the steady state unfinished workload distribution. There are other quantities of interest such as the waiting time and queue length processes that may also be of interest.

Acknowledgments. We would like to thank Bill Cleveland, Don Sun and Dong Lin for their contributions to the empirical work that suggested this result. We would also like to thank Matthew Andrews for giving a pointer to the website [30].

References

[1] P. Billingsley. *Convergence of Probability Measures.* John Wiley, New York, 1968.

[2] D. D. Botvich and N. G. Duffield. Large deviations, the shape of the loss curve, and economies of scale in large multiplexers. *Queueing Systems*, pages 293–320, 1995.

[3] O. J. Boxma and V. Dumas. Fluid queues with long-tailed activity period distributions. *Computer Communications*, 21:1509–1529, 1998.

[4] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. Internet Traffic Tends to Poisson and Independence as the Load Increases. Technical report, Bell Labs, Murray Hill, NJ, 2001.

[5] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. On the Nonstationarity of Internet Traffic. *Proc. ACM SIGMETRICS*, 2001.

[6] G. L. Choudhury, D. M. Lucantoni, and W. Whitt. Squeezing the Most out of ATM. *IEEE transactions on communications*, 44(2), 1996.

[7] W. S. Cleveland and M. Fuentes. Trellis Display: Modeling Data from Designed Experiments. Technical report, Bell Labs, 1997. black and white postscript: http://cm.bell-labs.com/stat/doc/trellis.doe.bw.ps. color postscript: http://cm.bell-labs.com/stat/doc/trellis.tour.col.ps.

[8] C. Courcoubetis and R. Weber. Buffer overflow asymptotics for a buffer handling many traffic sources. *Journal of Applied Probability*, 33:886–903, 1996.

[9] D. R. Cox and P. A. W. Lewis. *The Statistics Analysis of Series of Events.* Methuen, London, England, 1966.

[10] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes.* Springer-Verlag, New York, 1988.

[11] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications.* Jones and Bartlett, Boston, 1992.

[12] N. G. Duffield. Economies of scale in queues with sources having power-law large deviations scalings. *Journal of Applied Probability*, 33:840–857, 1996.

[13] A. Elwalid, D. Heyman, T.V.Lakshman, D. Mitra, and A. Weiss. Fundamental Bounds and Approximations to ATM Multiplexers with Applications to Video Teleconferencing. *IEEE J. Sel. Areas in Comm.*, 13(6):1004–1016, 1995.

[14] A. O. Erramilli, O. Narayan, and W. Willinger. Experimental Queueing Analysis with Long-Range Dependent Packet Traffic. *IEEE/ACM Transactions on Networking*, 4:209–223, 1996.

[15] D. Gross and C. M. Harris. *Fundamentals of Queueing Theory.* Wiley, 1974.

[16] U. M. J. Roberts and J. Virtamo. *Broadband Network Teletraffic, Final Report of Action COST 242.* Springer, Berlin, 1996.

[17] T. Konstantopoulos and S.-J. Lin. High variability vs long-range dependence for network performance. In *Proc. 35th. IEEE CDC*, pages 1354–1359, Kobe, Dec. 1996.

[18] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the Self-Similar Nature of Ethernet Traffic. *IEEE/ACM Transactions on Networking*, 2:1–15, 1994.

[19] N. Likhanov and R. Mazumdar. Cell loss asymptotics in buffers fed with a large number of independent stationary sources. In *Proceedings IEEE Infocom*, pages 339–346, 1998.

[20] R. M. Loynes. The stability of a queue with non-independent inter-arrival and service times. *Proc. Camb. Philos. Soc.*, pages 58:497–520, 1962.

[21] M. Mandjes and S. Borst. Overflow behavior in queues with many long-tailed inputs. *Advances in Applied Probability*, 32:1150–1167, 2001.

[22] M. Mandjes and J. H. Kim. Large deviations for small buffers: an insensitivity result. *Queuing Systems*, 2001. to appear.

[23] V. Paxson and S. Floyd. Wide-Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM Transactions on Networking*, 3:226–244, 1995.

[24] B. K. Ryu and A. Elwalid. The Importance of Long-Range Dependence of VBR Traffic in ATM Traffic Engineering: Myths and Realities. In *Proc. ACM SIGCOMM '96*, pages 3–14, 1996.

[25] B. K. Ryu and S. B. Lowen. Point process approaches to the modeling and analysis of self-similar traffic: Part I–Model constrction. In *Proc. IEEE Infocom '96*, pages 1468–1475, 1996.

[26] A. Simonian and J. Guibert. Large deviations approximation for fluid queues fed by a large number of on/off sources. *IEEE Journal of Selected Area in Communications*, 13:1017–1027, 1995.

[27] K. Sriram and W. Whitt. Charactering superposition arrival processes in packet multiplexers for voice and data. *IEEE Journal on Selected Areas in Communications*, SAC-4:833–846, 1986.

[28] W. Whitt. Some useful functions for functional limit theorems. *Mathematics of Operations Research*, 5(1):67–85, 1980.

[29] 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 Transactions on Networking*, 5:71–86, 1997.

[30] World of Computers. Rambus Memory, 2001. http://www.woc.co.uk/pricelist.asp.