

Title	Large deviations, the shape of the loss curve, and economies of scale in large multiplexers
Creators	Botvich, D. D. and Duffield, N. G.
Date	1994
Citation	Botvich, D. D. and Duffield, N. G. (1994) Large deviations, the shape of the loss curve, and economies of scale in large multiplexers. (Preprint)
URL	https://dair.dias.ie/id/eprint/700/
DOI	DIAS-STP-94-12

Large deviations, the shape of the loss curve, and economies of scale in large multiplexers

D.D. Botvich * N.G. Duffield†

May 13, 1994

Abstract

We analyse the queue Q^L at a multiplexer with L inputs. We obtain a large deviation result, namely that under very general conditions

$$\lim_{L \rightarrow \infty} L^{-1} \log P[Q^L > Lb] = -I(b)$$

provided the offered load is held constant, where the shape function I is expressed in terms of the cumulant generating functions of the input traffic. This provides an improvement on the usual effective bandwidth approximation $P[Q^L > b] \approx e^{-\delta b}$, replacing it with $P[Q^L > b] \approx e^{-LI(b/L)}$. The difference $I(b) - \delta b$ determines the economies of scale which are to be obtained in large multiplexers. If the limit $\nu = -\lim_{t \rightarrow \infty} t\lambda_t(\delta)$ exists (here λ_t is the finite time cumulant of the workload process) then $\lim_{b \rightarrow \infty} (I(b) - \delta b) = \nu$. We apply this idea to a number of examples of arrivals processes: heterogeneous superpositions, Gaussian processes, Markovian additive processes and Poisson processes. We obtain expressions for ν in these cases. ν is zero for independent arrivals, but positive for arrivals with positive correlations. Thus economies of scale are obtainable for highly bursty traffic expected in ATM multiplexing.

Keywords: Large deviations, scaling limits, ATM multiplexers, heterogeneous superpositions.

DIAS-APG-94-12

*School of Electronic Engineering, Dublin City University, Dublin 9, Ireland; E-mail: botvichd@dcu.ie

†Corresponding Author. School of Mathematical Sciences, Dublin City University, Dublin 9, Ireland; Dublin Institute for Advanced Studies, 10 Burlington Road, Dublin 4, Ireland. E-mail: duffieldn@dcu.ie

1 Introduction.

The problem of determining loss probabilities in queueing systems is crucial in the development of emergent technology of telecommunications networks using the Asynchronous Transfer Mode (ATM). Much recent work has focused on the analysis of the single server queue with general arrivals. This enables one to analyse queues with correlated arrivals, such as those which occur in the buffer of an ATM multiplexer whose input is a superposition of highly bursty sources.

Consider a general single server queue. For $t \in T$ (here $T = \mathbf{R}_+$ or \mathbf{Z}_+) denote by A_t the amount of work which arrives to be processed in the interval $[-t, 0)$ and by S_t the amount which can be processed in the same interval. If more work arrived than can be processed, the surplus waits in the queue. The workload process W is defined by $W_0 = 0$ and

$$W_t = A_t - S_t, \quad (1.1)$$

and the queue of unprocessed work at time zero is

$$Q = \sup_{t \geq 0} W_t. \quad (1.2)$$

The relation between the tail of the queue length distribution and the large deviation properties of the workload processes has been established in progressive degrees of generality. Following a heuristic proposal by Kesidis *et al* [20] (see also [29, 5] for further bibliographical details), Glynn and Whitt [14] showed for $T = \mathbf{Z}_+$ that if the pair $(W_t/t, t)$ satisfy a *large deviation principle* then

$$\lim_{b \rightarrow \infty} b^{-1} \log P[Q > b] = -\delta, \quad (1.3)$$

where

$$\delta = \sup\{\theta \mid \lambda(\theta) \leq 0\}, \quad (1.4)$$

and λ is the cumulant generating function of the workload process defined by

$$\lambda(\theta) = \lim_{t \rightarrow \infty} t^{-1} \log E[e^{\theta W_t}]. \quad (1.5)$$

Alternatively, δ can be expressed through

$$\delta = \inf_{t > 0} t \lambda^*(t^{-1}) \quad (1.6)$$

where λ^* , the Legendre-Fenchel transform of λ , is defined through

$$\lambda^*(x) := \sup_{\theta} (x\theta - \lambda(\theta)). \quad (1.7)$$

(We refer the reader to the book Dembo and Zeitouni [8] as a comprehensive reference for large deviations, that of Bucklew [3] for a more heuristic approach, and the article of Lewis and Pfister [24] for a general introduction).

Recently, Duffield and O'Connell have extended this result in two directions [10]. Firstly, the same result is shown to hold when $T = \mathbf{R}_+$, subject to a local growth condition on W . Secondly an analogous result holds with large deviation scalings more general than b^{-1}, t^{-1} in (1.3) and (1.5). These are appropriate for treating, for example, the case where the workload is fractional Brownian motion: this has been proposed as a model for the workload by Leland *et al* [23], based on observations of Ethernet traffic.

The relation (1.3) is the basis of the *effective bandwidth approximation* to the queue length distribution:

$$P[Q > b] \approx e^{-\delta b}. \quad (1.8)$$

(See for example, [5, 16, 17, 13, 19, 30] for development, applications and further references). The motivation here is that for ATM multiplexers one wants to estimate exponentially small loss probabilities, which in practice are to be as small as 10^{-9} . However, there is already theoretical and numerical work indicating that (1.8) is insufficient for this purpose. For a queue where the input is an L -fold superposition of markovian sources served at constant rate s , Duffield [9] proves the upper bound

$$P[Q > b] \leq e^{-\mu L} e^{-\delta b} \quad (1.9)$$

δ is as before and μ and δ depend only the the traffic due to a single source, and on the offered load through the ratio s/L . In the example of on-off markov sources with positive autocorrelation, μ is strictly positive (see Buffet and Duffield [4]). Thus in a large superposition, the loss probabilities may be exponentially small even for small b : the effective bandwidth approximation (1.8) can be extremely conservative through over-estimating the loss probabilities. On the other hand, if μ were negative, then (1.9) suggests that (1.8) will under-estimate the loss probabilities at large L , even for large b . Moreover, both types of error become more severe as L increases at constant load. Both these types of behaviour have been observed though numerical studies of queueing models by Choudhury *et al* [7], the effective bandwidth approximation, when compared with these results, is shown to over-estimate the loss probabilities in examples of bursty sources, and under-estimate them in examples of sub-bursty sources.

Thus we are led to investigate the large deviation properties of the queue length distribution in the size L . Apart from the above considerations, we are motivated by the observation in examples that the broad features of the queue length distribution remain roughly invariant when both the size L and the queue length b are jointly scaled. (See simulation results in the thesis of Corcoran [6] and heuristic arguments by Rasmussen *et al* [26]). For example, let Q^L be the queue due to a superposition of L identical sources, served at constant rate sL (s fixed), and denote by λ_t the finite time cumulant due to a single source with arrival process A served at rate s :

$$\lambda_t(\theta) = t^{-1} \log E[e^{\theta(A_t - st)}]. \quad (1.10)$$

The workload W_t^L of the superposition is the sum of L independent copies of $W_t = A_t - st$ and $Q_t^L = \sup_{t \geq 0} W_t^L$. As a consequence of Theorem 1 of the next section it follows (under suitable local regularity conditions on the workload) that for $b \geq 0$

$$\lim_{L \rightarrow \infty} L^{-1} \log P[Q^L > Lb] = -I(b) \quad (1.11)$$

where I is the *shape function* defined by

$$I(b) = \inf_{t > 0} t\lambda_t^*(b/t). \quad (1.12)$$

Such a result is shown to hold for heterogeneous superpositions also. The heuristic interpretation of (1.11) is that for each t , the family $(W_t^L/L, L)$ of superposed workloads satisfies a large deviation principle with rate function $(t\lambda_t)^*(\cdot) = t\lambda_t^*(\cdot/t)$. The large deviation heuristic is that *rare events occur in the most likely way*: the most likely way for a (rescaled) queue length b to be exceeded corresponds to the infimum of $t \rightarrow t\lambda_t^*(b/t)$.

The shape function can be seen to give the large scale corrections to the effective bandwidth approximation: (1.8) is replaced by

$$P[Q^L > b] \approx e^{-LI(b/L)}. \quad (1.13)$$

(Observe that if one replaces λ_t^* by λ^* in (1.12), where $\lambda = \lim_{t \rightarrow \infty} \lambda_t$, then $I(b)$ reverts to $b\delta$). $I(b) - \delta b$ determines the error incurred by using the effective bandwidth at large L , or to make a more positive statement, it determines the *economies of scale* to be obtained in multiplexers of a large number of sources.

We examine the initial value $I(0)$ and asymptotics as $b \rightarrow \infty$ of the shape function I in Theorems 2 and 3. For $T = \mathbf{Z}_+$, a workload with stationary increments, $I(0) = \lambda_1^*(0)$. This is just the large deviation result for the loss probability in a bufferless resource as found by Hui [16, 17]. The asymptotics of I are

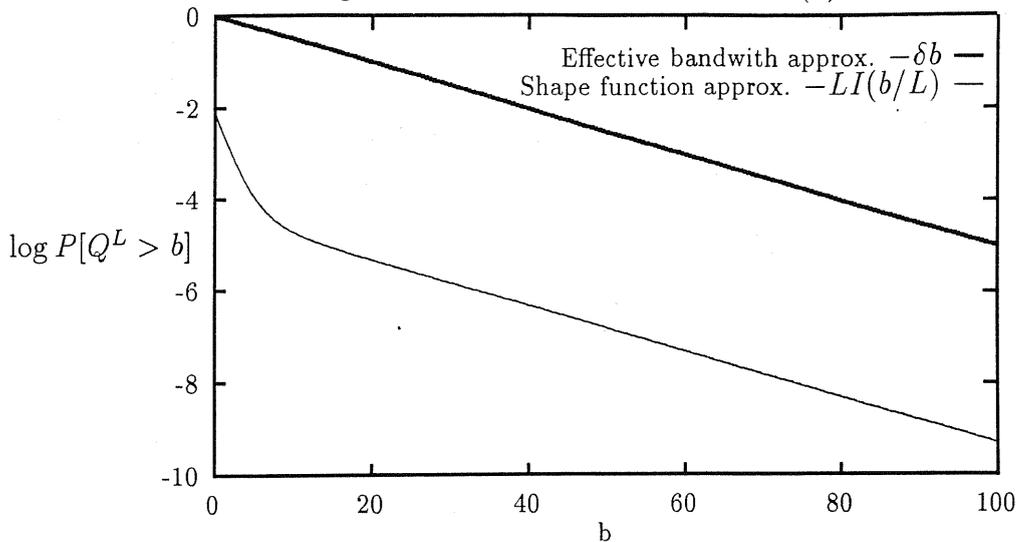
$$\lim_{b \rightarrow \infty} (I(b) - \delta b) = \nu \quad (1.14)$$

where

$$\nu = - \lim_{t \rightarrow \infty} t\lambda_t(\delta) \quad (1.15)$$

provided this limit exists, and subject to some regularity conditions in the case $T = \mathbf{Z}_+$. One sees immediately that $\nu = 0$ for uncorrelated arrivals, since then $\lambda_t(\delta) = \lambda(\delta) = 0$. Thus there are no economies of scale to be gained at large (rescaled) buffer sizes for uncorrelated arrivals, since then $-\delta b$ is asymptotic to $I(b)$ for large b . On the other hand, a sufficient condition for ν to be positive is that the workloads on disjoint time intervals are positively correlated (Theorem 4). This is typically the case for highly bursty sources. A generic configuration with $\nu > I(0) > 0$ is illustrated in Figure 1.

Figure 1: Economies of Scale with $\nu > I(0) > 0$



It is interesting to note that ν depends on finer details of the workload process than those which determine the asymptotic slope δ : it depends not only on the limiting cumulant λ but rather on the manner in which the λ_t approach λ as $t \rightarrow \infty$. To borrow from the terminology of physics, ν is not a thermodynamic quantity.

The paper is organised in the following way. The foregoing results are fully stated in section 2 and proved in section 3. In section 4 we apply them to a number of examples. The case of heterogeneous superpositions is worked out in 4.1. Gaussian workload processes are covered in 4.2 and the specific example of Ornstein-Uhlenbeck processes in 4.3, including a calculation of the shape function for a heterogeneous superposition. Markov Additive Processes are treated in 4.4. In this case we can express ν in terms of the maximal eigenfunction of the (Laplace transform) of the Markov transition kernel (Corollary 5). Comparisons of the approximation (1.13) with simulation are made for superpositions of Markovian on-off sources. Finally, in 4.5 we apply the results to a very simple class of examples: independent Poissonian arrivals with general service distribution. In light of the explicit distribution for $\sum^L M(\kappa)/M(L\mu)/1$ it is not surprising that in this case $I(0) = \nu = 0$: there are no economies of scale to be obtained for Poissonian arrivals at any buffer size.

2 Large deviations.

We begin by stating our hypotheses concerning the workload processes, then give some examples which satisfy the hypotheses. For each $L \in \mathbf{N}$, $(W_t^L)_{t \in T}$ (where $T = \mathbf{Z}_+$ or \mathbf{R}_+) is

a stochastic process, and $W_0^L = 0$. The queue length at time zero is

$$Q^L = \sup_{t \in T} W_t^L. \quad (2.1)$$

(Note that if the increments of W^L are stationary, then the distribution of Q^L is also stationary). For $\theta \in \mathbf{R}$ define the cumulant generating function

$$\lambda_t^L(\theta) = (Lt)^{-1} \log E[e^{\theta W_t^L}]. \quad (2.2)$$

Hypothesis 1

(i) For each $\theta \in \mathbf{R}$, the limits

$$\lambda_t(\theta) = \lim_{L \rightarrow \infty} \lambda_t^L(\theta) \quad \text{and} \quad \lambda(\theta) = \lim_{t \rightarrow \infty} \lambda_t(\theta) \quad (2.3)$$

exist as extended real numbers. Moreover, the first limit exists uniformly for all t sufficiently large.

(ii) λ_t and λ are essentially smooth. (Both are automatically convex by Hölder's inequality).

(iii) There exists $\theta > 0$ for which $\lambda_t(\theta) < 0$ for all $t \in T$.

(iv) ($T = \mathbf{R}_+$) For all $t \geq r \geq 0$ define $\tilde{W}_{t,r}^L = \sup_{0 < r' < r} |W_{t-r'}^L - W_t^L|$. Then for all $\theta \in \mathbf{R}$

$$\limsup_{r \rightarrow 0} \limsup_{L \rightarrow \infty} L^{-1} \sup_{t \geq 0} \log E[e^{\theta \tilde{W}_{t,r}^L}] = 0. \quad (2.4)$$

Remark: if Hypotheses 1(i),(ii) are satisfied, then by the Gärtner-Ellis theorem, for each t the pair $(W_t^L/L, L)$ satisfies a large deviation principle with good rate function given by the Legendre-Fenchel transform of λ_t , which we denote by λ_t^* . In other words, for any Borel set Γ ,

$$\limsup_{t \rightarrow \infty} L^{-1} \log P(W_t^L/L \in \Gamma) \leq - \inf_{x \in \bar{\Gamma}} \lambda_t^*(x), \quad (2.5)$$

and

$$\liminf_{t \rightarrow \infty} L^{-1} \log P(W_t^L/L \in \Gamma) \geq - \inf_{x \in \Gamma^\circ} \lambda_t^*(x), \quad (2.6)$$

where

$$\lambda_t^*(x) := \sup_{\theta \in \mathbf{R}} \{\theta x - \lambda_t(\theta)\}. \quad (2.7)$$

For $x > x_t^*$ where $\lambda_t^*(x_t^*) = 0$, it follows that

$$\limsup_{t \rightarrow \infty} L^{-1} \log P(W_t^L/L > x) \leq -\lambda_t^*(x), \quad (2.8)$$

and

$$\liminf_{t \rightarrow \infty} L^{-1} \log P(W_t^L/L > x) \geq \lim_{y \searrow x} -\lambda_t^*(y). \quad (2.9)$$

Hypothesis 1(iii) is a stability condition: then there exists a strictly positive solution δ of the equation $\lambda(\delta) = 0$, which is the asymptotic decay rate of the queue length distribution. Hypothesis 1(iv) is a local regularity condition on the sample paths of the workload.

Examples

- **Homogeneous superpositions.** There are L identical sources whose backward arrival processes are independent copies of $(A_t)_{t \in T}$. The superposition is served at a constant rate sL : the offered load is independent of L . Then

$$\lambda_t(\theta) = \lambda_t^L(\theta) = t^{-1} \log E[e^{\theta A_t}] - s\theta, \quad (2.10)$$

independent of L . Thus λ_t is the cumulant corresponding to the workload $W_t = A_t - st$ of a single source served at rate s .

- **Heterogeneous superpositions.** Sources are classified by type j in some finite index set J . There are L_j sources of type j , with $L = \sum_j L_j$ sources in total. The backward arrival process for a source of type j is $(A_{j,t})_{t \in T}$. All sources are independent. Then

$$\lambda_t^L(\theta) = \sum_{j \in J} p_j^L c_{j,t}(\theta) - s, \quad (2.11)$$

where

$$c_{j,t}(\theta) = t^{-1} \log E[e^{\theta A_{j,t}}] \quad \text{and} \quad p_j^L = L_j/L. \quad (2.12)$$

The limits $c_j(\theta) = \lim_{t \rightarrow \infty} c_{j,t}(\theta)$ are assumed to exist with $c_j(\cdot)$ essentially smooth. Then for any $\theta \in \mathbf{R}$, $L \rightarrow \lambda_t^L(\theta)$ is convergent provided the limits $p_j = \lim_{L \rightarrow \infty} p_j^L$ exist, and the convergence is uniform in t since J is finite. Heterogeneous superpositions have been previously analysed through the effective bandwidth approximation (see references above) and through eigenfunction expansions for classes of Markovian fluid models by Kosten [21] (building on the early work of Anick *et al*[1] on homogeneous superpositions) and finite state models by Elwalid *et al* [12].

- **Time rescalings.** The single source arrival process is $A_{L,t}$ where and the process $(A_{L,Lt})$ is convergent in distribution to some process (A_t) as $L \rightarrow \infty$. The superposition is served at a fixed rate s . Thus with \sum^L denoting an L -fold superposition:

$$Q_t^L = \sup_{t \geq 0} ((\sum^L A_{L,t}) - st) \quad (2.13)$$

$$= \sup_{t \geq 0} ((\sum^L A_{L,Lt}) - Lst). \quad (2.14)$$

We shall take the limit $L \rightarrow \infty$ and so assume $\lambda_t^L(\theta) = (Lt)^{-1} \log E[\exp \theta \sum^L (A_{L,Lt} - st)]$ satisfies Hypothesis 1. This class of models is motivated by examples of rescaled Markovian sources simulated by Corcoran [6], and of rescaled renewal processes examined by Sriram and Whitt [28] and Rasmussen *et al* [26].

When $T = \mathbf{R}_+$, if the arrival process has stationary increments, then for homogeneous superpositions a sufficient condition for Hypothesis 1(iv) to be satisfied is that

$$\limsup_{r \rightarrow 0} \log E[e^{\theta \tilde{A}_r}] = 0 \quad (2.15)$$

for all $\theta \in \mathbf{R}$, where $\tilde{A}_r = \sup_{0 < r' < r} |A_{r'}|$; for heterogeneous superpositions each arrival process $A_{i,t}$ can be taken to satisfy such a condition.

Theorem 1 *Under Hypothesis 1, for each $b > 0$*

$$\lim_{L \rightarrow \infty} L^{-1} \log P[\sup_{t > 0} W_t^L > Lb] = -I(b) \quad (2.16)$$

$$\text{where } I(b) := \inf_{t > 0} t\lambda_t^*(b/t) \quad (2.17)$$

For $t > 0, \theta \in \mathbf{R}$ define $\Lambda_t(\theta) = t\lambda_t(\theta/t)$.

Hypothesis 2

(i) W_t^L has stationary increments.

(ii) ($T = \mathbf{R}_+$) The limit $\Lambda(\theta) = \lim_{t \rightarrow 0} \Lambda_t(\theta)$ exists as an extended real number for all $\theta \in \mathbf{R}$.

Theorem 2 *Under Hypotheses 1 and 2*

$$(T = \mathbf{Z}_+) \quad I(0) = \lambda_1^*(0).$$

$$(T = \mathbf{R}_+) \quad I(0) = \Lambda^*(0).$$

Remark: Theorem 2 says that a stable workload with stationary increments is (asymptotically) most likely to exceed 0 at the smallest times. But this need not be the case for non-stationary workloads. The identification of the asymptotics of I requires some technical conditions, as follows.

Hypothesis 3

(i) ($T = \mathbf{Z}_+$ or $T = \mathbf{R}_+$) The following limit exists:

$$\nu := - \lim_{t \rightarrow \infty} t\lambda_t(\delta). \quad (2.18)$$

(ii) ($T = \mathbf{Z}_+$) λ_t and λ are strictly convex and $t \mapsto (t+1)\lambda'_{t+1}(\delta) - t\lambda'_t(\delta)$ is bounded above.

(ii') ($T = \mathbf{Z}_+$) λ_t and λ are strictly convex; $(\lambda_t^*)'$ and $(\lambda^*)'$ are uniformly Lipschitz continuous on some neighbourhood of $\lambda'(\delta)$; and

$$(t+1)\lambda'_{t+1}(\delta) - t\lambda'_t(\delta) = o(\sqrt{t}). \quad (2.19)$$

Remark: Hypothesis 3(i) can be understood as follows. Let $\lambda_t(\delta_t) = 0$. Then $\lambda_t(\delta) \approx \lambda'_t(\delta)(\delta - \delta_t)$. So the existence of a finite limit ν means that $\delta_t - \delta \sim t^{-1}$.

Theorem 3 Under Hypothesis 3(i), and with the addition of Hypotheses 3(ii) or 3(ii') for $T = \mathbf{Z}_+$, then

$$\lim_{t \rightarrow \infty} (I(b) - \delta b) = \nu. \quad (2.20)$$

We shall say more concerning the existence of ν in the context of Markov Additive Processes in section 4.4. However, we can make a general statement concerning the sign of ν .

Theorem 4 Let W_t^L have stationary increments, and suppose for each L and for each $0 \leq t_1 \leq t_2 \leq t_3 \leq t_4$ that $W_{t_4} - W_{t_3}$ and $W_{t_2} - W_{t_1}$ are non-negatively correlated. Then if ν exists, it is non-negative.

3 Proofs of the Theorems.

Proof of Theorem 1 : Lower Bound

$$\liminf_{L \rightarrow \infty} L^{-1} \log P[\sup_{t>0} W_t^L > Lb] \geq \liminf_{L \rightarrow \infty} L^{-1} \sup_{t>0} \log P[W_t^L > Lb] \quad (3.1)$$

$$\geq \sup_{t>0} \liminf_{L \rightarrow \infty} L^{-1} \log P[W_t^L > Lb] \quad (3.2)$$

$$= \sup_{t>0} -t\lambda_t^*(b/t), \quad (3.3)$$

since by Hypothesis 1(i),(ii) and the Gärtner-Ellis Theorem, the pair $(W_t^L/L, L)$ satisfies a large deviation principle with good rate function $(t\lambda_t)^*(b) = t\lambda_t^*(b/t)$, and by Hypothesis 1(iii), x_t^* (as defined above (2.8)) can be chosen uniformly negative.

Upper Bound: $T = \mathbf{Z}_+$. For any $t, \theta > 0$ and $\theta_{t'} > 0, 0 < t' < t$,

$$P[\sup_{t'>0} W_{t'}^L > Lb] \leq t \max_{0 < t' < t} P[W_{t'}^L > Lb] + \sum_{t' \geq t} P[W_{t'}^L > Lb] \quad (3.4)$$

$$\leq t \max_{0 < t' < t} e^{-Lb\theta_{t'} + t'L\lambda_{t'}^L(\theta_{t'})} + e^{-\theta Lb} \sum_{t'' \geq t} e^{Lt''\lambda_{t''}^L(\theta)} \quad (3.5)$$

by Chebychev's inequality. Since $\lambda_t^L(\theta) \rightarrow \lambda_t(\theta)$ uniformly in t , $\lambda_t(\theta) \rightarrow \lambda(\theta)$ and $\lambda(\theta) < 0$ on $(0, \delta)$, we can find $\theta > 0$ and $\varepsilon < 0$ such that $\lambda_t^L(\theta) < \varepsilon$ for all L, t sufficiently large. This means that for such L and t the geometric series in (3.5) is summable, yielding

$$P[\sup_{t' > 0} W_{t'}^L > Lb] \leq t \max_{0 < t' < t} e^{-Lb\theta_{t'} + t'L\lambda_{t'}^L(\theta_{t'})} + e^{-\theta Lb + Lt\varepsilon} / (1 - e^{L\varepsilon}) \quad (3.6)$$

Taking logarithms, dividing by L , taking the limsup as $L \rightarrow \infty$ and finally taking the infimum over the $\theta_{t'}$ we obtain

$$\limsup_{L \rightarrow \infty} L^{-1} \log P[\sup_{t' > 0} W_{t'}^L > Lb] \leq \max \left(\max_{0 < t' < t} (-t'\lambda_{t'}^*(b/t')), -\theta b + t\varepsilon \right) \quad (3.7)$$

Recall $\varepsilon < 0$ so that taking the limit $t \rightarrow \infty$ we obtain (in conjunction with the lower bound) the stated result.

Upper Bound: $T = \mathbf{R}_+$. For any $\varepsilon > 0$ and $n \in \mathbf{N}$ define

$$\hat{W}_n^L = \sup_{(n-1)\varepsilon < t \leq n\varepsilon} W_t^L \quad \text{and} \quad \hat{\lambda}_n^L = (nL)^{-1} \log E[e^{\theta \hat{W}_n^L}]. \quad (3.8)$$

By Hölder's inequality then for any p in $(0, 1)$:

$$n\hat{\lambda}_n^L(\theta) \leq n\varepsilon p\lambda_{n\varepsilon}^L(\theta/p) + (1-p)L^{-1} \log E[e^{\theta \hat{W}_n^L / (1-p)}], \quad (3.9)$$

with \tilde{W}^L as in Hypothesis 1(iv). According to Hypothesis 1, for any $p \in (0, 1)$ we can make the second term of the right hand side of (3.9) as small as we like by choosing ε sufficiently small then L sufficiently large. Thus we can repeat the steps (3.4), (3.5) and (3.6) with ε and p fixed, take the limits $t \rightarrow \infty$ then $\varepsilon \rightarrow 0$ to obtain

$$\limsup_{L \rightarrow \infty} L^{-1} \log P[\sup_{t > 0} W_t^L > Lb] \leq \limsup_{L \rightarrow \infty} L^{-1} \log P[\sup_{n > 0} \hat{W}_n^L > Lb] \quad (3.10)$$

$$\leq p \sup_{t > 0} -t\lambda_t^*(b/t), \quad (3.11)$$

since $(p\lambda_t(\cdot/p))^* = p\lambda_t^*(\cdot)$, and finally let $p \nearrow 1$ to get the stated result. ■

Recall for $t > 0, \theta \in \mathbf{R}$ the definition $\Lambda_t(\theta) = t\lambda_t(\theta/t)$.

Lemma 1 For $n \in \mathbf{N}, r \in T$

$$\Lambda_r(\theta) \geq \Lambda_{nr}(\theta). \quad (3.12)$$

Proof: This follows from Hölder's inequality and the stationarity of the increments of W . For $(\mu_i)_{1 \leq i \leq n}$, with $\mu_i > 0$ and $\sum_{i=1}^n \mu_i = 1$

$$E[e^{\theta W_t/t}] \leq \prod_{i=1}^n E[e^{\theta W_{\mu_i t}/(\mu_i t)}]^{\mu_i} \quad (3.13)$$

and hence

$$\Lambda_t(\theta) \leq \sum_{i=1}^n \mu_i \Lambda_{\mu_i t}(\theta), \quad (3.14)$$

from which (3.12) follows by taking $\mu_i = n^{-1}$ and $t = nr$. ■

Proof of Theorem 2:

$$t\lambda_t^*(0) = \sup_{\theta} -t\lambda_t(\theta) = \sup_{\theta} -t\lambda_t(\theta/t) = \Lambda_t^*(0), \quad (3.15)$$

so that $I(0) = \inf_t \Lambda_t^*(0)$.

For $T = \mathbf{Z}_+$ observe from Lemma 1 that $\Lambda_1(\theta) \geq \Lambda_t(\theta)$ and hence $\Lambda_1^*(0) \leq \Lambda_t^*(0)$. Hence $I(0) = \Lambda_1^*(0) = \lambda_1^*(0)$.

For $T = \mathbf{R}_+$

$$\inf_t \Lambda_t^* \geq (\sup_t \Lambda_t)^* = \Lambda^*, \quad (3.16)$$

the second equality following from Hypothesis 2 and Lemma 1. To obtain the reverse inequality we specialize to a monotonic sequence. We bound above by $\inf_t \Lambda_t^* \leq \Lambda_* := \inf_n \Lambda_{1/2^n}^*$. By Lemma 1, $n \rightarrow \Lambda_{1/2^n}$ is increasing: so $n \rightarrow \Lambda_{1/2^n}^*$ is decreasing to its limit, namely Λ_* . Noting the equality

$$(\Lambda_*)^* = (\inf_n \Lambda_{1/2^n}^*)^* = \sup_n \Lambda_{1/2^n}^{**} = \sup_n \Lambda_{1/2^n} = \Lambda, \quad (3.17)$$

we see that as a limit of convex functions Λ_* is convex and so $\Lambda_* = (\Lambda_*)^{**} = \Lambda^*$, and so we are done. ■

According to Hypothesis 1(ii), λ_t and λ are differentiable, so the convergence of λ_t to λ implies the pointwise convergence of λ_t' to λ' . (See, for example, Lemma IV.6.3 of [11]).

Proof of Theorem 3: Define

$$\beta(t) := t\lambda_t'(\delta) \quad (3.18)$$

Since $\lambda_t'(\delta) \rightarrow \lambda'(\delta) > 0$ as $t \rightarrow \infty$, $t \mapsto \beta(t)$ is increasing for t sufficiently large and $\lim_{t \rightarrow \infty} \beta(t) = +\infty$. Set

$$\tau(b) := \sup\{t \in T \mid \beta(t) \leq b\}. \quad (3.19)$$

$\text{Ran}(\beta) \ni b \mapsto \tau(b)$ is increasing and $\lim_{b \rightarrow \infty} \tau(b) = +\infty$. By definition of the Legendre-Fenchel transform of λ_t^* and (3.18)

$$t\lambda_t^*(b/t) - \delta b \geq t\lambda_t^*(\beta(t)/t) - \delta\beta(t) = -t\lambda_t(\delta). \quad (3.20)$$

We obtain upper bounds for $\limsup_{b \rightarrow \infty} (I(b) - \delta b)$, first for $T = \mathbf{R}_+$, then for $T = \mathbf{Z}_+$. We then show these are equal to a lower bound for $\liminf_{b \rightarrow \infty} (I(b) - \delta b)$.

Upper Bound: ($T = \mathbf{R}_+$) $\text{Ran}(\tau) = \mathbf{R}_+$, and so for any $b \in \mathbf{R}_+$,

$$I(b) - \delta b = \inf_t t \lambda_t^*(b/t) - \delta b \quad (3.21)$$

$$\leq \tau(b) \lambda_{\tau(b)}^*(b/\tau(b)) - \delta b \quad (3.22)$$

$$= -\tau(b) \lambda_{\tau(b)}(\delta), \quad (3.23)$$

the last equality following from (3.20) because $\beta(\tau(b)) = b$ for $T = \mathbf{R}_+$. Since $\tau(b) \rightarrow \infty$ as $b \rightarrow \infty$, then by Hypothesis 3,

$$\limsup_{b \rightarrow \infty} (I(b) - \delta b) \leq \nu. \quad (3.24)$$

($T = \mathbf{Z}_+$) In this case $\text{Ran}(\beta)$ is a discrete set, but the conclusion (3.24) holds provided we take the limit along $\text{Ran}(\beta)$, since $\beta \circ \tau$ acts as the identity there. But for any $b \in \mathbf{R}_+$ we have

$$\tau(b) \lambda_{\tau(b)}^*(b/\tau(b)) - \delta b = \tau(b) \lambda_{\tau(b)}^*(\beta(\tau(b))/\tau(b)) - \delta \beta(\tau(b)) + E_b \quad (3.25)$$

where

$$E_b = \tau(b) \left(\lambda_{\tau(b)}^*(b/\tau(b)) - \lambda_{\tau(b)}^*(\beta(\tau(b))/\tau(b)) \right) - \delta (b - \beta(\tau(b))) \quad (3.26)$$

$$\leq (b - \beta(\tau(b))) (\delta_b - \delta) \quad (3.27)$$

where

$$\delta_b = (\lambda_{\tau(b)}^*)'(b/\tau(b)). \quad (3.28)$$

Here we have used the fact that since λ_t is strictly convex, λ_t^* is differentiable (see Theorem 26.3 in [27]).

The proof is complete if we can show that $\limsup_{b \rightarrow \infty} E_b \leq 0$, since then $\limsup_{b \rightarrow \infty} (I(b) - \delta b) \leq \limsup_{b \rightarrow \infty} -\tau(b) \lambda_{\tau(b)}(\delta) = \nu$. Note

$$\beta(\tau(b)) \leq b \leq \beta(\tau(b) + 1) \quad (3.29)$$

(for sufficiently large b) and the relations

$$\tau(b) \lambda'_{\tau(b)}(\delta_b) = b; \quad \tau(b) \lambda'_{\tau(b)}(\delta) = \beta(\tau(b)); \quad (\tau(b) + 1) \lambda'_{\tau(b)+1}(\delta) = \beta(\tau(b) + 1) \quad (3.30)$$

from which it follows that $\delta_b \geq \delta$ since λ'_t is increasing. (This means E_b is non-negative). Combining these gives

$$\frac{\beta(\tau(b) + 1) - \beta(\tau(b))}{\tau(b)} \geq \frac{b - \beta(\tau(b))}{\tau(b)} = \lambda'_{\tau(b)}(\delta_b) - \lambda'_{\tau(b)}(\delta) \geq 0. \quad (3.31)$$

We now proceed under Hypothesis 3(ii) Since, from (3.30),

$$\lim_{b \rightarrow \infty} \beta(\tau(b))/\tau(b) = \lim_{b \rightarrow \infty} \lambda'_{\tau(b)}(\delta) = \lambda'(\delta), \quad (3.32)$$

then by (3.31)

$$\lim_{b \rightarrow \infty} (\lambda'_{\tau(b)}(\delta_b) - \lambda'_{\tau(b)}(\delta)) = 0. \quad (3.33)$$

Hence

$$\lim_{b \rightarrow \infty} \delta_b = \delta; \quad (3.34)$$

for if not then $\delta_{b_i} > \delta + \varepsilon$ for some $\varepsilon > 0$ and subsequence $b_i \rightarrow \infty$. so $\limsup_{i \rightarrow \infty} \lambda'_{\tau(b_i)}(\delta_{b_i}) - \lambda'_{\tau(b_i)}(\delta) \geq \limsup_{i \rightarrow \infty} \lambda'_{\tau(b_i)}(\delta + \varepsilon) - \lambda'_{\tau(b_i)}(\delta) = \lambda'(\delta + \varepsilon) - \lambda'(\delta) > 0$, since λ is strictly convex, in contradiction with (3.33). Finally,

$$0 \leq b - \beta(\tau(b)) \leq \beta(\tau(b) + 1) - \beta(\tau(b)) = (\tau(b) + 1)\lambda'_{\tau(b)+1}(\delta) - \tau(b)\lambda'_{\tau(b)}(\delta) \quad (3.35)$$

which is bounded according to Hypothesis 3(ii). Combining with (3.34), then $\lim_{b \rightarrow \infty} E_b = 0$ as required.

Alternatively, under Hypothesis 3(ii'), $0 \leq \delta_b - \delta \leq k(b - \beta(\tau(b)))$ for some $k > 0$ independent of b and so

$$0 \leq E_b \leq k(b - \beta(\tau(b)))^2 / \tau(b) \quad (3.36)$$

$$\leq (\beta(\tau(b) + 1) - \beta(\tau(b)))^2 / \tau(b) \quad (3.37)$$

$$= ((\tau(b) + 1)\lambda'_{\tau(b)+1}(\delta) - \tau(b)\lambda'_{\tau(b)}(\delta))^2 / \tau(b) \quad (3.38)$$

which goes to 0 as $b \rightarrow \infty$ by (2.19).

Lower Bound: ($T = \mathbf{R}_+$ or \mathbf{Z}_+) Suppose first that $\inf_t t\lambda_t^*(b/t)$ is attained at $\hat{\tau}(b)$.

$$\inf_t t\lambda_t^*(b/t) - \delta b = \hat{\tau}(b)\lambda_{\hat{\tau}(b)}^*(b/\hat{\tau}(b)) - \delta b \quad (3.39)$$

$$\geq -\hat{\tau}(b)\lambda_{\hat{\tau}(b)}(\delta) \quad (3.40)$$

by (3.20) and so

$$\liminf_{b \rightarrow \infty} (I(b) - \delta b) \geq \nu \quad (3.41)$$

provided $\hat{\tau}(b) \rightarrow \infty$ as $t \rightarrow \infty$. But if this is not the case, $\hat{\tau}(b)$ is bounded. Thus we obtain a contradiction with the upper bound (3.24) if we can show that

$$\lim_{b \rightarrow \infty} (t\lambda_t^*(b/t) - \delta b) = +\infty \quad (3.42)$$

for any fixed t . But

$$\frac{d}{db} (t\lambda_t^*(b/t) - \delta b) = (\lambda_t^*)'(b/t) - \delta. \quad (3.43)$$

By (3.20) this is zero for $b = \beta(t)$ and strictly increasing on $(\beta(t), \infty)$ because λ_t^* is strictly convex due to the smoothness of λ_t . Thus (3.42) holds, as required.

If $\inf_t t\lambda_t(b/t)$ is not attained, then we can repeat the above arguments replacing $\hat{\tau}(b)$ with $\hat{\tau}_\varepsilon(b)$ for which the infimum is approximated to within ε , uniformly in b , then take $\varepsilon \rightarrow 0$ at the end. ■

Proof of Theorem 4: Since W_t and $W_{t+t'} - W_t$ are non-negatively correlated and $w \mapsto e^{\theta w}$ is non-decreasing for $\theta \geq 0$,

$$E[e^{\theta W_{t+t'}}] = E[e^{\theta W_t} e^{\theta(W_{t+t'} - W_t)}] \geq E[e^{\theta W_t}] E[e^{\theta(W_{t+t'} - W_t)}] = E[e^{\theta W_t}] E[e^{\theta W_{t'}}], \quad (3.44)$$

the last equality being due to the stationarity of the increments of W . Thus $t \mapsto t\lambda_t(\theta)$ is superadditive. Applying the sub-additivity theorem (see Lemma 6.1.11 of [8]) to $-t\lambda_t(\delta)$ we obtain:

$$\lim_{t \rightarrow \infty} \lambda_t(\delta) = \sup_{t > 0} \lambda_t(\delta) = \lambda(\delta) = 0. \quad (3.45)$$

Thus $\lambda_t(\delta)$ is non-positive for all t and so $\lim_{t \rightarrow \infty} t\lambda_t(\delta)$, if it exists, is also non-positive. ■

4 Applications and Examples.

4.1 Heterogeneous Superpositions

We examine $I(0)$ and ν for the class of heterogeneous superpositions described in the section 2. Recall there are $L = \sum_j L_j$ sources in total in the superposition, L_j of type j , each having backward arrival process $A_{j,t}$. We set $c_{j,t}(\theta) = t^{-1} \log E[e^{\theta A_{j,t}}]$ and assume the existence of the limits $c_j(\theta) = \lim_{t \rightarrow \infty} c_{j,t}(\theta)$ and $p_j = \lim_{L \rightarrow \infty} L_j/L$. The service rate is sL so $\lambda_t(\theta) = \sum_j p_j c_{j,t}(\theta) - s\theta$.

First, we examine $I(0)$. Consider the case $T = \mathbf{R}_+$. We assume the existence of the limits $C_i(\theta) = \lim_{t \rightarrow 0} t c_{i,t}(\theta/t)$. Then

$$\Lambda(\theta) = \sum_j p_j C_j(\theta) - s\theta. \quad (4.1)$$

Thus

$$I(0) = \Lambda^*(0) = \sum_j p_j C_j^*(s_j^{(0)}) \quad (4.2)$$

where $s_j^{(0)} = C_j'(\hat{\theta})$, $\hat{\theta}$ being the unique solution of the equation $\sum_j p_j C_j'(\hat{\theta}) = s$. For $T = \mathbf{Z}_+$, the same formulae hold, but with the C_j replaced by $c_{j,1}$.

Second, we examine ν . δ is the unique solution of the equation $\lambda(\delta) = 0$, i.e.

$$\sum_j p_j c_j(\delta) - s\delta = 0 \quad (4.3)$$

and so

$$\nu = \sum_j p_j \nu_j \quad (4.4)$$

where

$$\nu_i = - \lim_{t \rightarrow \infty} t (c_{j,t}(\delta) - s_j^{(\infty)} \delta), \quad (4.5)$$

and $s_j^{(\infty)} = c_j(\delta)/\delta$. Note $s_j^{(\infty)} > s_j^{(0)}$ due to the convexity of c_j , since $c_j(0) = 0$. $s_j^{(0)}$ and $s_j^{(\infty)}$ are the usual effective bandwidths of sources of type j for zero and infinite buffers respectively. (See, for example, [19] for further details).

4.2 Superpositions of Gaussian Arrival Processes.

In this section we take $W_t^L = A_t^L - sLt$ where for each L , A_t^L is an L -fold superposition of independent copies of A_t : a zero-mean Gaussian process with stationary increments and variance σ_t^2 . We make the following Hypotheses concerning A

Hypothesis 4

$$(i) \limsup_{r \rightarrow 0} r^{-1} E[\sup_{0 < r' < r} |A_{r'}|] < \infty.$$

The following limits exist as extended real numbers:

$$(ii) \sigma^2 := \lim_{t \rightarrow \infty} \sigma_t^2 / t.$$

$$(iii) \hat{\sigma}_0^2 := \lim_{t \rightarrow 0} \sigma_t^2 / t^2.$$

$$(iv) \hat{\sigma}_\infty^2 := \lim_{t \rightarrow \infty} (\sigma_t^2 - t\sigma^2).$$

Proposition 1 Under Hypothesis 4, Theorems 1, 2 and 3 hold with

$$\delta = 2s/\sigma^2; \quad I(0) = \frac{s^2}{2\hat{\sigma}_0^2} \quad \text{and} \quad \nu = -\frac{2s^2\hat{\sigma}_\infty^2}{\sigma^4}. \quad (4.6)$$

Proof: Hypothesis 4(ii) means that λ exists, and

$$\lambda_t(\theta) = \theta^2 \sigma_t^2 / (2t) - s\theta \quad \text{and} \quad \lambda(\theta) = \theta^2 \sigma^2 / 2 - s\theta \quad (4.7)$$

are clearly convex and differentiable. If (i) holds, then $a_r := E[\sup_{0 < r' < r} |A_{r'}|] < ra$ for r sufficiently small and some $a > 0$. So by Borell's inequality (see [8], Exercise 5.2.14),

$$E[e^{\theta \sup_{0 < r' < r} |A_{r'}|}] < e^{ra\theta} + 2\sqrt{\pi}\theta r e^{(\theta^2 + a)/2} \quad (4.8)$$

which goes to 1 as $r \rightarrow 0$. Hence Hypothesis 1(iv) is satisfied and Theorem 1 holds.

From (4.7) the decay rate is $\delta = 2s/\sigma^2$. Under (iii) the limit

$$\Lambda(\theta) = \lim_{t \rightarrow 0} t \lambda_t(\theta/t) = \lim_{t \rightarrow 0} \frac{\theta^2 \sigma_t^2}{2t^2} - \theta s = \theta^2 \hat{\sigma}_0^2 / 2 - \theta s \quad (4.9)$$

exists. Thus Hypothesis 2 is satisfied and $I(0) = \Lambda^*(0) = s^2/(2\hat{\sigma}_0^2)$.

Finally, under (iv) the limit

$$-\nu = \lim_{t \rightarrow \infty} t\lambda_t(\delta) = \lim_{t \rightarrow \infty} 2(\sigma_t^2 - \sigma^2 t)s^2/\sigma^4 = 2s^2\hat{\sigma}_\infty^2/\sigma^4 \quad (4.10)$$

exists, so Hypothesis 3(i) is satisfied. ■

The proof goes through for heterogeneous superpositions as described in sections 2 and 4.1, where A_t^L is a superposition of sums of L_j copies of independent Gaussian processes $A_{j,t}$ with mean zero and variance $\sigma_{j,t}^2$ satisfying Hypothesis 4, provided the limits $p_j = \lim_{L \rightarrow \infty} L_j/(\sum_j L_j)$ exist. In this case $\sigma_t^2 = \sum_j p_j \sigma_{j,t}^2$.

We note that when A is Brownian motion, $\hat{\sigma}_0 = +\infty$ and $\hat{\sigma}_\infty = 0$, so $I(0) = \nu = 0$. (Compare with the discussion in section 4.5).

4.3 Ornstein-Uhlenbeck Arrival Processes

An example where the workload is modelled by a Gaussian process with stationary increments is the following. Consider a queue with constant service rate, for which the workload W_t is the position component of a stationary Ornstein-Uhlenbeck process with added negative drift. Such an arrival process has been proposed by Norros *et al* [25] as a model of continuous correlated arrivals. It arises as the heavy traffic limit of superposed 2-state markov fluid sources under suitable rescaling of time and mean activity (see [22]).

We consider the stationary Ornstein-Uhlenbeck velocity process $(V_t, t \in \mathbf{R}_+)$, defined to be the solution of the stochastic differential equation

$$dV_t = -V_t dt + \sqrt{2}(s/\kappa) dB(t) \quad (4.11)$$

where V_0 is normally distributed with zero mean and variance $(s/\kappa)^2$. Here B is standard Brownian motion, $\kappa > 0$ is a load parameter (the case $\kappa = 0$ corresponding to unit load), and $s > 0$ can be viewed as a service rate. The corresponding position process (with zero initial condition) is

$$A_t = \int_0^t V_s ds, \quad (4.12)$$

and the workload is

$$W_t = A_t - st. \quad (4.13)$$

W_t is Gaussian with mean $-st$ and variance

$$\sigma_t^2 = 2(s/\kappa)^2(t + e^{-t} - 1). \quad (4.14)$$

Hence

$$\lambda_t(\theta) = \frac{\theta^2 \sigma_t^2}{2t} - s\theta, \quad \text{and} \quad \lambda(\theta) = \frac{\theta^2 s^2}{\kappa^2} - s\theta. \quad (4.15)$$

This gives $\sigma^2 = 2s^2/\kappa^2$, $\delta = \kappa^2/s$, $\hat{\sigma}_0^2 = s^2/\kappa^2$ and $\hat{\sigma}_\infty^2 = -2s^2/\kappa^2$. Thus items (ii), (iii) and (iv) of Hypothesis 4 is satisfied and

$$I(0) = \kappa^2/2 \quad \text{and} \quad \nu = \kappa^2. \quad (4.16)$$

Item (i) is satisfied if $E[\sup_{0 < t < 1} |V_t|] < \infty$: this is shown in section 3.3 of [10].

In fact we can easily calculate $I(b)$ numerically. Normalizing b by s , a routine calculation yields

$$t\lambda_t^*(sb/t) = \frac{s^2(b+t)^2}{2\sigma_t^2} = \kappa^2 \frac{(t+b)^2}{4(t-1+e^{-t})}. \quad (4.17)$$

We can also perform the same calculations for heterogeneous superpositions. Arrivals of type j are Ornstein-Uhlenbeck position processes with mean 0 and variance

$$\sigma_{j,t}^2 = 2(s/\kappa_j)^2(r_j t + e^{-r_j t} - 1), \quad (4.18)$$

and occur with limiting proportion p_j in the superposition. Here we have included a possible time rescalings r_j on each process. Thus the superposition has variance $\sigma_t^2 = \sum_j p_j \sigma_{j,t}^2$, and the analysis of the previous section gives:

$$\sigma^2 = 2s^2 \sum_j p_j r_j \kappa_j^{-2}. \quad (4.19)$$

$$\delta^{-1} = s^{-1} \sum_j p_j r_j \kappa_j^{-2} \quad (4.20)$$

$$I(0) = \left(2 \sum_j p_j r_j^2 \kappa_j^{-2} \right)^{-1} \quad (4.21)$$

$$\nu = \frac{\sum_j p_j r_j^2 \kappa_j^{-2}}{\left(\sum_j p_j r_j \kappa_j^{-2} \right)^2} \quad (4.22)$$

$$t\lambda_t^*(sb/t) = \frac{(b+t)^2}{4 \sum_j p_j \kappa_j^{-2} (r_j t + e^{-r_j t} - 1)}. \quad (4.23)$$

In Figure 2 the curve of $b \mapsto -I(sb)$ is plotted for two types with $r_1 = \kappa_1 = 1$, $r_2 = \kappa_2^2 = 2$ and $p_1 = p_2 = 1/2$. The curve lies between those obtained for homogeneous arrivals of each type separately: these are also plotted.

4.4 Markov Additive Arrival Processes.

In this section we obtain an expression for ν in the case that the increments of the workload W occur at integer times and are distributed according to the state of an underlying Markov

process X describing the configuration of the source of the arrivals. (Specifically, one could consider W to be the single source workload in a homogeneous superposition described in section 2; the corresponding results for heterogeneous superpositions follow from section 4.1). A convenient description for this is that of a Markov Additive Process.

To be precise, let $X = (X_t)_{t \in \mathbf{Z}_+}$ be a stationary ergodic Markov process on a state space E (with σ -field \mathcal{E}), and adjoin to it an additive component $W = (W_t)_{t \in \mathbf{Z}_+}$ with $W_0 = 0$ such that (X, W) is a Markov process on the state space $E \times \mathbf{R}_+$. Furthermore, for each $t \in \mathbf{N}$ the joint distribution of the increment $Z_{t+1} := W_{t+1} - W_t$ and X_{t+1} , conditioned on $(X_{t'}, W_{t'})_{0 \leq t' \leq t}$ depends only on X_t . This dependence can be expressed through the kernel

$$P(x, G \times B) := P[X_{t+1} \in G, Z_{t+1} \in B \mid X_t = x], \quad (4.24)$$

for $G \in \mathcal{E}$ and B a Borel set of \mathbf{R}_+ .

For $\theta \in \mathbf{R}$ define the transformed kernel $\hat{P}(\theta)$ by

$$\hat{P}(x, G; \theta) := \int P(x, G \times dz) e^{\theta z}, \quad (4.25)$$

and denote by \hat{P}^t its t -fold convolution. A technical recurrence condition for the kernel P (eq. (3.1) of [18]) is required for what follows, and we assume it to be satisfied. The main technical result we require concerning the kernel \hat{P} is an extension of the standard Perron-Frobenius to non-discrete state spaces: (see Lemma 3.1. and Lemma 3.4 of [18] and Theorem III.10.1 of [15]). Let q denote the stationary distribution of X .

Proposition 2 *λ is strictly convex and essentially smooth. For all θ in the effective domain of λ , $e^{\lambda(\theta)}$ is the simple maximal eigenvalue of $\hat{P}(\theta)$. The corresponding (right) eigenfunction $r(\cdot; \theta)$ and Radon-Nikodym derivative $d\ell(\cdot; \theta)/dq$ of the (left) eigenmeasure $\ell(\cdot; \theta)$ are uniformly bounded and positive. With the normalization $\int \ell(dx; \theta) r(x; \theta) = 1$*

$$\hat{P}^t(x, G; \theta) = r(x; \theta) \ell(G; \theta) e^{t\lambda(\theta)} \left(1 + \mathbf{O}(\varepsilon(\theta)^t) \right), \quad (4.26)$$

where $0 < \varepsilon(\theta) < 1$.

Corollary 5

$$\nu = - \lim_{t \rightarrow \infty} t \lambda_t(\delta) = - \log \left(\ell(E; \delta) \int q(dx) r(x; \delta) \right). \quad (4.27)$$

Proof: This follows since $t \lambda_t(\delta) = \log \int q(dx) \hat{P}^t(x, E; \delta)$ and $\lambda(\delta) = 0$. ■

To calculate $I(0)$ we note that

$$\lambda_1(\theta) = \log \int q(dx) P(x, dy \times dz) e^{\theta z}. \quad (4.28)$$

In the special case, frequent in modelling, that the increment Z_t is a non-random function ζ of X_t when both are conditioned on X_{t-1} we have

$$P(x, dy \times dz) = R(x, dy)\delta_{\zeta(y)}(dz) \quad (4.29)$$

where R is the transition kernel for X , and so (4.28) reduces to

$$\lambda_1(\theta) = \int q(dy)e^{\theta\zeta(y)}. \quad (4.30)$$

Note that the complexity of the calculation of ν and $I(0)$ (and indeed the whole curve of $I(b)$) is independent of the number of sources L in the superposition.

Application: two-state markov chain. We consider arrivals to be generated by a discrete time markov chain on two states: on and off. In the on state an arrival of unit length is generated; in the off state no arrival is generated. Transitions from off to on occur with probability a ; the reverse transition with probability d . The arrivals are serviced at constant rate $s < 1$. Thus, within the general framework above we have $E = \{0, 1\}$, $Z_t = \zeta(X_t)$ with $\zeta(0) = -s, \zeta(1) = 1 - s$, so that $P(x, dy \times dz) = R(x, dy)\delta_{\zeta(y)}(dz)$ where R is the transition matrix of X :

$$R = \begin{pmatrix} 1-a & a \\ d & 1-d \end{pmatrix}. \quad (4.31)$$

The stationary distribution of R is $q = (\frac{d}{a+d}, \frac{a}{a+d})$: the stability condition is $a/(a+d) < s$.

The eigenvector/eigenvalue analysis of \hat{P} yields the following. $e^{\lambda(\theta)}$ is the maximal eigenvalue of the matrix

$$\hat{P}(\theta) = R e^{\theta\zeta(\cdot)} = \begin{pmatrix} 1-a & ae^\theta \\ d & (1-d)e^\theta \end{pmatrix} e^{-s\theta}. \quad (4.32)$$

Let $y = e^\theta$ and set

$$x_\pm = (2a)^{-1} \left(y(1-d) - (1-a) \pm \sqrt{((1-a) - y(1-d))^2 + 4ady} \right). \quad (4.33)$$

The eigenvalues of $\hat{P}(\theta)$ are $v_\pm = y^{-s}(ax_\pm + 1 - a)$. Hence $\delta = \log y$ for y such that $1 = y^{-s}(ax_+ + 1 - a)$. The (unnormalised) eigenmeasures and eigenfunctions are

$$\ell_\pm(\cdot; \theta) = \ell_\pm := (d, ax_\pm) \quad \text{and} \quad r_\pm(\cdot; \theta) = r_\pm := (y, x_\pm) \quad (4.34)$$

respectively. Thus

$$e^{t\lambda_1(\theta)} = q \cdot \hat{P}^t(\cdot, E; \theta) = \frac{q \cdot r_+ \ell_+ \cdot (1, 1)v_+^t}{\ell_+ \cdot r_+} + \frac{q \cdot r_- \ell_- \cdot (1, 1)v_-^t}{\ell_- \cdot r_-}, \quad (4.35)$$

and

$$\nu = -\log \left(\frac{q \cdot r_+ \ell_+ \cdot (1, 1)}{\ell_+ \cdot r_+} \right) = -\log \frac{(ax_+ + dy)(ax_+ + d)}{(ax_+^2 + dy)(a + d)}, \quad (4.36)$$

using the values for x_+ obtained from (4.33) with $y = e^\delta$. One sees that from the spectral decomposition (4.35) that $t \mapsto t\lambda'_t(\delta) = A_t + tB_t$ where A_t is bounded and $t(B_{t+1} - B_t) \rightarrow 0$ as $t \rightarrow \infty$. Hence Hypothesis 3(ii) is satisfied in this model, and indeed in any Markovian model for which such a differentiable spectral decomposition exists.

At $b = 0$ we find

$$I(0) = \lambda_1^*(0) = -\inf_{\theta} \lambda_1(\theta) = -\log \inf_{\theta} \frac{ae^\theta + d}{(a+d)e^{\theta s}} = -\log \frac{a\hat{x} + d}{(a+d)\hat{x}^s} \quad (4.37)$$

where $\hat{x} = sd/(a(1-s))$. This agrees with the large deviation (upper) bound according to Hui [16] for the probability of overflow at a bufferless resource (i.e. with $b = 0$).

The sign of ν can be related to the sign of the correlations of the arrivals process. One sees from (4.36) that $\text{sgn}(\nu) = \text{sgn}((x_+ - 1)(x_+ - y))$. But it is shown from Proposition 3 of [4] that $1 - a - d > 0 \implies x_+ > y > 1$ while $1 - a - d < 0 \implies y > x_+ > 1$ and $1 - a - d = 0 \implies x_+ = y > 1$. Furthermore the covariance of successive arrivals $\text{Cov}(Z_t, Z_{t+1}) = ad(1 - a - d)/(a + d)^2$. Summarizing:

$$\text{sgn}(1 - a - d) = \text{sgn}(\nu) = \text{sgn}(\text{Cov}(Z_t, Z_{t+1})). \quad (4.38)$$

Bursty sources will modelled with $a + d < 1$: successive arrivals are positively correlated. A sub-bursty markov model (i.e. with negatively correlated arrivals) has been studied numerically by Choudhury *et al* [7]. It is found that the log-loss curves are concave, and asymptotic to a straight line with *positive* intercept at $b = 0$: correspondingly, our value for ν will be negative.

Comparisons and estimates. Theorem 1 can be used as a basis for approximation of superpositions of finitely many lines: we take

$$P[Q^L > b] \approx e^{-LI(b/L)}. \quad (4.39)$$

The λ_t are using (4.35) with $y = e^\theta$. The resulting approximation is compared with simulations in three cases. Figure 3 takes $a = 0.03$, $d = 0.045$, $L = 84$ and $s = 40/84$, a superposition of highly bursty sources. In Figure 4 the parameters $a = 0.3$, $d = 0.5344$, $L = 100$ and $s = 40/100$ are chosen to make $\nu = I(0)$: the curve is very close to linear. In Figure 5 a sub-bursty case $a = 0.55$, $d = 0.825$, $L = 84$ and $s = 40/84$ is shown. In these examples, the shape of the log-loss curve is closely reproduced by the approximation, but with a shift which makes the approximation conservative (in these cases). This may well be a limitation of the first-order large deviation method. In fact the discrepancy is well within the $O(\log L)$ corrections to the large deviation estimate of $P[W_1^L > 0]$ in (4.37), so some improvement may be possible with further work involving these corrections to the first-order large deviation result. As a final numerical example we take a large superposition of extremely bursty sources: $a = .0003$, $d = .0007$, $s = 400$, $L = 1000$. For these parameters $-LI(0)/\log 10 = 9.8$ and $-L\nu/\log 10 = 20.2$: the desired loss probabilities of 10^{-9} are already obtained at $b = 0$.

4.5 Poissonian Arrivals.

We conclude with a discussion of Poissonian arrivals. For simplicity we consider homogeneous superpositions. Each of L sources is a Poisson stream of intensity κ . The service requirement of each arrival is independently distributed with distribution G . The superposition is served at rate sL . Define

$$g(\theta) := \int e^{\theta x} G(dx). \quad (4.40)$$

Then

$$\lambda_t^L(\theta) = \lambda_t(\theta) = \lambda(\theta) = t^{-1} \log e^{-s\theta t} \sum_{n=0}^{\infty} \frac{e^{-\kappa t} (\kappa t \int e^{\theta x} G(dx))^n}{n!} \quad (4.41)$$

$$= \kappa (g(\theta) - 1) - s\theta. \quad (4.42)$$

Since λ_t is independent of t then $\nu = -\lim_{t \rightarrow \infty} t\lambda_t(\delta) = 0$. Furthermore, $I(0) = 0$ since for $\theta > 0, t < 1$

$$\Lambda_t(\theta) = t\lambda(\theta/t) \quad (4.43)$$

$$= t\kappa(g(\theta/t) - 1) - s\theta \quad (4.44)$$

$$\geq t\kappa(g(\theta)^{1/t} - 1) - s\theta \quad (4.45)$$

which goes to $+\infty$ as $t \rightarrow 0$ since $g(\theta) > 1$. Similarly one has $\lim_{t \rightarrow 0} tg(\theta/t) = 0$ for $\theta \leq 0$. Thus

$$\Lambda(\theta) = \begin{cases} -s\theta & \text{if } \theta \leq 0 \\ +\infty & \text{if } \theta > 0 \end{cases} \quad (4.46)$$

and so $I(0) = \Lambda^*(0) = 0$. From this we draw the conclusion that there are no economies of scale to be obtained from a superposition of Poissonian arrivals. In contrast, Bernoulli arrivals will generally give $I(0) > \nu = 0$: take $a = 1 - d$ in the on-off model as an example. The difference in $I(0)$ between the two cases can be shown to go to zero if one constructs the Poissonian arrivals process as a continuum limit of Bernoulli arrivals.

Acknowledgements. The authors thank Neil O'Connell, John Lewis, Fergal Toomey and particularly Raymond Russell for their comments. Raymond Russell wrote the software used in the simulations of the on-off Markov model.

References

- [1] D. Anick, D. Mitra and M.M. Sondhi (1982) Stochastic theory of a data-handling system with multiple sources. *Bell Sys. Tech. J.*, 61:1872–1894
- [2] C. Borell (1975). The Brunn-Minkowski inequality in Gauss space. *Invent. Math.*, 30:205–216.

- [3] J.A. Bucklew (1990). *Large deviation techniques in decision, simulation and estimation*. Wiley, New York.
- [4] E. Buffet and N.G. Duffield (1992) Exponential upper bounds via martingales for multiplexers with Markovian arrivals. *J. Appl. Prob.* to appear.
- [5] Cheng-Shang Chang (1993). Stability, queue length and delay of deterministic and stochastic queueing networks. Preprint.
- [6] T.J. Corcoran (1994) Prediction of ATM multiplexer performance by simulation and analysis of a model of packetized voice traffic. M.Sc. Thesis, Dublin City University.
- [7] Gagan L. Choudhury, David M. Lucantoni and Ward Whitt (1993). Squeezing the most out of ATM. Preprint.
- [8] Amir Dembo and Ofer Zeitouni (1993). *Large Deviation Techniques and Applications*. Jones and Bartlett, Boston-London.
- [9] N.G. Duffield (1993). Exponential bounds for queues with Markovian arrivals. *Queueing Systems* , to appear.
- [10] N.G. Duffield and Neil O'Connell (1993). Large deviations and overflow probabilities for the general single-server queue, with applications. Preprint DIAS-APG-93-30
- [11] R.S. Ellis (1985). *Entropy, Large Deviations, and Statistical Mechanics*, Springer, New York.
- [12] A.I. Elwalid, D. Mitra and T.E. Stern (1991). Statistical multiplexing of Markov modulated sources: theory and computational algorithms. In: *Teletraffic and Datatraffic in a period of change, ITC-13*, A. Jensen & V.B. Iversen (Eds.) Elsevier Science Publishers B.V. (North-Holland).
- [13] R.J. Gibbens and P.J. Hunt (1991). Effective Bandwidths for the multi-type UAS channel *Queueing Systems*, 9:17-28
- [14] Peter W. Glynn and Ward Whitt (1993). Logarithmic asymptotics for steady-state tail probabilities in a single-server queue. *J. Appl. Prob.*, to appear.
- [15] T.E. Harris (1963), *The theory of branching processes*. Springer, Berlin.
- [16] J.Y. Hui (1988). Resource allocation for broadband networks. *IEEE J. Selected Areas in Commun.* 6:1598-1608
- [17] J.Y. Hui (1990). *Switching and traffic theory for integrated broadband networks*. Kluwer. Boston .

- [18] I. Iscoe, P. Ney and E. Nummelin (1985). Large deviations of uniformly recurrent Markov additive processes. *Adv. in Appl. Math.* 6:373–412
- [19] F.P. Kelly (1991). Effective bandwidths at multi-type queues. *Queueing Systems* 9:5–16
- [20] G. Kesidis, J. Walrand and C.S. Chang (1993). Effective bandwidths for multiclass Markov fluids and other ATM Sources. Preprint.
- [21] L. Kosten (1988) Stochastic Theory of data handling systems with groups of multiple sources. *Proc. 2nd Int. Symp. on the Performance of Computer Communication Systems*, eds. H. Rudin & W. Bux, North-Holland.
- [22] V. Kulkarni and T. Rolski (1993). Fluid Model Driven by an Ornstein-Uhlenbeck Process. Preprint.
- [23] Will E. Leland, Murad S. Taqqu, Walter Willinger and Daniel V. Wilson (1993). Ethernet traffic is self-similar: stochastic modelling of packet traffic data. Preprint.
- [24] J.T. Lewis and C.-E. Pfister (1994). Thermodynamic probability theory: some aspects of large deviations. *Theor. Prob. Appl.*, to appear.
- [25] I. Norros, J.W. Roberts, A. Simonian and J. Virtamo (1991). The superposition of variable bitrate sources in ATM multiplexers. *IEEE J. Selected Areas in Commun.* 9:378–387.
- [26] Carsten Rasmussen, Jens H. Sorensen, Kenn S. Kvols and Soren B. Jacobsen (1991). Source-independent call acceptance procedures in ATM networks. *IEEE J. Selected Areas in Commun.* 9:351–358.
- [27] R. Tyrrell Rockafellar (1970) *Convex Analysis*. Princeton University Press, Princeton.
- [28] K. Sriram and W. Whitt (1986). Characterizing superposition arrival processes in packet multiplexers for voice data. *IEEE J. Select. Areas Commun.* 4:833–846
- [29] G. de Veciana, C. Courcoubetis and J. Walrand (1993) Decoupling bandwidths for networks: a decomposition approach to resource management. Preprint.
- [30] W. Whitt (1992) Tail probabilities with statistical multiplexing and effective bandwidths in multi-class queues. Preprint.

Figure 2: Ornstein-Uhlenbeck superpositions

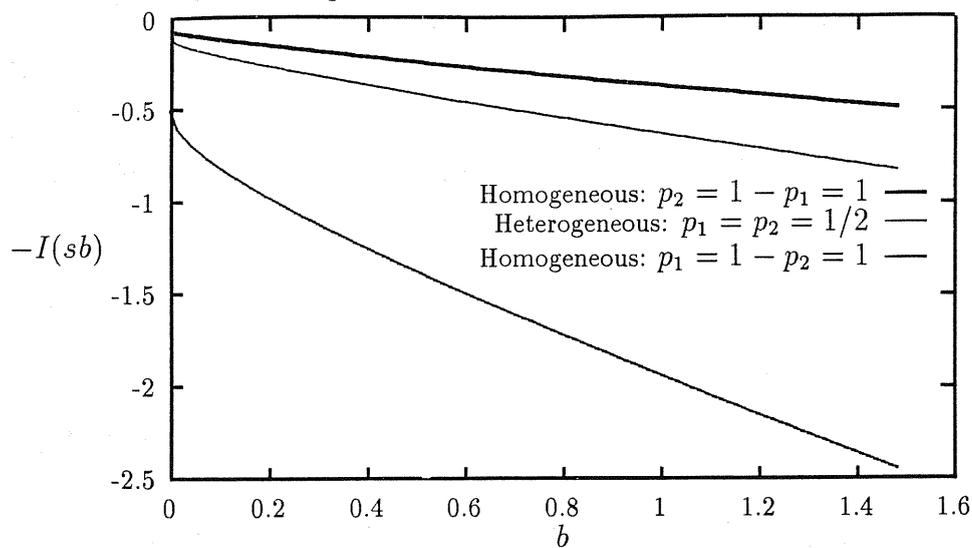


Figure 3: $a + d < 1$ with $a = 0.03$, $d = 0.045$, $s = 40$ and $L = 84$.

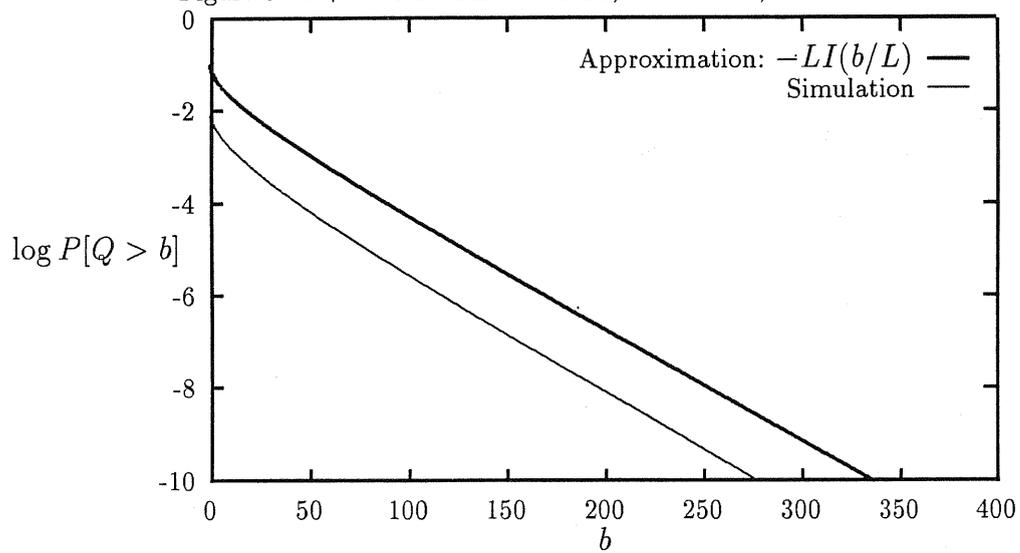


Figure 4: $I(0) = \nu$ with $a = 0.3$, $d = .5344$, $s = 40$ and $L = 100$

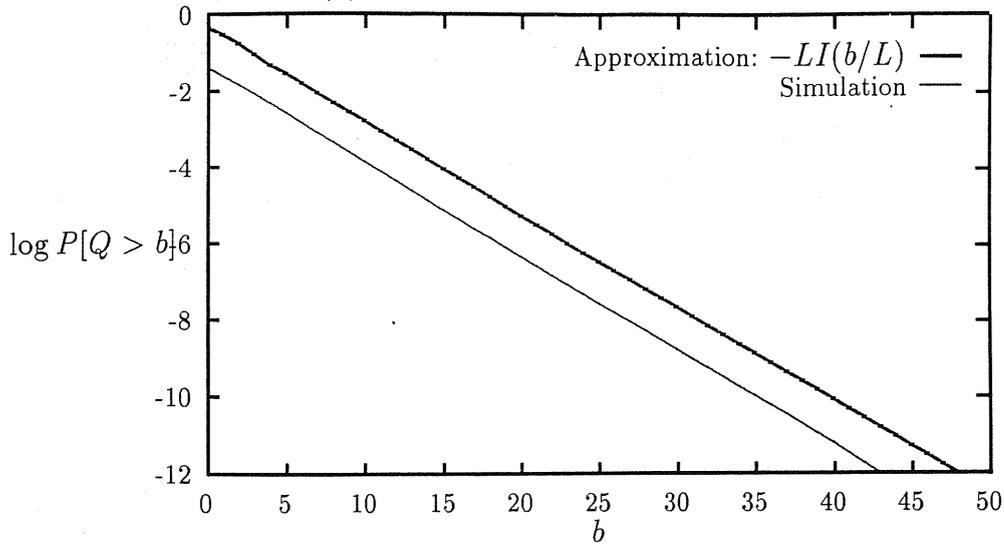


Figure 5: $a + d > 1$ with $a = 0.55$, $d = 0.825$, $s = 40$ and $L = 84$.

