Centrum voor Wiskunde en Informatica

REPORT *RAPPORT*

PNA

Probability, Networks and Algorithms

Probability, Networks and Algorithms

Heavy-traffic approximations for linear networks operating under α-fair bandwidth-sharing policies

P.M.D. Lieshout, S.C. Borst, M.R.H. Mandjes

REPORT PNA-R0609 AUGUST 2006

Centrum voor Wiskunde en Informatica (CWI) is the national research institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organisation for Scientific Research (NWO).
CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

ISSN 1386-3711

Heavy-traffic approximations for linear networks operating under alpha-fair bandwidth-sharing policies

ABSTRACT
We consider the flow-level performance of a linear network supporting elastic traffic, where the service capacity is shared among the various classes of users according to a weighted alpha-fair policy. Assuming Poisson arrivals and exponentially distributed service requirements for each class, the dynamics of the user population may be described by a Markov process. While valuable stability results have been established for the family of alpha-fair policies, the distribution of the number of active users has remained intractable in all but a few special cases. In order to gain further insight in the flow-level performance in more general scenarios, we develop approximations for the mean number of users based on the assumption that one or two of the nodes experience heavy-traffic conditions. In case of just a single `bottleneck' node, we exploit the fact that this node approximately behaves as a two-class Discriminatory Processor-Sharing model. In the case that there are two nodes critically loaded, we rely on the observation that the joint workload process at these nodes is asymptotically independent of the fairness coefficient alpha, provided all classes have equal weights. In particular, the distribution of the joint workload process is roughly equal to that for an unweighted Proportional Fair policy, which is exactly known. In both cases, the numbers of users at non-bottleneck nodes can be approximated by that in an M/M/1 queue with reduced service capacity. Extensive numerical experiments indicate that the resulting approximations tend to be reasonably accurate across a wide range of parameters, even at relatively moderate load values. The approximations for the mean number of users also provide useful estimates for the mean transfer delays and user throughputs.

2000 Mathematics Subject Classification: 60K25; 68M20
1998 ACM Computing Classification System: D.4.8 [Performance]: Queueing Theory
Keywords and Phrases: Simultaneous resource possession; alpha-fair sharing; Heavy traffic; Linear network; Flow-level performance
Note: This research has been funded by the Dutch Bsik/BRICKS (Basic Research in Informatics for Creating the Knowledge Society) project.

Heavy-Traffic Approximations for Linear Networks Operating under α-fair Bandwidth-Sharing Policies *

P. Lieshout, S.C. Borst and M. Mandjes

CWI
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
Email: {lieshout, sem, michel}@cwi.nl

August 14, 2006

Abstract

We consider the flow-level performance of a linear network supporting elastic traffic, where the service capacity is shared among the various classes of users according to a weighted alpha-fair policy. Assuming Poisson arrivals and exponentially distributed service requirements for each class, the dynamics of the user population may be described by a Markov process. While valuable stability results have been established for the family of alpha-fair policies, the distribution of the number of active users has remained intractable in all but a few special cases. In order to gain further insight in the flow-level performance in more general scenarios, we develop approximations for the mean number of users based on the assumption that one or two of the nodes experience heavy-traffic conditions.

In case of just a single 'bottleneck' node, we exploit the fact that this node approximately behaves as a two-class Discriminatory Processor-Sharing model. In the case that there are two nodes critically loaded, we rely on the observation that the joint workload process at these nodes is asymptotically independent of the fairness coefficient alpha, provided all classes have equal weights. In particular, the distribution of the joint workload process is roughly equal to that for an unweighted Proportional Fair policy, which is exactly known. In both cases, the numbers of users at non-bottleneck nodes can be approximated by that in an M/M/1 queue with reduced service capacity. Extensive numerical experiments indicate that the resulting approximations tend to be reasonably accurate across a wide range of parameters, even at relatively moderate load values. The approximations for the mean number of users also provide useful estimates for the mean transfer delays and user throughputs.

Categories and Subject Descriptors
D.4.8 [**Performance**]: Queueing Theory

General Terms
Performance

Keywords
Simultaneous resource possession, Alpha-fair sharing, Heavy traffic, Linear network, Flow-level performance

*This research has been funded by the Dutch BSIK/BRICKS (Basic Research in Informatics for Creating the Knowledge Society) project. S.C. Borst is also affiliated with Bell Labs, Lucent Technologies, Murray Hill NJ, USA, and Eindhoven University of Technology, Eindhoven, the Netherlands. M. Mandjes is also affiliated with the Korteweg-de Vries Institute, University of Amsterdam, Amsterdam, the Netherlands, and EURANDOM, Eindhoven, the Netherlands.

1 Introduction

Over the past several years, the Processor-Sharing discipline has emerged as a useful paradigm for evaluating the flow-level performance of elastic data transfers competing for bandwidth on a single bottleneck link. Bandwidth-sharing networks as considered by Massoulié & Roberts [9] provide a natural extension for modeling the dynamic interaction among competing elastic flows that traverse several links.

It is well known that the queue length distribution in a single-server Processor-Sharing system with Poisson arrivals has a simple geometric distribution that only depends on the service requirement distribution through its mean. In contrast, the distribution of the number of active users in bandwidth-sharing networks with several nodes has remained generally intractable, even for exponentially distributed service requirements. Bonald & Massoulié [1] established the crucial result that a wide family of alpha-fair bandwidth-sharing policies as introduced by Mo & Walrand [10] achieve stability under the simple (and necessary) condition that no individual link is overloaded and $\alpha \in (0, \infty)$. The family of alpha-fair policies covers several common notions of fairness as special cases, such as max-min fairness ($\alpha \to \infty$), Proportional Fairness ($\alpha \to 1$), and maximum throughput ($\alpha \downarrow 0$). In [11] it has also been shown that the case $\alpha = 2$, with additional class weights set inversely proportional to the respective round trip times, provides a reasonable modeling abstraction for the bandwidth sharing realized by TCP (Transmission Control Protocol) in the Internet.

The above-mentioned stability results imply that flow-level performance measures such as transfer delays are finite provided that no individual link is overloaded. However, the derivation of the exact transfer delays and actual user throughputs has proven largely elusive, except in the special case of an unweighted Proportional Fair bandwidth-sharing policy in certain topologies, such as linear networks. In particular, it is not well understood how the flow-level performance measures depend on the specific choice of the fairness coefficient alpha and the possible additional weight factors associated with the various classes.

In order to gain further insight in the latter issues, we develop in the present paper approximations for the mean number of users in linear networks operating under alpha-fair bandwidth-sharing policies. The approximations are based on the assumption that one or two of the nodes experience heavy-traffic conditions. In case of just a single 'bottleneck' node, we exploit the fact that this node approximately behaves as a two-class Discriminatory Processor-Sharing queue. The mean number of users can thus be calculated from the results of Fayolle *et al.* [5]. In the case that there are two nodes critically loaded, we rely on the following two observations. First, the heavy-traffic results of Kang *et al.* [7] and Kelly & Williams [8] show that with equal class weights, the joint workload process is asymptotically independent of the fairness coefficient alpha. Second, the joint workload process for a Proportional Fair policy can be exactly computed from the known distribution of the number of users [9]. Combining these two observations, we obtain simple explicit estimates for the workloads at the two bottleneck nodes, which we also numerically validate. We then develop various approximation methods by using the latter estimates in conjunction with characterizations of invariant states from [7, 8] that relate the number of users of the various classes to the workloads at the various nodes.

Extensive numerical experiments indicate that the resulting approximations tend to be reasonably accurate across a wide range of parameters, even at relatively moderate load values. The approximations for the mean number of users also provide useful estimates for the mean transfer delays and user throughputs. In addition, the numerical results offer valuable insight into the impact of the choice of the fairness coefficient alpha and the possible additional weight factors, and how the performance impact depends on the traffic characteristics.

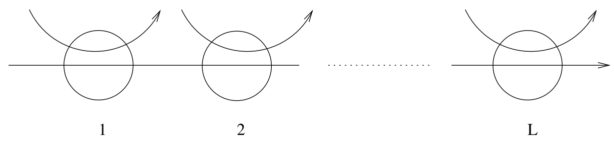

Figure 1: A linear bandwidth-sharing network.

The remainder of the paper is organized as follows. In Section 2 we provide a detailed model description and discuss some preliminaries. In Section 3 we present some results for the known distribution of the user population for a Proportional Fair policy, and use these to obtain the Laplace-Stieltjes Transform (LST) of the joint workload process at the various nodes. Section 4 reviews the heavy-traffic results of [7, 8], which provide the basis for the approximations that we develop subsequently. In Section 5 we focus on the case of a single bottleneck node, and exploit the fact that this node approximately behaves as a two-class Discriminatory Processor-Sharing model to obtain approximations for the mean number of users. Next, in Section 6 we turn the attention to a scenario with two bottleneck nodes, and invoke the principle that the joint workload process can be approximated by the known behavior for a Proportional Fair policy, provided all classes have equal weights. In conjunction with a few equivalent characterizations of invariant states from [7, 8], the latter principle is then leveraged in Section 7 to devise various approximation methods. Section 8 concludes with some final observations.

2 Model description

We consider a linear network as depicted in Figure 1. The network consists of L nodes, each with unit service rate. There are $L+1$ classes of users: each class corresponds to a specific route in the network. Class-i users require service at node i only, $i = 1, \ldots, L$, whereas class-0 users require service at all L nodes simultaneously.

We assume that class-i users arrive according to a Poisson process of rate λ_i, and have exponentially distributed service requirements with mean μ_i^{-1}, $i = 0, \ldots, L$. The arrival processes are all independent. The traffic load of class i is then $\rho_i = \lambda_i \mu_i^{-1}$. Note that the traffic load at node i is given by $\rho_0 + \rho_i$, $i = 1, \ldots, L$. Let $n = (n_0, n_1, \ldots, n_L)$ be the state of the network, with n_i representing the number of class-i users.

The network operates under a so-called alpha-fair sharing policy, as introduced in [10]. When the network is in state $n \neq 0$, the service rate x_i^* allocated to each of the class-i users is obtained by solving the optimization problem:

$$\begin{array}{ll} \max & \sum_{i=0}^{L} \kappa_i n_i \frac{x_i^{1-\alpha}}{1-\alpha} \\ \text{s.t.} & n_0 x_0 + n_i x_i \leq 1, \quad i = 1, \ldots, L, \end{array}$$

where the κ_is are non-negative class weights, and $\alpha \in (0, \infty)/\{1\}$ may be interpreted as a fairness coefficient. The cases $\alpha \downarrow 0$, $\alpha \to 1$ and $\alpha \to \infty$ correspond to allocations which achieve *maximum throughput*, *proportional fairness*, and *max-min fairness*, respectively. Let $s_i(n) := x_i^* n_i$ denote the total service rate allocated to class i. In [1] it was shown that

$$s_0(n) = \frac{(\kappa_0 n_0^{\alpha})^{1/\alpha}}{(\kappa_0 n_0^{\alpha})^{1/\alpha} + (\sum_{j=1}^{L} \kappa_j n_j^{\alpha})^{1/\alpha}}; \quad s_i(n) = 1 - s_0(n), \quad i = 1, \ldots, L, \tag{1}$$

if $n \neq 0$.

Let $N(t)$ denote the state of the network at time t. Then $N(t)$ is a Markov process with transition rates:

$$q(n, n+e_i) = \lambda_i; \quad q(n, n-e_i) = \mu_i s_i(n), \quad i = 0, \ldots, L,$$

where e_i denotes the $(i+1)$th unit vector in $\mathbb{R}^{L+1}$. Evidently, $\rho_0 + \rho_i < 1$, $i = 1, \ldots, L$, is a necessary condition for the process $N(t)$ to be ergodic. In [1] it was shown that this condition is in fact also sufficient for every $\alpha \in (0, \infty)$.

In general there are no closed-form expressions available for the steady-state distribution of $N(t)$. However, for the case $\alpha = 1$ and $\kappa_i = \kappa$ an explicit expression has been derived in [9], as will be presented in the next section.

3 Unweighted Proportional Fairness

In this section we consider the case $\alpha = 1$ and $\kappa_i = \kappa$, $i = 0, \ldots, L$. The following theorem appeared in slightly different form in [9].

Theorem 3.1 *Under the stability condition* $\max_{1 \leq i \leq L} \rho_0 + \rho_i < 1$, *the process* $N(t)$ *is reversible, with steady-state distribution given by*

$$\pi(n) = C^{-1} \binom{\sum_{i=0}^{L} n_i}{n_0} \prod_{i=0}^{L} \rho_i^{n_i}, \tag{2}$$

where the normalization constant C *equals*

$$C = \frac{(1-\rho_0)^{L-1}}{\prod_{i=1}^{L}(1-\rho_0-\rho_i)}. \tag{3}$$

The mean number of class-0 users in steady state is given by

$$\mathbb{E}(N_0) = \frac{\rho_0}{1-\rho_0}\left(1 + \sum_{i=1}^{L} \frac{\rho_i}{1-\rho_0-\rho_i}\right)$$

and for $i = 1, \ldots, L$,

$$\mathbb{E}(N_i) = \frac{\rho_i}{1-\rho_0-\rho_i}.$$

Let $W_i(t)$ denote the workload, i.e., the unfinished amount of work at node i at time t, $i = 1, \ldots, L$. Thus $W_i(t)$ consists of the remaining service requirements of all class-0 and class-i users at time t. Theorem 3.1 enables us to derive the Laplace-Stieltjes Transform (LST) of the joint distribution of $W(t) = (W_1(t), \ldots, W_L(t))$ in steady state.

Theorem 3.2 *Under the stability condition* $\max_{1 \leq i \leq L} \rho_0 + \rho_i < 1$, *the LST of* $W(t)$ *in steady state is given by*

$$\tilde{W}(r) \equiv \tilde{W}(r_1, \ldots, r_L) = \left(\frac{1 - \frac{\lambda_0}{\mu_0 + \sum_{j=1}^{L} r_j}}{1-\rho_0}\right)^{L-1} \prod_{i=1}^{L} \frac{1-\rho_0-\rho_i}{1 - \frac{\lambda_0}{\mu_0+\sum_{j=1}^{L} r_j} - \frac{\lambda_i}{\mu_i + r_i}}. \tag{4}$$

Proof: Due to the memoryless property of the exponential distribution, the residual service requirement of a class-i user is also exponentially distributed with mean μ_i^{-1}, $i = 0,\dots,L$. Therefore $W_i(t)$ is distributed as $\sum_{j=1}^{N_0(t)} B_{0,j} + \sum_{j=1}^{N_i(t)} B_{i,j}$, where $B_{i,j}$ are i.i.d. copies of an exponentially distributed variable with mean μ_i^{-1}, $i = 1,\dots,L$. Now

$$\tilde{W}(r) = \mathbb{E}\left(e^{-\sum_{i=1}^{L} r_i W_i}\right) = \mathbb{E}\left(e^{-\sum_{i=1}^{L} r_i \sum_{j=1}^{N_0} B_{0,j} - \sum_{i=1}^{L}\left(r_i \sum_{j=1}^{N_i} B_{i,j}\right)}\right).$$

Conditioning on the values of N_i, $i = 0,\dots,L$, we obtain that $\tilde{W}(r)$ equals

$$\sum_{n_0=0}^{\infty}\cdots\sum_{n_L=0}^{\infty}\pi(n)\mathbb{E}\left(e^{-\sum_{i=1}^{L} r_i \sum_{j=1}^{n_0} B_{0,j} - \sum_{i=1}^{L}\left(r_i \sum_{j=1}^{n_i} B_{i,j}\right)}\right)$$

$$= \sum_{n_0=0}^{\infty}\cdots\sum_{n_L=0}^{\infty}\pi(n)\left(\frac{\mu_0}{\mu_0 + \sum_{i=1}^{L} r_i}\right)^{n_0}\prod_{i=1}^{L}\left(\frac{\mu_i}{\mu_i + r_i}\right)^{n_i}.$$

Substituting (2) and invoking that $\rho_i = \lambda_i \mu_i^{-1}$, we obtain that $\tilde{W}(r)$ is equivalent to

$$C^{-1}\prod_{i=1}^{L}\sum_{n_i=0}^{\infty}\left(\frac{\rho_i\mu_i}{\mu_i + r_i}\right)^{n_i}\sum_{n_0=0}^{\infty}\binom{\sum_{j=0}^{L} n_j}{n_0}\left(\frac{\rho_0\mu_0}{\mu_0 + \sum_{j=1}^{L} r_j}\right)^{n_0}$$

$$= C^{-1}\prod_{i=1}^{L}\sum_{n_i=0}^{\infty}\left(\frac{\lambda_i}{\mu_i + r_i}\right)^{n_i}\left(1 - \frac{\lambda_0}{\mu_0 + \sum_{j=1}^{L} r_j}\right)^{-1-\sum_{j=1}^{L} n_j}$$

$$= C^{-1}\frac{1}{1 - \frac{\lambda_0}{\mu_0 + \sum_{j=1}^{L} r_j}}\prod_{i=1}^{L}\sum_{n_i=0}^{\infty}\left(\frac{\frac{\lambda_i}{\mu_i + r_i}}{1 - \frac{\lambda_0}{\mu_0 + \sum_{j=1}^{L} r_j}}\right)^{n_i}$$

$$= C^{-1}\frac{1}{1 - \frac{\lambda_0}{\mu_0 + \sum_{j=1}^{L} r_j}}\prod_{i=1}^{L}\frac{1}{1 - \frac{\frac{\lambda_i}{\mu_i + r_i}}{1 - \frac{\lambda_0}{\mu_0 + \sum_{j=1}^{L} r_j}}}$$

$$= \frac{1}{\left(1 - \frac{\lambda_0}{\mu_0 + \sum_{j=1}^{L} r_j}\right)(1-\rho_0)^{L-1}}\prod_{i=1}^{L}\frac{1 - \rho_0 - \rho_i}{\left(1 - \frac{\frac{\lambda_i}{\mu_i + r_i}}{1 - \frac{\lambda_0}{\mu_0 + \sum_{j=1}^{L} r_j}}\right)}. \tag{5}$$

The second equality above follows by applying the negative binomial formula: $(1-x)^{-d} = \sum_{n=0}^{\infty}\binom{d-1+n}{n}x^n$. The final equality follows by substituting (3). Rearranging (5) finally gives (4), and completes the proof. □

Remark. We now provide some interpretation for the expression for $\tilde{W}(r)$ given in Theorem 3.2. Consider an M/H$_2$/1 queue with arrival rate $\tilde{\lambda}_0 + \lambda_i$ and service requirements that are exponentially distributed with mean $1/\tilde{\mu}_0$ $(1/\mu_i)$ with probability $\frac{\tilde{\lambda}_0}{\tilde{\lambda}_0+\lambda_i}$ $\left(\frac{\lambda_i}{\tilde{\lambda}_0+\lambda_i}\right)$, where $\tilde{\lambda}_0 := \lambda_0/L$

and $\tilde{\mu}_0 := \mu_0/L$. The LST of the workload $V_i(t)$ in steady state of this M/H$_2$/1 queue is given by the well-known Pollaczek-Khinchin formula

$$\tilde{V}_i(r_i) = \frac{(1-\rho_0-\rho_i)r_i}{(\tilde{\lambda}_0+\lambda_i)\tilde{B}(r_i)+r_i-(\tilde{\lambda}_0+\lambda_i)},$$

where

$$\tilde{B}(r_i) := \frac{\tilde{\lambda}_0}{\tilde{\lambda}_0+\lambda_i}\frac{\tilde{\mu}_0}{\tilde{\mu}_0+r_i}+\frac{\lambda_i}{\tilde{\lambda}_0+\lambda_i}\frac{\mu_i}{\mu_i+r_i}.$$

Substituting $\tilde{B}(r_i)$ we find

$$\tilde{V}_i(r_i) = \frac{1-\rho_0-\rho_i}{1-\frac{\tilde{\lambda}_0}{\tilde{\mu}_0+r_i}-\frac{\lambda_i}{\mu_i+r_i}}.$$

Let us assume we have L of these M/H$_2$/1 queues, all independent, indexed by i, $i=1,\ldots,L$. Then the joint LST of the workload $V(t)$ is given by

$$\tilde{V}(r) \equiv \tilde{V}(r_1,\ldots,r_L) = \prod_{i=1}^{L}\tilde{V}_i(r_i) = \prod_{i=1}^{L}\frac{1-\rho_0-\rho_i}{1-\frac{\tilde{\lambda}_0}{\tilde{\mu}_0+r_i}-\frac{\lambda_i}{\mu_i+r_i}}. \tag{6}$$

Comparing (6) with (4) indeed shows some similar terms. Obviously, the two expressions cannot be expected to be identical, because the linear network is different from L independent M/H$_2$/1 queues. Taking $r_i=r$, $i=1,\ldots,L$, (4) can however be rewritten as

$$\tilde{W}(r,\ldots,r) = \left(\frac{1-\frac{\tilde{\lambda}_0}{\tilde{\mu}_0+r}}{1-\rho_0}\right)^{L-1}\prod_{i=1}^{L}\frac{1-\rho_0-\rho_i}{1-\frac{\tilde{\lambda}_0}{\tilde{\mu}_0+r}-\frac{\lambda_i}{\mu_i+r}}.$$

The above provides some interpretation for the LST (4). It says that $\tilde{V}(r,\ldots,r)=\tilde{W}(r,\ldots,r)\tilde{U}(r)$, where

$$\tilde{U}(r) := \left(\frac{1-\rho_0}{1-\frac{\tilde{\lambda}_0}{\tilde{\mu}_0+r}}\right)^{L-1}$$

is a term that accounts for the dependence and interaction among the L M/H$_2$/1 queues. Note that the LST of the workload $S(t)$ in steady state in an M/M/1 queue with arrival rate $\tilde{\lambda}_0$ and service rate $\tilde{\mu}_0$ is given by $\tilde{S}(r)=(1-\rho_0)/\left(1-\frac{\tilde{\lambda}_0}{\tilde{\mu}_0+r}\right)$. Hence, $\tilde{U}(r)$ is the LST of the sum of the workloads in $L-1$ of these M/M/1 queues (all independent). The above shows that

$$\sum_{i=1}^{L}W_i+\sum_{i=1}^{L-1}U_i \stackrel{\mathrm{d}}{=} \sum_{i=1}^{L}V_i,$$

where U_i, $i=1,\ldots,L-1$ are i.i.d. copies of U, and $\stackrel{\mathrm{d}}{=}$ indicates that both sides are equal in distribution.

If $\alpha \neq 1$ or $\kappa_i \neq \kappa$, then there are no explicit expressions available for the steady-state distribution of $N(t)$.

4 Fluid and diffusion models

In this section we discuss the heavy-traffic results of [7, 8], which provide the basis for the approximations developed in Sections 6-7. Define the following fluid scaled processes:

$$\overline{N}^k(t) := N(kt)/k \quad \text{and} \quad \overline{W}^k(t) := W(kt)/k,$$

where $W_i(t) = N_0(t)/\mu_0 + N_i(t)/\mu_i$, $i = 1, \dots, L$. The fluid model can then be obtained from the original model by letting $k \to \infty$. For ease of notation, let $\overline{N}^\infty(t)$ be denoted by $\overline{N}(t)$, and $\overline{W}^\infty(t)$ by $\overline{W}(\overline{N}(t))$. Define

$$s_0(t) := \frac{\left(\kappa_0 \overline{N}_0(t)^\alpha\right)^{1/\alpha}}{\left(\kappa_0 \overline{N}_0(t)^\alpha\right)^{1/\alpha} + (\sum_{l=1}^L \kappa_l \overline{N}_l(t)^\alpha)^{1/\alpha}}; \quad s_i(t) := 1 - s_0(t), \quad i = 1, \dots, L,$$

i.e., $s_i(t)$ denotes the total service rate allocated to class i at time t, $i = 0, \dots, L$. Then the evolution of the workload process can be described as follows:

$$\frac{\mathrm{d}}{\mathrm{d}t}\overline{N}_i(t) = \lambda_i - \mu_i s_i(t), \qquad \text{for} \quad i = 0, \dots, L;$$
$$\overline{N}_i(t) \geq 0, \qquad \text{for} \quad i = 0, \dots, L.$$

A fluid model solution is an absolutely continuous function $\overline{N} : [0, \infty) \to \mathbb{R}_+^{L+1}$, such that at each regular point t for $\overline{N}(\cdot)$, we have that for $i = 0, \dots, L$,

$$\frac{\mathrm{d}}{\mathrm{d}t}\overline{N}_i(t) = \begin{cases} \lambda_i - \mu_i s_i(t) & \text{if } \overline{N}_i(t) > 0; \\ 0 & \text{if } \overline{N}_i(t) = 0, \end{cases}$$

and for $i = 1, \dots, L$,

$$s_0(t) I_0(t) + \rho_0(1 - I_0(t)) + s_i(t) I_i(t) + \rho_i(1 - I_i(t)) \leq 1,$$

where $I_i(t) = 1$ if $\overline{N}_i(t) > 0$, and 0 otherwise. A state $\overline{N}$ is called invariant if there is a fluid model solution such that $\overline{N}(t) = \overline{N}$ for all $t \geq 0$. Let $J := \{j \in \{1, \dots, L\} : \rho_0 + \rho_j = 1\} \neq \emptyset$ be the set of nodes that are critically loaded.

The following theorem appeared in slightly different form in [8].

Theorem 4.1 *The following statements are equivalent:*
(i) $\overline{N}$ is an invariant state;
(ii) $s_i(t) = \rho_i$ for all i such that $\overline{N}_i > 0$;
(iii) There is a $q \in \mathbb{R}_+^L$ such that

$$\overline{N}_0 = \rho_0 \left(\frac{\sum_{j \in J} q_j}{\kappa_0}\right)^{1/\alpha},$$

for $i \in J$,

$$\overline{N}_i = \rho_i \left(\frac{q_i}{\kappa_i}\right)^{1/\alpha},$$

and for $i \notin J$, $\overline{N}_i = 0$;
(iv) $\overline{N} = \triangle(\overline{W}(\overline{N}))$, where $\triangle(x)$ is the unique value of $\overline{N} \in \mathbb{R}_+^{L+1}$ that solves the optimization problem:

$$\begin{aligned} \min \quad & F(\overline{N}) = \tfrac{1}{\alpha+1} \textstyle\sum_{i=0}^L \lambda_i \kappa_i \mu_i^{\alpha-1} \left(\tfrac{\overline{N}_i}{\lambda_i}\right)^{\alpha+1} \\ \text{s.t.} \quad & \overline{N}_0/\mu_0 + \overline{N}_i/\mu_i \geq x_i, \quad i \in J; \\ & \overline{N}_i \geq 0, \quad i = 0, \dots, L. \end{aligned}$$

In the remainder of this section we assume that there are $L = 2$ nodes, and that $\kappa_0 = \kappa_1 = \kappa_2 = \kappa$. Furthermore, we assume heavy-traffic conditions at both nodes, i.e., $J = \{1, 2\}$. Define the diffusion scaled processes:

$$\hat{N}^k(t) := N(k^2 t)/k \text{ and } \hat{W}^k(t) := W(k^2 t)/k,$$

where $W_i(t) = N_0(t)/\mu_0 + N_i(t)/\mu_i$, $i = 1, 2$, as before. In [7] the authors show (under the assumptions mentioned above) that $\hat{W}^k(t)$ converges in distribution to a continuous process $\breve{W}(t)$ as $k \to \infty$. The process $\breve{W}(t)$ is a so-called Semimartingale Reflecting Brownian Motion (SRBM) that lives in the cone

$$\left\{ w : w_i = \frac{\rho_0}{\mu_0} \left(\frac{q_1 + q_2}{\kappa} \right)^{1/\alpha} + \frac{\rho_i}{\mu_i} \left(\frac{q_i}{\kappa} \right)^{1/\alpha}, \ q_1, q_2 \geq 0, \ i = 1, 2 \right\}.$$

In [7] it was shown that for all $\alpha \in (0, \infty)$ this is the same as the cone

$$\left\{ (w_1, w_2) : w_1 \geq 0, \quad w_1 \frac{\rho_0/\mu_0}{(1-\rho_0)/\mu_1 + \rho_0/\mu_0} \leq w_2 \leq w_1 \frac{(1-\rho_0)/\mu_2 + \rho_0/\mu_0}{\rho_0/\mu_0} \right\},$$

as depicted in Figure 2. The state space is an infinite two-dimensional wedge, and the process behaves in the interior of the wedge like a two-dimensional Brownian motion with zero drift and covariance matrix

$$\begin{pmatrix} 2\left(\frac{\rho_0}{\mu_0} + \frac{\rho_1}{\mu_1}\right) & 2\frac{\rho_0}{\mu_0} \\ 2\frac{\rho_0}{\mu_0} & 2\left(\frac{\rho_0}{\mu_0} + \frac{\rho_2}{\mu_2}\right) \end{pmatrix}.$$

The process reflects instantaneously at the boundary of the wedge, the angle of reflection being constant along each side. Vertical (horizontal) reflection on the bounding face $w_2 = w_1 \frac{\rho_0/\mu_0}{(1-\rho_0)/\mu_1+\rho_0/\mu_0}$ $\left(w_2 = w_1 \frac{(1-\rho_0)/\mu_2+\rho_0/\mu_0}{\rho_0/\mu_0}\right)$ can be interpreted as a manifestation of so-called entrainment: congestion at node 1 (node 2) prevents node 2 (node 1) from utilizing the full service rate. In [12] it was shown that the process is transient in the cone, i.e., no steady-state distribution exists.

5 Single bottleneck node

In this section we propose a method for approximating $\mathbb{E}N_i$, $i = 0, \ldots, L$, in case of a single bottleneck node, i.e., $|J| = 1$. In case just a single node, say z, $z \in \{1, \ldots, L\}$, is critically loaded, statement (iii) of Theorem 4.1 suggests that the number of class-i users, $i = 1, \ldots, L$, $i \neq z$, will be negligible compared to the number of class-0 and class-z users. Hence, the service rates allocated to the various classes will be predominantly determined by the number of class-0 and class-z users, and approximately equal

$$s_0(n) = \frac{(\kappa_0 n_0^\alpha)^{1/\alpha}}{(\kappa_0 n_0^\alpha)^{1/\alpha} + \left(\sum_{j=1}^L \kappa_j n_j^\alpha\right)^{1/\alpha}} \approx \frac{(\kappa_0 n_0^\alpha)^{1/\alpha}}{(\kappa_0 n_0^\alpha)^{1/\alpha} + (\kappa_z n_z^\alpha)^{1/\alpha}} = \frac{\kappa_0^* n_0}{\kappa_0^* n_0 + \kappa_z^* n_z};$$

$$s_i(n) = \frac{\left(\sum_{j=1}^L \kappa_j n_j^\alpha\right)^{1/\alpha}}{(\kappa_0 n_0^\alpha)^{1/\alpha} + \left(\sum_{j=1}^L \kappa_j n_j^\alpha\right)^{1/\alpha}} \approx \frac{\kappa_z^* n_z}{\kappa_0^* n_0 + \kappa_z^* n_z},$$

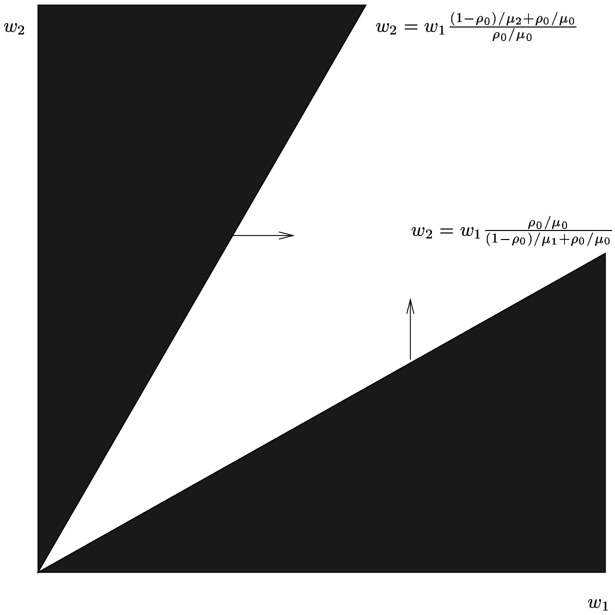

Figure 2: The workload cone.

for $i = 1, \ldots, L$, where $\kappa_0^* = \kappa_0^{1/\alpha}$ and $\kappa_z^* = \kappa_z^{1/\alpha}$. Thus, node z roughly behaves as a Discriminatory Processor-Sharing model with relative weights κ_0^* and κ_z^* for classes 0 and z, respectively. The results of [5] then imply that $\mathbb{E}N_0$ and $\mathbb{E}N_z$ satisfy the set of linear equations

$$\mathbb{E}N_0 - \rho_0 \mathbb{E}N_0 - \kappa_z^* \frac{\lambda_z \mathbb{E}N_0 + \lambda_0 \mathbb{E}N_z}{\kappa_0^* \mu_0 + \kappa_z^* \mu_z} \approx \rho_0;$$

$$\mathbb{E}N_z - \rho_z \mathbb{E}N_z - \kappa_0^* \frac{\lambda_z \mathbb{E}N_0 + \lambda_0 \mathbb{E}N_z}{\kappa_0^* \mu_0 + \kappa_z^* \mu_z} \approx \rho_z,$$

from which we deduce that

$$\mathbb{E}N_0 \approx \frac{\rho_0}{1 - \rho_0 - \rho_z} \left(1 + \frac{\mu_0 \rho_z (\kappa_z^* - \kappa_0^*)}{\kappa_0^* \mu_0 (1 - \rho_0) + \kappa_z^* \mu_z (1 - \rho_z)} \right);$$

$$\mathbb{E}N_z \approx \frac{\rho_z}{1 - \rho_0 - \rho_z} \left(1 + \frac{\mu_z \rho_0 (\kappa_0^* - \kappa_z^*)}{\kappa_0^* \mu_0 (1 - \rho_0) + \kappa_z^* \mu_z (1 - \rho_z)} \right).$$

Let $\mathbb{E}\tilde{N}_i$ denote the approximation for $\mathbb{E}N_i$, $i = 0, z$. Then

$$\bar{s} := \frac{\kappa_z^* \mathbb{E}\tilde{N}_z}{\kappa_0^* \mathbb{E}\tilde{N}_0 + \kappa_z^* \mathbb{E}\tilde{N}_z}$$

can be regarded as an approximation for the service rate allocated to classes $i = 1, \ldots, L$, $i \neq z$. The number of class-i users, $i = 1, \ldots, L$, $i \neq z$, will approximately behave as in an M/M/1 queue with arrival rate λ_i and service rate $\mu_i \bar{s}$. This gives the approximation

$$\mathbb{E}N_i \approx \frac{\rho_i}{\bar{s} - \rho_i}, \quad i = 1, \ldots, L, \quad i \neq z.$$

Note that the values of κ_i, $i = 1, \ldots, L$, $i \neq z$, do not appear in this approximation. This suggests that the weights of classes that do not traverse the bottleneck node, will tend to have limited impact on the flow-level performance.

We now discuss the numerical experiments that we conducted to examine the accuracy of the above-described method. We first test this approach for a linear network with $L = 2$ nodes, $\alpha = 1$, and $\kappa_i = \kappa$, $i = 0, 1, 2$, for which we have exact expressions for $\mathbb{E}N_i$, $i = 0, 1, 2$, see Theorem 3.1. We fix $\rho_0 = 0.6$ and $\rho_1 = 0.39$, so that node 1 is highly loaded ($z = 1$), and vary the value of ρ_2. Note that in case of equal weights, the approximations only depend on the traffic characteristics through the class loads, and not on the specific values of the λ_is and μ_is. The results are presented in Table 1, and indicate that the approximations are remarkably accurate. As could be expected, the smaller the value of ρ_2, the better the approximations.

In case $\alpha \neq 1$ or $\kappa_i \neq \kappa$, there are no exact expressions available for $\mathbb{E}N_i$, $i = 0, 1, 2$, and we need to resort to simulation experiments to investigate the accuracy of the approximations. Throughout this paper, the simulation numbers are obtained as averages over 10000 busy periods. We choose the same setting as above, but with $\kappa_0 = 2$, $\kappa_1 = 0.5$ and $\kappa_2 = 1$. In this case the approximations do depend on the specific values of the μ_is. We consider two scenarios: in Scenario 1 we take $\mu_0 = \mu_1 = \mu_2 = 1$, while in Scenario 2 we set $\mu_0 = 0.75$, $\mu_1 = 1$ and $\mu_2 = 1.5$. The results are presented in Tables 2-3.

Note that the approximations for $\mathbb{E}N_0$ and $\mathbb{E}N_1$ do not depend on the presence of class-2 users, and are in particular independent of the value of ρ_2. Further observe that if $\alpha \to \infty$, then $\kappa_0^*, \kappa_1^* \to 1$, and as a consequence $\mathbb{E}N_i \approx \rho_i/(1 - \rho_0 - \rho_i)$, $i = 0, 1$. The results are surprisingly accurate, even if node 2 is also relatively highly loaded ($\rho_0 + \rho_2 = 0.9$). Note that $\mathbb{E}N_2^{sim}$ is increasing in ρ_2, as could be expected. The influence of ρ_2 on $\mathbb{E}N_0^{sim}$ and $\mathbb{E}N_1^{sim}$ is more subtle, as closer inspection of Tables 2-3 demonstrates. It might be natural to expect that increasing ρ_2 would also have an adverse impact on $\mathbb{E}N_0^{sim}$ and $\mathbb{E}N_1^{sim}$. As the value of ρ_2 and the number of class-2 users increases, however, the service rate $s_0(n)$ will decrease, whereas the service rate $s_2(n)$ will increase. The resulting increase in the number of class-0 users will have the counteracting effect of decreasing $s_2(n)$, and conversely the expected decrease in the number of class-2 users will have the opposite effect of increasing $s_0(n)$. Because of these interacting effects, the net impact basically remains unpredictable, and as Tables 2-3 reveal, $\mathbb{E}N_0^{sim}$ and $\mathbb{E}N_1^{sim}$ do not necessarily change in a monotone manner as the value of ρ_2 increases.

6 Two bottleneck nodes and equal weights: workload invariance

In this section we consider the scenario that there are two nodes critically loaded, i.e., $|J| = 2$. Since the nodes can be indexed arbitrarily, we may assume without loss of generality that $J = \{1, 2\}$. Also, suppose that $\kappa_i = \kappa$, $i = 0, \dots, L$.

Let $W(t)$ be the workload process associated with the two bottleneck nodes. The results from [7, 8] as reviewed in Section 4 indicate that the behavior of $W(t)$ is asymptotically independent of the value of α. In particular, this suggests that the behavior of the workload process can be approximated by the known distribution for $\alpha = 1$. In order to examine this hypothesis, we calculated the mean workload (using Theorem 3.2)

$$\mathbb{E}W_i^{exact}(\alpha = 1) \equiv \mathbb{E}W_i^{exact}(1) = \frac{\lambda_i/\mu_i^2}{1 - \rho_0 - \rho_i} + \frac{\lambda_0/\mu_0^2}{1 - \rho_0}\left(1 + \sum_{j=1}^{L} \frac{\rho_j}{1 - \rho_0 - \rho_j}\right), \tag{7}$$

with $i \in J$, and compared it with simulation for the case of $L = 2$ nodes, $\rho_0 + \rho_1 = \rho_0 + \rho_2 = 0.99$, i.e., $|J| = 2$, and $\mu_i = \kappa_i = 1$, $i = 0, 1, 2$. We also considered the asymmetric case $\rho_0 + \rho_1 = \rho_0 + \rho_2 = 0.99$, $\kappa_i = 1$, $i = 0, 1, 2$, $\mu_0 = 0.75$, $\mu_1 = 1$ and $\mu_2 = 1.5$.

Define

$$X_i := \mathbb{E}W_i^{sim}(\alpha)/\mathbb{E}W_i^{exact}(1) - 1, \quad i \in J.$$

The results, summarized in Table 4, indicate that the mean workload for $\alpha = 1$ indeed provides a reasonably accurate approximation for a wide range of α values. Note that X_i should be equal to 0 for all cases with $\alpha = 1$. In most cases with $\alpha > 1$, $\mathbb{E}W_i^{exact}(1)$ is larger than $\mathbb{E}W_i^{sim}(\alpha)$, and thus seems to yield a conservative approximation. Below we provide an explanation for this observation. In preparation for that, we first present the following proposition.

Proposition 6.1 *For fixed $n = (n_0, \ldots, n_L)$ and $\kappa_i = \kappa$, $i = 0, \ldots, L$, the service rate $s_0(n)$ allocated to class-0 users is increasing in α.*

Proof: For fixed $n = (n_0, \ldots, n_L)$ and $\kappa_i = \kappa$, $i = 0, \ldots, L$, we obtain from (1) that

$$s_0(n) = \frac{n_0}{n_0 + (\sum_{j=1}^L n_j^\alpha)^{1/\alpha}}. \tag{8}$$

Equivalently, we have to prove that $(\sum_{j=1}^L n_j^\alpha)^{1/\alpha}$ is decreasing in α. First note that

$$n_1^{\alpha r} + \cdots + n_L^{\alpha r} \leq (n_1^\alpha + \cdots + n_L^\alpha)^r$$

for all $r \geq 1$. Therefore,

$$\left(\sum_{i=1}^L n_i^\beta\right)^{1/\beta} = \left(\sum_{i=1}^L n_i^{\alpha r}\right)^{1/\alpha r} = (n_1^{\alpha r} + \cdots + n_L^{\alpha r})^{1/\alpha r} \leq (n_1^\alpha + \cdots + n_L^\alpha)^{r/\alpha r} = \left(\sum_{i=1}^L n_i^\alpha\right)^{1/\alpha},$$

for all $\beta \geq \alpha$, which proves the stated. □

Now observe that the workload at each of the nodes is minimized (sample-path-wise in fact) when class 0 receives priority over classes $i \in J$. Since the capacity allocated to class-0 users is increasing in α, it is thus plausible that more generally the mean workload $\mathbb{E}W_i^{exact}(\alpha)$ decreases as function of α, which implies that X_i is smaller than 0 for $\alpha > 1$, $i \in J$. This provides an explanation for the negative values in Table 4. The latter property can in fact be rigorously proved using Proposition 6.1 and stochastic coupling arguments. Due to page limitations, we omit the proof.

7 Two bottleneck nodes and equal weights: approximation methods

In this section we develop three methods for approximating $\mathbb{E}N_i$, $i = 0, 1, 2$. Recall that we suppose that $J = \{1, 2\}$ and $\kappa_i = \kappa$, $i = 0, \ldots, L$. The various methods differ in some technical details, but they all rely on the insights from the heavy-traffic results as reviewed in Section 4. In Section 7.4 we present approximations for $\mathbb{E}N_i$, $i = 3, \ldots, L$.

7.1 Method 1

The numerical results presented in the previous section indicate that $\mathbb{E}W_i^{exact}(\alpha)$ is nearly constant in $\alpha \in (0, \infty)$, provided that the load at all nodes $i \in J$ is sufficiently high. In particular, it is approximately equal to the known value for $\alpha = 1$ as given by (7). Further observe that $\mathbb{E}W_i^{exact}(\alpha) = \mathbb{E}N_0/\mu_0 + \mathbb{E}N_i/\mu_i$, $i \in J$. Thus, we obtain

$$\mathbb{E}N_0/\mu_0 + \mathbb{E}N_i/\mu_i \approx \frac{\lambda_i/\mu_i^2}{1 - \rho_0 - \rho_i} + \frac{\lambda_0/\mu_0^2}{1 - \rho_0}\left(1 + \sum_{j=1}^L \frac{\rho_j}{1 - \rho_0 - \rho_j}\right), \quad i \in J, \tag{9}$$

i.e., a set of two linear equations with three unknowns. If we can find one additional constraint, then we should be able to determine $\mathbb{E}N_i$, $i = 0, 1, 2$ (as long as the resulting system of equations is non-singular).

Now observe that Theorem 4.1 shows that an invariant state $\overline{N}$ in the fluid model can be expressed as

$$\overline{N}_0 = \rho_0 \left(\frac{\sum_{i \in J} q_i}{\kappa_0} \right)^{1/\alpha}; \quad \overline{N}_i = \rho_i \left(\frac{q_i}{\kappa_i} \right)^{1/\alpha}, \quad i \in J, \quad q \in \mathbb{R}_+^2.$$

This suggests the following approximation:

$$\begin{aligned}
&(\mathbb{E}N_0, \mathbb{E}N_1, \mathbb{E}N_2) \approx \\
&\int_{q_1=0}^{\infty} \int_{q_2=0}^{\infty} \left(\rho_0 \left(\frac{\sum_{i \in J} q_i}{\kappa} \right)^{1/\alpha}, \rho_1 \left(\frac{q_1}{\kappa} \right)^{1/\alpha}, \rho_2 \left(\frac{q_2}{\kappa} \right)^{1/\alpha} \right) \mathrm{d}\mathbb{P}(Q_1 < q_1, Q_2 < q_2) = \\
&\frac{1}{\kappa^{1/\alpha}} \left(\rho_0 \mathbb{E} \left(\left(\sum_{i \in J} Q_i \right)^{1/\alpha} \right), \rho_1 \mathbb{E} \left(Q_1^{1/\alpha} \right), \rho_2 \mathbb{E} \left(Q_2^{1/\alpha} \right) \right).
\end{aligned}$$

Using the additional approximation

$$(\mathbb{E}N_0, \mathbb{E}N_1, \mathbb{E}N_2) \approx \frac{\gamma}{\kappa^{1/\alpha}} \left(\rho_0 \left(\sum_{i \in J} \mathbb{E}Q_i \right)^{1/\alpha}, \rho_1 \left(\mathbb{E}Q_1 \right)^{1/\alpha}, \rho_2 \left(\mathbb{E}Q_2 \right)^{1/\alpha} \right), \tag{10}$$

with γ some multiplicative constant, and substituting (10) in (9) then yields a system of two equations with two unknowns. Numerically solving this system yields $\gamma^\alpha \mathbb{E}Q_i$, $i = 1, 2$, from which we can obtain $\mathbb{E}N_i$, $i = 0, 1, 2$, using (10). Note that

$$\mathbb{E} \left(\left(\sum_{i \in J} Q_i \right)^{1/\alpha} \right) \leq \sum_{i \in J} \mathbb{E} \left(Q_i^{1/\alpha} \right)$$

if $\alpha \in (1, \infty)$, which would provide an upper bound for $\mathbb{E}N_0$ relative to $\mathbb{E}N_i$, $i \in J$. Likewise, if $\alpha \in (0, 1)$, then this would give a lower bound for $\mathbb{E}N_0$ relative to $\mathbb{E}N_i$, $i \in J$.

We tested this approach by comparing the results with simulation figures. We took the same simulation parameters as in the previous section. The results are presented in Tables 5-6. Throughout this paper, $\mathbb{E}N_i^{Mj}$ denotes the approximation of $\mathbb{E}N_i$ that is obtained by using Method j. Note that in Table 5 we have $\mathbb{E}N_1^{M1} = \mathbb{E}N_2^{M1}$ by symmetry. The tables indicate that Method 1 gives reasonably accurate estimates for $\mathbb{E}N_i$, particularly $\mathbb{E}N_0$. Note that Method 1 is fast as well: it suffices to solve a system of two equations with two unknowns.

7.2 Method 2

We now discuss a second method for approximating $\mathbb{E}N_i$, $i = 0, 1, 2$. Again, we start from Equation (9) as in Method 1. The difference with Method 1 is that we now use statement *(iv)* (instead of *(iii)*) of Theorem 4.1. Statement *(iv)* implies that a workload vector $w = (w_1, w_2)$ uniquely determines a state vector n that solves the optimization problem:

$$\min \quad F(n_0, n_1, \ldots, n_L) = \frac{1}{\alpha + 1} \sum_{i=0}^{L} \lambda_i \kappa_i \mu_i^{\alpha - 1} \left(\frac{n_i}{\lambda_i} \right)^{\alpha + 1}$$

$$\text{s.t.} \quad n_0/\mu_0 + n_i/\mu_i \geq w_i, \quad i \in J; \tag{11}$$
$$n_i \geq 0, \quad i = 0, \dots, L.$$

The method now works as follows. We determine the vector $(\mathbb{E}N_0, \mathbb{E}N_1, \dots, \mathbb{E}N_L)$ that minimizes the function $F(\mathbb{E}N_0, \mathbb{E}N_1, \dots, \mathbb{E}N_L)$ subject to the constraints in Equation (9). Note that $\mathbb{E}N_i = 0$, $i = 3, \dots, L$.

As it turns out, Methods 1 and 2 result in similar approximations for $\mathbb{E}N_i$, $i = 0, 1, 2$. This is not too surprising: the only difference between the methods is that we use statement (iii) in one case, and (iv) in the other. However, statements (iii) and (iv) are in fact equivalent in case of heavy traffic, so both methods should roughly agree when the load is sufficiently high.

Remark. Method 2 uses the *mean* workloads to approximate the mean number of users. However, we can potentially improve the accuracy of the approximation if we use the *distribution* of the workloads, which is also asymptotically independent of α in heavy traffic. The resulting approximation is then given by

$$\mathbb{E}N_i = \sum_{n \geq 0} \triangle_i(w(n))\pi(n), \quad i = 0, 1, 2,$$

where $w_i(n) = n_0/\mu_0 + n_i/\mu_i$, $i \in J$, $\triangle(x)$ is as in Theorem 4.1, and $\pi(n)$ is given by (2). This will typically result in a different approximation for $\mathbb{E}N_i$ than Method 2, since the optimization problem (11) is non-linear. The disadvantage is that it is very time-consuming.

7.3 Method 3

This method is similar to both previous methods, i.e., we again start from Equation (9) to obtain a set of two equations with three unknowns $\mathbb{E}N_i$, $i = 0, 1, 2$. Statement (ii) of Theorem 4.1 provides an additional equation, which allows us to numerically solve to above system of equations. First note from (8) that

$$\sum_{n \geq 0} \frac{n_0}{n_0 + (\sum_{l=1}^L n_l^\alpha)^{1/\alpha}} \tilde{\pi}(n) = \sum_{n \geq 0} s_0(n)\tilde{\pi}(n) = \rho_0,$$

where $\tilde{\pi}(n)$ is the steady-state distribution of $N(t)$ in case $\alpha \in (0, \infty)/\{1\}$. The additional equation is then obtained by replacing the latter equation by the approximation

$$\frac{\mathbb{E}N_0}{\mathbb{E}N_0 + (\sum_{l=1}^L \mathbb{E}N_l^\alpha)^{1/\alpha}} \approx \frac{\mathbb{E}N_0}{\mathbb{E}N_0 + (\sum_{l \in J} \mathbb{E}N_l^\alpha)^{1/\alpha}} = \rho_0.$$

We numerically solved the above system of equations for both the symmetric and asymmetric scenarios considered in the previous section. The results are presented in Tables 7-8. Note that the approximations obtained from Method 3 slightly differ from those of Methods 1 and 2. This may be explained from the fact that statement (ii) of Theorem 4.1 (for $i = 0$) is only partly satisfied.

7.4 Approximation for non-bottleneck nodes

In the previous subsections we presented three methods for approximating the mean number of users at the bottleneck nodes. We now provide an approximation for the number of users at the remaining nodes, i.e., $\mathbb{E}N_i$, $i = 3, \dots, L$. The method is similar in nature as the one presented in

Section 5 for the case of a single bottleneck node. Let $\mathbb{E}\tilde{N}_i$ denote the approximations obtained for $\mathbb{E}N_i$, $i = 0, 1, 2$. In view of (8), define

$$\bar{s}_0 := \frac{\mathbb{E}\tilde{N}_0}{\mathbb{E}\tilde{N}_0 + \left(\sum_{j \in J} \mathbb{E}\tilde{N}_j^{\alpha}\right)^{1/\alpha}}$$

as an approximation for the service rate allocated to class 0. As before, the number of class-i users, $i = 3, \dots, L$, will roughly behave as in an M/M/1 queue with arrival rate λ_i and service rate $\mu_i(1 - \bar{s}_0)$. This gives the approximation

$$\mathbb{E}N_i \approx \frac{\rho_i}{1 - \bar{s}_0 - \rho_i}, \quad i = 3, \dots, L.$$

Remark. In [3] the authors characterized the class of allocation functions that have the property that the steady-state distribution of the number of users is insensitive, i.e., only depends on the traffic characteristics of the various classes through their respective loads. In addition, they introduced the notion of *balanced fairness*, which refers to the most efficient insensitive allocation. For the linear network (see Figure 1) it was shown in [4] that balanced fairness is equivalent to unweighted proportional fairness, i.e., $\alpha = 1$ and $\kappa_i = \kappa$, $i = 0, \dots, L$. Note that the steady-state distribution in Theorem 3.1 indeed only depends on the loads, and not on any higher-order traffic characteristics. In [2] it was shown that balanced fairness provides a good approximation for unweighted proportional fairness and unweighted max-min fairness. The results of this section, though, illustrate that the accuracy of the balanced fairness approximation for unweighted max-min fairness degrades in heavy-traffic conditions.

8 Conclusion

We analyzed the flow-level performance of a linear network carrying elastic traffic, where the service capacity is shared according to a weighted alpha-fair policy. We devised approximations for the mean number of users based on the assumption that one or two of the nodes operate under heavy-traffic conditions. The approximations for the mean number of users also yield simple estimates for the mean transfer delays and user throughputs.

In case just a single node operates in heavy traffic, we exploited the fact that the user population at this node approximately evolves as in a two-class Discriminatory Processor-Sharing model. Numerical tests indicate that the resulting approximations based on the results of [5] are remarkably accurate.

In the case that there are two nodes critically loaded, we relied on the principle that the joint workload process at these nodes is asymptotically independent of the fairness coefficient alpha, provided all classes have equal weights. In particular, we derived the joint workload process for an unweighted Proportional Fair policy, and used this as an approximation for all unweighted alpha-fair policies. We then presented three approximation methods for the mean number of users based on the results of [7, 8], which relate the mean number of users to the mean workloads at the critically loaded nodes. Extensive numerical tests demonstrated that the resulting approximations are similar for the three methods, and tend to be reasonably accurate across a wide range of settings, even at relatively moderate load values. Irrespective of the number of critically loaded nodes, the numbers of users at the other (less loaded) nodes roughly behave as in an M/M/1 queue with reduced service capacity, and are hardly influenced by the corresponding class weights.

It is substantially more difficult to handle the cases in which there are 1) two nodes critically loaded and not all class weights are equal, or 2) more than two bottleneck nodes. Although the mean number of users can still be related to the mean workloads in these scenarios, the joint workload process at these nodes is no longer independent of the fairness coefficient alpha. In addition, even for a weighted Proportional Fair policy the workload distribution is no longer known. Hence, we cannot apply the three methods mentioned above for approximating the mean number of users. One option to obtain conservative estimates in case 2) would be to use the property that the mean workload for an unweighted alpha-fair policy, with alpha larger than one, is smaller than for an unweighted Proportional Fairness policy as mentioned in Section 6. Alternatively, as in Section 4, we can approximate the workload process by an SRBM living in a cone that now does depend on the fairness coefficient alpha. Subsequently, we can derive the steady-state distribution of the process, thus having an approximation for the mean workloads. If we succeeded in this, then we could obtain approximations for the mean number of users by applying one of the three methods. However, it turns out to be extremely hard to derive the steady-state distribution for an SRBM living in a multi-dimensional cone, see [6]. The latter suggests that it is also hard to determine the steady-state distribution of the approximation for the workload process, if possible at all.

A natural extension of the present work is to consider generally distributed service requirements. One can expect that some of the results carry over in heavy-traffic conditions. Another future research direction includes extending the results to general network topologies, e.g., cyclic networks and grid networks. The method presented for the single bottleneck node in the linear network, should also be effective in network topologies that have multiple bottleneck nodes, but where classes of users require service at just one bottleneck node.

Acknowledgment The authors thank Maaike Verloop for providing the simulation program. Also, the authors wish to thank three anonymous referees for providing helpful comments.

References

[1] Bonald, T., Massoulié, L. (2001). Impact of fairness on Internet performance. *Proceedings of the ACM SIGMETRICS/Performance 2001 Conference*, Boston MA, USA, 82-91.

[2] Bonald, T., Massoulié, L., Proutière, A., Virtamo, J. (2006). A queueing analysis of max-min fairness, proportional fairness and balanced fairness. To appear in: *Queueing Systems and Applications.*

[3] Bonald, T., Proutière, A. (2003). Insensitive bandwidth sharing in data networks. *Queueing Systems*, 44: 69-100.

[4] Bonald, T., Proutière, A., Roberts, J.W., Virtamo, J. (2003). Computational aspects of balanced fairness. *Proceedings of the 18th International Teletraffic Congress*, Berlin, Germany, 801-810.

[5] Fayolle, G., Mitrani, I., Iasnogorodski, R. (1980). Sharing a processor among many job classes. *Journal of the ACM*, 27: 519-532.

[6] Harrison, J.M., Williams, R.J. (1987). Multidimensional reflected Brownian motions having exponential stationary distributions. *Annals of Probability*, 15: 115-137.

[7] Kang, W., Kelly, F.P., Lee, N., Williams, R.J. (2004). Fluid and Brownian approximations for an internet congestion control model. *Proceedings of the 43rd IEEE Conference on Decision and Control*, the Bahamas, 3938-3943.

[8] Kelly, F.P., Williams, R.J. (2004). Fluid model for a network operating under a fair bandwidth-sharing policy. *Annals of Applied Probability*, 14: 1055-1083.

[9] Massoulié, L., Roberts, J.W. (2000). Bandwidth sharing and admission control for elastic traffic. *Telecommunication Systems*, 15: 185-201.

[10] Mo, J., Walrand, J. (2000). Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8: 556-567.

[11] Padhye, J., Firoiu, V., Towsley, D., Kurose, J. (2000). Modeling TCP Reno performance: A simple model and its empirical validation. *IEEE/ACM Transactions on Networking*, 8: 133-145.

[12] Varadhan, S., Williams, R.J. (1985). Brownian motion in a wedge with oblique reflection. *Communications on Pure and Applied Mathematics*, 38: 405-443.

ρ_2	$\mathbb{E}N_0^{exact}$	$\mathbb{E}N_0^{method}$	$\mathbb{E}N_1^{exact}$	$\mathbb{E}N_1^{method}$	$\mathbb{E}N_2^{exact}$	$\mathbb{E}N_2^{method}$
0.1	60.50	60.00	39.00	39.00	0.33	0.34
0.2	61.50	60.00	39.00	39.00	1.00	1.03
0.3	64.50	60.00	39.00	39.00	3.00	3.19

Table 1: Results for $\alpha = 1$ and $\kappa_0 = \kappa_1 = \kappa_2 = \kappa$.

ρ_2	α	$\mathbb{E}N_0^{sim}$	$\mathbb{E}N_0^{method}$	$\mathbb{E}N_1^{sim}$	$\mathbb{E}N_1^{method}$	$\mathbb{E}N_2^{sim}$	$\mathbb{E}N_2^{method}$
0.1	1	27.88	28.24	69.12	70.76	0.35	0.35
0.2	1	27.54	28.24	67.25	70.76	1.06	1.08
0.3	1	34.51	28.24	80.19	70.76	3.34	3.52
0.1	2	49.76	43.40	63.75	55.60	0.36	0.34
0.2	2	45.08	43.40	57.55	55.60	1.08	1.05
0.3	2	41.59	43.40	52.34	55.60	3.46	3.32
0.1	5	61.35	53.43	52.03	45.57	0.35	0.34
0.2	5	54.37	53.43	46.40	45.57	1.08	1.04
0.3	5	50.23	53.43	42.52	45.57	3.47	3.24
0.1	∞	60.09	60.00	39.20	39.00	0.36	0.34
0.2	∞	60.68	60.00	39.52	39.00	1.08	1.03
0.3	∞	63.72	60.00	40.99	39.00	3.52	3.19

Table 2: Results for Scenario 1.

ρ_2	α	$\mathbb{E}N_0^{sim}$	$\mathbb{E}N_0^{method}$	$\mathbb{E}N_1^{sim}$	$\mathbb{E}N_1^{method}$	$\mathbb{E}N_2^{sim}$	$\mathbb{E}N_2^{method}$
0.1	1	35.72	30.91	89.57	77.78	0.34	0.35
0.2	1	26.21	30.91	63.59	77.78	1.06	1.07
0.3	1	34.84	30.91	81.73	77.78	3.37	3.48
0.1	2	45.85	45.60	58.73	58.34	0.35	0.34
0.2	2	50.17	45.60	63.96	58.34	1.09	1.04
0.3	2	58.57	45.50	74.97	58.34	3.51	3.31
0.1	5	54.47	54.43	46.59	46.43	0.36	0.34
0.2	5	63.51	54.43	53.81	46.43	1.09	1.04
0.3	5	55.19	54.43	46.88	46.43	3.59	3.24
0.1	∞	44.16	60.00	28.72	39.00	0.36	0.34
0.2	∞	78.66	60.00	50.98	39.00	1.12	1.03
0.3	∞	53.80	60.00	34.60	39.00	3.72	3.19

Table 3: Results for Scenario 2.

ρ_0	$\rho_1 = \rho_2$	α	X_1	X_2
0.3	0.69	1	0.001	0.019
0.5	0.49	1	0.006	0.020
0.7	0.29	1	-0.015	-0.024
0.3	0.69	2	-0.042	-0.047
0.5	0.49	2	-0.041	-0.056
0.7	0.29	2	-0.039	-0.040
0.3	0.69	5	-0.027	-0.065
0.5	0.49	5	-0.005	0.003
0.7	0.29	5	-0.058	-0.069
0.3	0.69	∞	-0.007	0.011
0.5	0.49	∞	-0.043	-0.063
0.7	0.29	∞	-0.061	-0.055

ρ_0	$\rho_1 = \rho_2$	α	X_1	X_2
0.3	0.69	1	0.006	-0.009
0.5	0.49	1	-0.046	-0.033
0.7	0.29	1	0.048	0.042
0.3	0.69	2	-0.065	-0.077
0.5	0.49	2	-0.025	-0.038
0.7	0.29	2	-0.039	-0.049
0.3	0.69	5	-0.040	-0.036
0.5	0.49	5	-0.055	-0.057
0.7	0.29	5	-0.037	-0.035
0.3	0.69	∞	-0.028	-0.022
0.5	0.49	∞	-0.048	-0.076
0.7	0.29	∞	-0.003	-0.009

Table 4: Testing whether $W(t)$ is independent of α. Left (Right): the symmetric (asymmetric) case.

ρ_0	$\rho_1 = \rho_2$	α	$\mathbb{E}N_0^{sim}$	$\mathbb{E}N_0^{M1}$	$\mathbb{E}N_1^{sim}$	$\mathbb{E}N_2^{sim}$	$\mathbb{E}N_1^{M1}$=$\mathbb{E}N_2^{M1}$
0.3	0.69	1	60.20	59.80	68.58	70.81	68.77
0.5	0.49	1	100.27	99.33	48.61	50.66	48.67
0.7	0.29	1	135.01	138.07	29.19	27.67	28.60
0.3	0.69	2	50.21	48.95	72.98	72.36	79.62
0.5	0.49	2	86.98	87.42	54.90	52.79	60.58
0.7	0.29	2	126.62	128.91	33.59	33.38	37.76
0.3	0.69	5	47.86	42.83	77.25	72.40	85.75
0.5	0.49	5	88.24	79.86	58.97	60.25	68.14
0.7	0.29	5	120.76	122.49	36.22	34.38	44.18
0.3	0.69	∞	49.75	39.02	77.94	80.25	89.55
0.5	0.49	∞	82.62	74.83	59.02	56.06	73.17
0.7	0.29	∞	120.81	117.92	35.64	36.67	48.75

Table 5: Results for Method 1: the symmetric case.

ρ_0	$\rho_1 = \rho_2$	α	$\mathbb{E}N_0^{sim}$	$\mathbb{E}N_0^{M1}$	$\mathbb{E}N_1^{sim}$	$\mathbb{E}N_1^{M1}$	$\mathbb{E}N_2^{sim}$	$\mathbb{E}N_2^{M1}$
0.3	0.69	1	59.40	59.75	70.06	68.77	67.58	68.65
0.5	0.49	1	95.19	99.22	45.80	48.70	48.46	48.55
0.7	0.29	1	143.78	137.93	31.01	28.66	29.58	28.48
0.3	0.69	2	50.17	51.12	71.83	80.27	73.27	85.91
0.5	0.49	2	90.58	90.72	55.61	60.04	56.49	65.56
0.7	0.29	2	127.79	131.91	33.86	36.68	33.78	40.51
0.3	0.69	5	49.89	45.99	75.98	87.11	81.50	96.17
0.5	0.49	5	85.94	85.16	56.51	67.46	61.03	76.68
0.7	0.29	5	127.15	127.71	35.20	42.27	39.39	48.90
0.3	0.69	∞	50.20	43.75	77.31	90.08	83.68	100.64
0.5	0.49	∞	84.83	82.88	59.16	70.49	58.67	81.23
0.7	0.29	∞	130.52	126.06	37.72	44.48	40.41	52.22

Table 6: Results for Method 1: the asymmetric case.

ρ_0	$\rho_1 = \rho_2$	α	$\mathbb{E}N_0^{sim}$	$\mathbb{E}N_0^{M3}$	$\mathbb{E}N_1^{sim}$	$\mathbb{E}N_2^{sim}$	$\mathbb{E}N_1^{M3}$=$\mathbb{E}N_2^{M3}$
0.3	0.69	1	60.20	59.34	68.58	70.81	69.23
0.5	0.49	1	100.27	98.67	48.61	50.66	49.33
0.7	0.29	1	135.01	137.26	29.19	27.67	29.41
0.3	0.69	2	50.21	48.52	72.98	72.36	80.05
0.5	0.49	2	86.98	86.70	54.90	52.79	61.30
0.7	0.29	2	126.62	127.91	33.59	33.38	38.76
0.3	0.69	5	47.86	42.41	77.25	72.40	86.16
0.5	0.49	5	88.24	79.12	58.97	60.25	68.88
0.7	0.29	5	120.76	121.38	36.22	34.38	45.29
0.3	0.69	∞	49.75	38.63	77.94	80.25	89.94
0.5	0.49	∞	82.62	74.08	59.02	56.06	73.92
0.7	0.29	∞	120.81	116.74	35.64	36.67	49.93

Table 7: Results for Method 3: the symmetric case.

ρ_0	$\rho_1 = \rho_2$	α	$\mathbb{E}N_0^{sim}$	$\mathbb{E}N_0^{M3}$	$\mathbb{E}N_1^{sim}$	$\mathbb{E}N_1^{M3}$	$\mathbb{E}N_2^{sim}$	$\mathbb{E}N_2^{M3}$
0.3	0.69	1	59.40	59.40	70.06	69.24	67.58	69.35
0.5	0.49	1	95.19	98.77	45.80	49.31	48.46	49.46
0.7	0.29	1	143.78	137.40	31.01	29.35	29.58	29.53
0.3	0.69	2	50.17	50.76	71.83	80.75	73.27	86.63
0.5	0.49	2	90.58	90.18	55.61	60.76	56.49	66.64
0.7	0.29	2	127.79	131.24	33.86	37.57	33.78	41.86
0.3	0.69	5	49.89	45.63	75.98	87.58	81.50	96.88
0.5	0.49	5	85.94	84.59	56.51	68.21	61.03	77.82
0.7	0.29	5	127.15	126.97	35.20	43.27	39.39	50.40
0.3	0.69	∞	50.20	43.42	77.31	90.54	83.68	101.31
0.5	0.49	∞	84.83	82.33	59.16	71.22	58.67	82.33
0.7	0.29	∞	130.52	125.31	37.72	45.37	40.41	53.71

Table 8: Results for Method 3: the asymmetric case.