



Centrum voor Wiskunde en Informatica

REPORT*RAPPORT*

PNA

Probability, Networks and Algorithms



Probability, Networks and Algorithms

Bandwidth-sharing networks under a diffusion scaling

U. Ayesta, M.R.H. Mandjes

REPORT PNA-R0810 JUNE 2008

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)

Copyright © 2008, Stichting Centrum voor Wiskunde en Informatica
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

Bandwidth-sharing networks under a diffusion scaling

ABSTRACT

This paper considers networks operating under alpha-fair bandwidth sharing. When imposing a peak rate (i.e., an upper bound on the users' transmission rates, which could be thought of as access rates), the equilibrium point of the fluid limit is explicitly identified, for both the single-node network as well as the linear network. More specifically, a criterion is derived that indicates, for each specific class, whether or not it is essentially transmitting at peak rate. Knowing the equilibrium point of the fluid limit, the steady-state behavior under a diffusion scaling is determined. This allows an explicit characterization of the correlations between the number of flows of the various classes.

2000 Mathematics Subject Classification: 90B18, 60K25

Keywords and Phrases: bandwidth sharing, fluid limit, diffusion

Bandwidth-sharing networks under a diffusion scaling*

Urtzi Ayesta[†] and Michel Mandjes*

[†]LAAS-CNRS

7 Avenue Colonel Roche
31077 Toulouse Cedex, France
`urtzi@laas.fr`

*University of Amsterdam, Korteweg-de Vries Institute

Plantage Muidergracht 24
1018 TV Amsterdam, the Netherlands
`mmandjes@science.uva.nl`

Abstract

This paper considers networks operating under α -fair bandwidth sharing. When imposing a peak rate (i.e., an upper bound on the users' transmission rates, which could be thought of as access rates), the equilibrium point of the fluid limit is explicitly identified, for both the single-node network as well as the linear network. More specifically, a criterion is derived that indicates, for each specific class, whether or not it is essentially transmitting at peak rate. Knowing the equilibrium point of the fluid limit, the steady-state behavior under a diffusion scaling is determined. This allows an explicit characterization of the correlations between the number of flows of the various classes.

Keywords: Bandwidth-Sharing network, linear-network, access-rate limit, diffusion scaling

*Part of this research was done while both authors were with CWI, Amsterdam, the Netherlands

1 Introduction

Over the past decade, bandwidth-sharing models have been proposed for describing the flow-level performance of elastic jobs competing for resources. The framework presented by Massoulié and Roberts (2002) covers a broad range of interesting allocations — see also the class of α -fair sharing policies introduced by Mo and Walrand (2000); by varying a single parameter (i.e., $\alpha \in (0, \infty)$), one retrieves for instance max-min fairness and proportional fairness as special cases. Interestingly, for single-hop scenarios these policies reduce to discriminatory processor sharing (which is an important generalization of the traditional, egalitarian processor sharing discipline). For multi-hop scenarios, the interaction between the flows (with different routes through the network, and the flows along each route have their own specific stochastic characteristics, such as the flow arrival-rate and flow-size distribution) is intrinsically more complex.

It is not straightforward to evaluate the performance of networks operating under α -fair bandwidth sharing. A first key result was obtained by Bonald and Massoulié (2001), who use fluid-limit techniques [Dai (1996)] to show that these networks are stable under the (plausible) condition that, on each link, the load imposed is strictly smaller than the link capacity. A limited number of results is available on the (joint) distribution of the number of flows present on each route; in this respect we mention the bounds found by Bonald and Proutière (2004), relying on the balanced-fairness concept [Bonald and Proutière (2003)], as well as the approximations of Kelly and Williams (2004) for the regime in which at least one resource is critically loaded. The latter study considers the case of Poisson arrival of flows and exponentially distributed flow sizes (which are quite common assumptions in the literature on α -fair networks), whereas the bounds of Bonald and Proutière (2004) are insensitive (i.e, they apply for any flow-size distribution with the same mean). Partial results and approximations for (mean) flow-level delays have been reported for instance in Lieshout, Borst and Mandjes (2006).

The framework of α -fair sharing policies is often used to describe flow-level behavior in communication networks, for instance those operating under TCP (take $\alpha = 2$). The approach assumes that there is a *separation of timescales*, in that the transmission rates adapt instantly as soon as the network population changes. Packet-level effects are not taken into account, but one could do so by relying on throughput models for TCP [Abendroth, van den Berg and Mandjes (2006), Gibbens *et al.* (2000)].

As the critically loaded regime that was considered in Kelly and Williams (2004) is not always realistic, one may wonder whether any analysis is possible when relaxing this assumption. Under exponentiality assumptions, the network population follows a Markov chain, but an explicit solution of its equilibrium distribution is not known. Therefore, a natural next step

is to find out whether analysis of some specific asymptotic regime is possible; one could for instance study a so-called *fluid scaling* in which the arrival as well as service processes is sped up by a factor L (corresponding to multiplying the Poisson arrival rate as well as the link capacities by L), where L is some large number. At first glance, however, this fluid scaling does not look very promising for analyzing α -fair networks, as can be illustrated by the easiest example possible: a single node used by a single route. Then the model reduces to a M/M/1 queue operating under processor sharing (PS); the flow arrival rate is $L\lambda$, the flows have sizes that are i.i.d. samples from an exponential distribution with mean μ^{-1} , and the link has capacity LC (where we assume that $\rho := \lambda/\mu < C$). The number of flows present is a Markov chain, with rate $L\lambda$ upwards and rate $L\mu C$ downwards. Popularly speaking, the equilibrium of the fluid limit corresponds to the state of the Markov chain in which the upward rate and the downward rate match, which obviously does not happen for this model: for any state there is a net drift downwards. (To compare, consider also the M/M/ ∞ queue under the same scaling. Then the rate upwards is $L\lambda$, and the rate downwards, from state n , is $Ln\mu$. This yields the equilibrium point $n^* = \lambda/\mu$, as could be expected from the well-known fact that the number of flows present has, in equilibrium, a Poisson distribution with mean $L\lambda/\mu$. In other words: in the case of the M/M/ ∞ system the fluid limit approach gives more meaningful information than for M/M/1-PS).

There is, however, an interesting remedy to the shortcomings of the fluid-limit approach. The fact that in the above example (of a single M/M/1-PS queue) no equilibrium point of the fluid limit exists is a direct consequence of the unlimited service rate that each user can potentially claim; when there is just one flow present in the system, it is served at the full rate LC . Suppose, however, that the transmission rate of a single user is restricted to some rate r (which could be thought of as an access rate), the equilibrium of the fluid limit follows from the equation $L\lambda = \mu \min\{LC, Lrn\}$, which has solution $n^* = \lambda/(r\mu)$. Informally, this means that the average number of flows present is about Ln^* . We observe that imposing a maximum access rate (a ‘peak rate’) has the effect that the equilibrium point of the fluid limit becomes well-defined. It is noted that, from a more practical point of view, the assumption of finite access rates is very natural; in fact, the possibility of a single user claiming the entire link capacity was a less realistic feature of standard α -fair networks.

One of the main contributions of the present paper is the characterization of the equilibrium point for the fluid limit, for the situation in which all the routes are imposed a peak rate. For the case of a single node, we prove the existence and uniqueness of the equilibrium point and we explicitly characterize it. For the important class of linear networks, we prove existence of the equilibrium point. In the above example, with a single node used by a single route, the equilibrium point of the fluid limit does not depend on the link capacity C . When considering a scenario with multiple routes, however, our results indicate that for some routes the link

capacity does not play a role (the ‘peak-rate constrained routes’), while for others it does.

Having identified the equilibrium point of the fluid limit, a next step is to analyze the deviations around it. The general procedure is to consider the so-called *diffusion scaling*: consider the sped-up Markov chain, subtract L times the equilibrium point of the fluid limit, and divide by \sqrt{L} . One expects from central-limit type of arguments that the resulting process has a stationary distribution that converges to a multivariate Normal distribution (with mean 0, and a certain covariance matrix) when L grows large. A second significant contribution of our paper is that, for the models mentioned above (single node and linear network), we explicitly calculate the corresponding covariance matrix.

Knowledge of the covariance matrix provides the answer to several intriguing questions. Consider for instance a linear network with two nodes shared by three routes: route 0 goes through both, while route i uses just resources at node i , $i = 1, 2$. Suppose we observe that many flows are present on route 1, and we wonder what this says about the number of flows on route 2. One possible line of reasoning is: the fact that there are many flows on route 1 is a consequence of an unusually large number of route-0 flows, and therefore the number of route-2 flows will be large as well. One could also argue, however, that the fact that there are many route-1 flows implies that the route-0 flows obtain less resources than usual at node 1, and consequently also at node 2, which is beneficial for route 2 and leads to a relatively low number of route-2 flows. The results of this paper show which of these two effects is dominant.

Fluid limits and the diffusion scaling have been used extensively before, also in the context of the performance analysis of communication networks; we mention, for example, studies on the interaction between elastic and streaming traffic [Bonald and Massoulié (2001), Kelly and Williams (2004), Kumar and Massoulié (2005)] and also (measurement-based) admission control schemes [Key and Massoulié (2003)].

The rest of the paper is organized as follows. In Section 2 we present the general methodology for the single-link case with different access rates. More specifically, we identify the equilibrium point of the fluid limit, and determine the covariance matrix corresponding to the diffusion scaling. We also derive an explicit criterion that reveals which routes are peak-rate constrained and which are not. In Section 3 we answer the same set of questions for a linear network. We also shed light on the question raised above on the correlations in a two-node network shared by three routes. Section 4 concludes the paper.

2 Single node

In this section we consider a single network resource of capacity \mathcal{C} shared by M routes; as we are focusing on a single node, the word ‘route’ is perhaps somewhat unnatural, and therefore we will use ‘class’ instead. We first introduce the model without peak rates, which turns out to be equivalent to discriminatory processor sharing, and then impose the access rate limitations.

Let $N_i(t)$ be the number of class- i jobs present at time t , for $t \geq 0$ and $i = 1, \dots, M$. Flows of class i arrive according to a Poisson process with rate λ_i , $i = 1, \dots, M$. The transfer sizes are exponentially distributed with mean μ_i^{-1} , $i = 1, \dots, M$. The α -fair allocation at time t follows from, for $\alpha \in (0, \infty)$ and weights $\bar{g}_1, \dots, \bar{g}_M > 0$,

$$\max_{\vec{y}} \sum_{i=1}^M N_i(t) \bar{g}_i \frac{y_i^{1-\alpha}}{1-\alpha}, \quad \text{under} \quad \sum_{i=1}^M N_i(t) y_i \leq \mathcal{C}; \quad (1)$$

here y_i is the amount of bandwidth allocated to a single flow of class i (for $\alpha = 1$ we optimize $\sum_{i=1}^M N_i(t) \bar{g}_i \log y_i$). The solution of the optimization problem is the allocation

$$y_i = \frac{\bar{g}_i^{1/\alpha} \mathcal{C}}{\sum_{j=1}^M \bar{g}_j^{1/\alpha} N_j(t)}.$$

We observe that we can equivalent say that the resource is shared according to a vector of weights $g_i > 0$, $i = 1, \dots, M$ in a discriminatory-processor-sharing fashion [Kleinrock (1967), Fayolle, Mitrani and Iasnogorodski (1980)], with $g_i = \bar{g}_i^{1/\alpha}$. In other words, the total capacity attributed to a class- i transfer is

$$C_i(t) = \frac{g_i \mathcal{C}}{\sum_{j=1}^M g_j N_j(t)},$$

at time t ; without loss of generality we assume $\sum_{j=1}^M g_j = 1$. Hence, such a discriminatory-processor-sharing model could correspond to the interaction of flows of several classes in one router of a network. It is noted that this model could for instance be used to examine the impact of heterogeneous round-trip-times on bandwidth sharing [Altman, Jimenez and Kofman (2004)]: in TCP the classes with lower round-trip-times obtain a higher share of the bandwidth, which could be modelled by granting them a higher weight.

As argued in the introduction, the above rate-allocation mechanism has some flaws, at least from a practical perspective: in principle any source can grab the entire bandwidth \mathcal{C} , while in practice access rates impose an upper bound on the rate allocated to a single flow. Therefore, we suppose that in addition the rate of each class- i job is constrained by an access-link rate

limitation $r_i = 1, \dots, M$. As a consequence, the *effective* rate of a class- i transfer becomes

$$C_i(t) = \min \left\{ \frac{g_i \mathcal{C}}{\sum_{j=1}^M g_j N_j(t)}, r_i \right\}, \quad (2)$$

$i = 1, \dots, M$.

Let $\varrho_i := \lambda/(\mu_i \mathcal{C})$ denote the load imposed by class i . We assume that the total load is smaller than the resource capacity, i.e., $\sum_{i=1}^M \varrho_i < 1$. The process

$$\vec{N} \equiv (\vec{N}(t))_{t \geq 0} = (N_1(t), \dots, N_M(t))_{t \geq 0}$$

is a continuous-time Markov chain with state space \mathbb{N}^M , equipped with transition rates

$$\begin{cases} (\vec{N}(t)) \rightarrow (\vec{N}(t) + \vec{e}_i) : \lambda_i, \\ (\vec{N}(t)) \rightarrow (\vec{N}(t) - \vec{e}_i) : \mu_i N_i(t) C_i(t), \text{ if } N_i(t) \geq 1, \end{cases}$$

where \vec{e}_i denotes the unit vector with 1 in component i , and 0 otherwise, $i = 1, \dots, M$.

Remark 2.1 It is noted that $\sum_{i=1}^M \varrho_i < 1$ implies that the continuous-time Markov chain $(\vec{N}(t))_{t \geq 0}$ is ergodic. This can be shown by using the result that states that if the transition probabilities of an ergodic Markov chain are modified on a finite subset, then the resulting Markov chain remains ergodic; this ‘folk theorem’ was rigorously proved by Leskelä (2004) (in his Lemma 3), but see also Meyn and Tweedie (1993). Note that

$$\frac{g_i N_i(t)}{\sum_{j=1}^M g_j N_j(t)} \leq 1,$$

$i = 1, \dots, M$, thus from (2) it follows that if $n_i > \mathcal{C}/r_i$, class- i will not be peak-rate limited; conclude that the transition rates differ from the DPS rates only on a finite set of states. Since the single-server DPS system is stable if $\sum_{i=1}^M \varrho_i < 1$, it follows that the peak-rate limited system will be stable as well under the same condition on the loads.

Similarly, for any peak-rate limited variant of an α -fair network it can be shown that there is stability as long as *per node* the load imposed is smaller than the available link capacity. For the system without peak-rate limitation this was proved by Bonald and Massoulié (2001), and given the fact that peak-rate limitation corresponds to the modification of the rates on a finite subset of the state space, the result in Leskelä (2006) entails stability of the peak-rate constrained system. In particular, the linear network discussed in Section 3 is stable under the ‘usual’ condition.

Remark 2.2 The peak-rate restriction can be imposed in several ways. Above we first computed the optimal allocation (in that the objective function in (1) was maximized), and then

truncated the resulting rates at the access rates r_i . An alternative, for instance, is to determine the rates by solving

$$\max_{\vec{y}} \sum_{i=1}^M N_i(t) \bar{g}_i \frac{y_i^{1-\alpha}}{1-\alpha}, \quad \text{under} \quad \sum_{i=1}^M N_i(t) y_i \leq \mathcal{C}, \quad y_i \in [0, r_i], \quad i = 1, \dots, M. \quad (3)$$

We note that the allocation (3) will waste resources only if $\sum_{i=1}^M N_i(t) r_i \leq \mathcal{C}$, whereas allocation (1) may waste resources even when $\sum_{i=1}^M N_i(t) r_i > \mathcal{C}$. Therefore, allocation (3) is Pareto efficient, while (1) is not. In this paper we chose to use (1) rather than (3); it is not in the scope of the paper to verify which of these alternatives is closest to reality. For instance, is not clear whether algorithms such as TCP work in a Pareto-efficient manner. Below we will see that the system under criterion (1) allows fairly explicit analysis; this turns out to be considerably harder under (3).

2.1 Fluid scaling

When considering the fluid scaling process, one essentially speeds up time. In our case it means that the arrivals of flows occur more frequently, but at the same time the link rate becomes faster, thus maintaining the same load. More concretely, we replace the arrival rates λ_i by $L\lambda_i$, and the service rate is sped up in the same way, i.e., \mathcal{C} is replaced by $L\mathcal{C}$. Calling the resulting Markov chain $(\vec{N}^{(L)}(t))_{t \geq 0}$, we assume that the normalized process $(L^{-1} \cdot \vec{N}^{(L)}(t))_{t \geq 0}$ converges to a deterministic limit, characterized by the differential equations, $i = 1, \dots, M$,

$$n'_i(t) = \lambda_i - \phi_i(\vec{n}(t)), \quad \text{where} \quad \phi_i(\vec{n}) := \mu_i \min \left\{ \frac{g_i n_i \mathcal{C}}{\sum_{j=1}^M g_j n_j}, r_i n_i \right\}; \quad (4)$$

see Hunt and Kurtz (1991) and Kumar and Massoulié (2005) for further background. Therefore, the ‘equilibrium point’, i.e., the vector \vec{n}^* that solves $\lambda_i = \phi_i(\vec{n})$ for all $i = 1, \dots, M$, is the long-run limit (with probability 1). Our first goal is to prove that there is one such a vector.

As mentioned above, the equilibrium point of the fluid limit solves $\lambda_i = \phi_i(\vec{n})$, for $i = 1, \dots, M$. First observe that this system of equations is equivalent to, for $i = 1, \dots, M$,

$$n_i = \max \left\{ \frac{\lambda_i}{g_i} \sum_{j=1}^M g_j n_j, \frac{\lambda_i}{r_i \mu_i} \right\}. \quad (5)$$

Define $\gamma_k := g_k / r_k$, $k = 1, \dots, M$. Without loss of generality, we relabel the classes in a nondecreasing way with respect to γ_k , that is, $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_M$.

We call a class i *peak-rate constrained* if the vector \vec{n}^* is such that

$$\frac{g_i \mathcal{C}}{\sum_{j=1}^M g_j n_j^*} > r_i.$$

Let \mathcal{S} denote the subset of classes that are peak-rate constrained and let \mathcal{S}^C denote the other classes. In Proposition 2.3 we show that there is a unique solution to the system of equations (5), and, interestingly, that the classes with the highest γ_i are the candidates for being peak-rate constrained.

Proposition 2.3 *There exists a unique $s \in \{1, \dots, M\}$ such that $\mathcal{S} = \{1, \dots, s\}$. The solution to the system of equations (5) is unique and is given by $\vec{n}^* = (n_1^*, n_2^*, \dots, n_M^*)$, where for $i \in \mathcal{S}$,*

$$n_i^* \equiv n_i^*(s) = \lambda_i / (r_i \mu_i),$$

and for $i \in \mathcal{S}^C$,

$$n_i^* \equiv n_i^*(s) = \frac{\varrho_i}{g_i} \left(\frac{1}{1 - \sum_{j=s+1}^M \varrho_j} \right) \left(\sum_{j=1}^s \frac{g_j \lambda_j}{r_j \mu_j} \right). \quad (6)$$

Proof We prove the result in three steps.

Step 1. We first show that if there is a solution to (5), then \mathcal{S} is non-empty. Suppose that this would not hold, i.e., for all $i = 1, \dots, M$,

$$\frac{g_i n_i^*}{\sum_j g_j n_j^*} = \frac{\lambda_i}{\mu_i C} = \varrho_i.$$

Summing the left-most terms over $i = 1, \dots, M$, we obviously obtain 1, whereas the right-most terms sum to a value strictly smaller than 1 (which is due to the stability constraint $\sum_i \varrho_i < 1$). We thus obtain a contradiction. We conclude that at least one class is peak-rate constrained.

Step 2. Suppose for the moment that \mathcal{S} is of the form $\{1, \dots, s\}$, for some $s \in \{1, \dots, M\}$. Then, for any $i \in \mathcal{S}$ it is easy to see that $n_i^*(s) = \lambda_i / (r_i \mu_i)$. For $i \in \mathcal{S}^C$, we have that $g_i n_i^* = \varrho_i \sum_{j=1}^M g_j n_j^*$. Now summing over all $i \in \mathcal{S}^C$ yields

$$\sum_{i=s+1}^M g_i n_i^* = \sum_{i=s+1}^M \varrho_i \left(\sum_{j=1}^M g_j n_j^* \right) = \sum_{i=s+1}^M \varrho_i \left(\sum_{j=s+1}^M g_j n_j^* + \sum_{j=1}^s \frac{g_j \lambda_j}{r_j \mu_j} \right).$$

This immediately implies that

$$\sum_{i=s+1}^M g_i n_i^* = \left(\frac{\sum_{j=s+1}^M \varrho_j}{1 - \sum_{j=s+1}^M \varrho_j} \right) \left(\sum_{j=1}^s \frac{g_j \lambda_j}{r_j \mu_j} \right).$$

Then equation (6) follows from $n_i^* = (\varrho_i / g_i) \sum_{j=1}^M g_j n_j^*$, for all $i \in \mathcal{S}^C$.

Step 3. Now it remains to prove that \mathcal{S} is indeed of the form $\{1, \dots, s\}$ for a unique $s \in \{1, \dots, M\}$. We do this by showing that there is unique s for which, for all $i \in \{1, \dots, s\}$ and

$j \in \{s+1, \dots, M\}$,

$$\gamma_i > \frac{\sum_{k=1}^M g_k n_k^*(s)}{\mathcal{C}} > \gamma_j; \quad (7)$$

if $s = M$, then evidently only the first inequality needs to be met. Using equation (6), we have that

$$\frac{\sum_{j=1}^M g_j n_j^*(s)}{\mathcal{C}} = \frac{\sum_{j=1}^s \varrho_j \gamma_j}{1 - \sum_{j=s+1}^M \varrho_j}. \quad (8)$$

Because the γ_i are non-increasing in i , and in view of identity (8), requirement (7) is equivalent to proving that there is a unique $s \in \{1, \dots, M\}$ such that

$$\gamma_s \stackrel{(i)}{>} \frac{\sum_{j=1}^s \varrho_j \gamma_j}{1 - \sum_{j=s+1}^M \varrho_j} \stackrel{(ii)}{>} \gamma_{s+1};$$

again, if $s = M$, then only inequality (i) needs to be met. Inequality (i) can be rewritten as

$$\gamma_s \left(1 - \sum_{j=s}^M \varrho_j \right) > \sum_{j=1}^{s-1} \varrho_j \gamma_j,$$

whereas inequality (ii) obviously reads

$$\gamma_{s+1} \left(1 - \sum_{j=s+1}^M \varrho_j \right) < \sum_{j=1}^s \varrho_j \gamma_j.$$

In other words, if we denote, for given s , condition (i) by $C(s)$ (so $C(s) \in \{\text{TRUE}, \text{FALSE}\}$), then condition (ii) corresponds to $\neg C(s+1)$. In other words: we have to prove that there is a unique s such that both $C(s)$ and $\neg C(s+1)$.

To this end, first observe that $C(1)$ reduces to the condition $\gamma_1 \varrho_1 < \gamma_1 (1 - \sum_{j=2}^M \varrho_j)$, which is true due to the stability condition. This means that $C(1)$ holds, but this could be expected as we already proved that \mathcal{S} is non-empty. Now we show that $\neg C(s)$ implies $\neg C(s+1)$. If $\neg C(s)$, then, by definition, $\gamma_s (1 - \sum_{j=s}^M \varrho_j) < \sum_{j=1}^{s-1} \varrho_j \gamma_j$. This implies that

$$\gamma_s \left(1 - \sum_{j=s}^M \varrho_j \right) + \varrho_s \gamma_s < \sum_{j=1}^s \varrho_j \gamma_j.$$

Now $\neg C(s+1)$ follows immediately from $\gamma_{s+1} \leq \gamma_s$.

The above arguments imply that there cannot be more than just one s for which $C(s)$ and $\neg C(s+1)$, which concludes the proof of the uniqueness. \square

Interestingly, the above result entails that the order in which classes are candidate for becoming peak-rate constrained is exclusively determined by the γ_i , i.e., the ratio of the weight and access rate. In other words: we find the (at first glance perhaps somewhat surprising) result that neither the arrival rates nor the mean flow-size plays a role here.

We now prove that the process returns to \vec{n}^* when it is sufficiently close to the equilibrium point. We rely on the concept of the ‘linearized system’. The idea is that we linearize the system $n'_i(t) = \lambda_i - \phi_i(\vec{n}(t))$ around \vec{n}^* , and argue that the process returns to \vec{n}^* after small perturbations. We introduce the M -dimensional function $\vec{m}(t) = \vec{n}(t) - \vec{n}^*$, and determine a matrix $P \equiv (p_{ij})_{i,j=1}^M$ such that

$$\vec{m}'(t) = -P\vec{m}(t). \quad (9)$$

By linearizing $\vec{n}(t)$ we obtain

$$m'_i(t) = -\sum_{j=1}^M m_j(t) p_{ij} = -\sum_{j=1}^M m_j(t) \left. \frac{\partial \phi_i(\vec{n})}{\partial n_j} \right|_{\vec{n}:=\vec{n}^*}.$$

(Throughout we assume, for $i = 1, \dots, M$, that $g_i C / \sum_j g_j n_j^* \neq r_i$, so that this derivative is well-defined.) We have that for $i \in \mathcal{S}$,

$$\frac{\partial \phi_i}{\partial n_i} = r_i \mu_i; \quad \frac{\partial \phi_i}{\partial n_j} = 0 \text{ for } j \neq i. \quad (10)$$

Similarly, for $i \in \mathcal{S}^C$ we get

$$\frac{\partial \phi_i}{\partial n_i} = \left(\frac{g_i}{\sum_j g_j n_j} - \frac{g_i^2 n_i}{(\sum_j g_j n_j)^2} \right) \mu_i C,$$

and

$$\frac{\partial \phi_i}{\partial n_j} = -\frac{g_i g_j n_i}{(\sum_j g_j n_j)^2} \mu_i C \quad \text{for } j \neq i.$$

Using that, for $i \in \mathcal{S}^C$, it holds that that $\varrho_i \sum_{j=1}^M g_j n_j^* = g_i n_i^*$, we obtain $i \in \mathcal{S}^C$

$$\left. \frac{\partial \phi_i}{\partial n_i} \right|_{\vec{n}:=\vec{n}^*} = \frac{\lambda_i(1 - \varrho_i)}{n_i^*}; \quad \left. \frac{\partial \phi_i}{\partial n_j} \right|_{\vec{n}:=\vec{n}^*} = -\frac{g_j}{g_i} \frac{\lambda_i \varrho_i}{n_i^*}. \quad (11)$$

Notice that the diagonal elements of P are positive. In the next proposition we show that all the eigenvalues of the matrix P are positive (or, more precisely, have a positive real part), which in particular implies the stability of the linearized system (9) [Khalil, (2001)]. As a consequence of this stability, the dynamics of the linearized system will converge towards 0, and thus $\vec{n}(t)$ will converge towards the unique solution \vec{n}^* identified in Proposition 2.3. We note that it generally does *not* imply stability of the *original* system, but that we already established in another way in Remark 2.1.

Proposition 2.4 *All eigenvalues of P have a positive real part.*

Proof Recall that \mathcal{S} is of the form $\{1, \dots, s\}$ and that $\mathcal{S}^C = \{s+1, \dots, M\}$. First observe that $r_i \mu_i$ are eigenvalues (real and positive) for $i = 1, \dots, s$. From the structure of P it immediately follows that the other eigenvalues are the eigenvalues of $P^{(N)} = (p_{ij}^{(N)})_{i,j=1}^N$, where $N := M - s$, with

$$p_{ij}^{(N)} := \begin{cases} \lambda_i(1 - \varrho_i)/n_i^* & \text{for } i = j; \\ -(g_j/g_i)\lambda_i\varrho_i/n_i^* & \text{for } i \neq j. \end{cases}$$

Consider the eigenvalues of this matrix. They solve

$$\begin{aligned} \theta x_i &= - \sum_{j \neq i} \frac{g_j}{g_i} \frac{\lambda_i \varrho_i}{n_i^*} x_j + \frac{\lambda_i(1 - \varrho_i)}{n_i^*} x_i \\ &= - \sum_{j=1}^N \frac{g_j}{g_i} \frac{\lambda_i \varrho_i}{n_i^*} x_j + \frac{\lambda_i}{n_i^*} x_i, \end{aligned}$$

where \vec{x} is the eigenvector corresponding to eigenvalue θ . Now multiply with g_i . Then it follows that θ is also eigenvalue of $\bar{P}^{(N)} = (\bar{p}_{ij}^{(N)})_{i,j=1}^N$, where

$$\bar{p}_{ij}^{(N)} := \begin{cases} \lambda_i(1 - \varrho_i)/n_i^* & \text{for } i = j; \\ -\lambda_i \varrho_i/n_i^* & \text{for } i \neq j. \end{cases}$$

with eigenvector \vec{y} satisfying $g_j x_j = y_j$. Hence we have to verify that all eigenvalues of $\bar{P}^{(N)}$ have a positive real part.

The matrix $\bar{P}^{(N)}$ equals DQ , where

$$Q := \begin{pmatrix} 1 - \varrho_1 & -\varrho_1 & \dots & -\varrho_1 \\ -\varrho_2 & 1 - \varrho_2 & \dots & -\varrho_2 \\ \vdots & \vdots & & \vdots \\ -\varrho_N & -\varrho_N & \dots & 1 - \varrho_N \end{pmatrix},$$

and

$$D := \text{diag} \left\{ \frac{\lambda_1}{n_1^*}, \dots, \frac{\lambda_N}{n_N^*} \right\}.$$

Now note that $\det(DQ - \theta I) = \det(D) \det(Q - \theta D^{-1}) = \det(QD - \theta I)$. Thus it follows that we have to check whether the real part of the eigenvalues of QD , or, equivalently, $(QD)^T$, is positive. The latter matrix is (strictly) diagonally dominant due to the stability condition. As a consequence of Geršgorin's circle theorem [p.344, Horn and Johnson (1985)], it has only

eigenvalues with a positive real part; more precisely, each eigenvalue of $(QD)^T$ (and hence also of QD) is in at least one of the N disks

$$\left\{ z \in \mathbb{C} : \left| z - \frac{\lambda_i(1 - \varrho_i)}{n_i^*} \right| < \frac{\lambda_i}{n_i^*} \sum_{j \neq i} \varrho_j \right\},$$

and hence all eigenvalues are in the right half plane. This proves the stated. \square

Example 2.5 We consider a single server with two classes. For class 1 we choose $\lambda_1 = 0.75$, $\mu_1 = 2$, $g_1 = 2$, and $r_1 = 0.1$, while for class 2 we have $\lambda_2 = 1.5$, $\mu_2 = 4$, $g_2 = 1$ and $r_2 = 0.8$. Note that $\gamma_1 = 20$ and $\gamma_2 = 1.25$. Furthermore, observe that \mathcal{C} should be at least 0.75 to ensure stability.

We saw that in the equilibrium point of the fluid limit always at least one class is peak-rate constrained, and this is the class i with the highest γ_i ; in this case class 1 is apparently peak-rate constrained, and $n_1^* = \lambda_1/(r_1\mu_1) = \frac{15}{4} = 3.75$, irrespective of \mathcal{C} . The value of n_2^* , however, *does* depend on \mathcal{C} . It can be calculated (for instance by equating, for $i = 2$, the two terms between the brackets in (4), with $n_1^* = 3.75$) that for $\mathcal{C} < \frac{51}{8} = 6.375$, we have that class 2 is not peak-rate constrained, and

$$n_2^*(\mathcal{C}) = \frac{45}{16\mathcal{C} - 6};$$

for $\mathcal{C} > 6.375$ both classes are peak-rate constrained, and $n_2^* = \lambda_2/(r_2\mu_2) = \frac{3}{8} = 0.375$.

2.2 Diffusion scaling

In this subsection we develop the diffusion approximation for the steady-state distribution of the normalized process $(L^{-1} \cdot \vec{N}^{(L)}(t))_{t \geq 0}$ around the equilibrium point of the fluid limits \vec{n}^* . Interestingly, the diffusion scaling allows us to determine the correlations between the number of flows present of different classes.

To motivate our diffusion approach, we first return to the example mentioned in the introduction: a one-dimensional birth death process with birth rate $L\lambda$ and death rate (when the network occupancy is k) $\min\{LC, rk\}$; assume for ease that $k^* := L \cdot \mathcal{C}/r \in \mathbb{N}$. It is easily verified that the probability distribution of the steady-state the number of jobs in the system, say $N^{(L)}$, is

$$\mathbb{P}(N^{(L)} = k) =: \pi_k = \left(\frac{\lambda L}{\mu r} \right)^k \frac{1}{k!} \cdot \pi_0$$

for $0 \leq k \leq k^*$, and

$$\pi_k = \left(\frac{\lambda}{\mu \mathcal{C}} \right)^k \frac{1}{(k^*)!} \left(\frac{\mathcal{C} L}{r} \right)^{k^*} \cdot \pi_0$$

for $k > k^*$, where π_0 is obtained through normalization. We mentioned in the introduction that most of the probability mass is around $L \cdot \lambda/(r\mu)$. This can be made precise, in the sense that it can be verified that $\mathbb{E}N^{(L)}/L \rightarrow \lambda/(r\mu)$ as $L \rightarrow \infty$; notice that under $\varrho < \mathcal{C}$ we have that the equilibrium point is *smaller* than k^* . Similarly, through direct arguments it can be seen that

$$Y^{(L)} := \frac{1}{\sqrt{L}}(N^{(L)} - L \cdot \lambda/(r\mu))$$

converges to a Normally distributed random variable Y with mean 0 and variance $\lambda/(r\mu)$ as $L \rightarrow \infty$. The important conclusion here is that, in this scaling, the distribution of Y is essentially determined by the system dynamics around the ‘equilibrium point’ $\lambda/(r\mu)$. Also, the larger L , the less frequent the process will attain values larger than k^* .

Let us now return to our model, and see how to translate the above properties for the one-dimensional case into our multi-dimensional setting. We first introduce the perturbation process

$$\vec{Y}^{(L)}(t) := \frac{1}{\sqrt{L}}(\vec{N}^{(L)}(t) - L \cdot \vec{n}(t)),$$

where $\vec{n}(t)$ is the solution of (4). We are particularly interested in the steady-state version ($t \rightarrow \infty$).

Based on the above considerations for the one-dimensional case, we expect this process to have a diffusion limit as $L \rightarrow \infty$. With the matrix P as defined above, as $L \rightarrow \infty$, this limiting process, say $\vec{Y}(t)$, approximately satisfies the following stochastic differential equation:

$$d\vec{Y}(t) = -P\vec{Y}(t)dt + d\vec{W}(t),$$

where $\vec{W}(t) = A\vec{B}(t)$, with $\vec{B}(t)$ being an M -dimensional vector of independent standard Brownian motions, and A a diagonal matrix with, for $i = 1, \dots, M$,

$$A_{ii} := \sqrt{\lambda_i + \phi_i(\vec{n}^*)} = \sqrt{2\lambda_i};$$

cf. for instance [Section II.C, Key and Massoulié (2003)] and Kelly, Maulloo and Tan (1998). This is an M -dimensional Ornstein-Uhlenbeck process; we note that the process $\vec{Y}(t)$ is a Gauss-Markov process [Section 5.6, Karatzas and Shreve (1991)] and obeys the alternative representation $\vec{Y}(t) = -\int_{-\infty}^t e^{-P(t-s)} A d\vec{B}(s)$.

Let us consider the steady-state version of $\vec{Y}(t)$, i.e., $\vec{Y} := \lim_{t \rightarrow \infty} \vec{Y}(t)$. The covariance matrix is then given by [Equation (6.19), Karatzas and Shreve (1991)]

$$\mathbb{E}(\vec{Y}\vec{Y}^T) = \Sigma := \int_0^\infty e^{-Pt} A A^T e^{-P^T t} dt;$$

where $AA^T = 2\text{diag}\{\vec{\lambda}\}$.

We now compute the matrix Σ for $M = 2$; for higher dimensions the formulas are less clean. First, recall that $\mathcal{S} = \emptyset$ does not occur. On the other hand, the case $\mathcal{S} = \{1, 2\}$ is trivial: $\Sigma = \text{diag}\{\lambda_1/(r_1\mu_1), \lambda_2/(r_2\mu_2)\}$. We therefore concentrate on the more challenging case $\mathcal{S} = \{1\}$. Hence, from Proposition 2.3 we have that $n_1^* = \lambda_1/(r_1\mu_1)$, and $n_2^* = \lambda_1/(r_1\mu_1) \cdot (\rho_2/(1 - \rho_2)) \cdot (g_1/g_2)$. Then applying equations (10) and (11) the P -matrix reads

$$P = \begin{pmatrix} p_{11} & 0 \\ p_{21} & p_{22} \end{pmatrix}, \quad \text{with}$$

$$p_{11} := r_1\mu_1, \quad p_{21} := -(1 - \varrho_2) \cdot \frac{r_1\mu_1}{\lambda_1} \cdot \lambda_2, \quad p_{22} := (1 - \varrho_2)^2 \frac{r_1\mu_1}{\lambda_1} \cdot \mu_2 \mathcal{C} \cdot \frac{g_2}{g_1}.$$

where obviously p_{11} and p_{22} are positive, and p_{21} is negative. Straightforward algebra yields

$$P^k = \begin{pmatrix} p_{11}^k & 0 \\ \sum_{i=0}^{k-1} p_{11}^i p_{21} p_{22}^{k-i-1} & p_{22}^k \end{pmatrix},$$

for $k \geq 1$. Thus, we have that $(e^{-Pt})_{11} = e^{-p_{11}t}$, $(e^{-Pt})_{22} = e^{-p_{22}t}$, and

$$(e^{-Pt})_{21} = \sum_{k=1}^{\infty} \sum_{i=0}^{k-1} p_{11}^i p_{21} p_{22}^{k-i-1} \frac{(-t)^k}{k!} = \frac{p_{21}}{p_{22} - p_{11}} (e^{-p_{22}t} - e^{-p_{11}t}).$$

We note that the matrix e^{-Pt} could also be calculated by the eigendecomposition of the matrix P ; see for example [Equation (A.19), Asmussen (2000)]. We note that $e^{-P^T t} = (e^{-Pt})^T$ and that $AA^T = \text{diag}\{2\lambda_1, 2\lambda_2\}$. Let $\bar{\Sigma}(t)$ denote the symmetric matrix $\bar{\Sigma}(t) := e^{-Pt} AA^T e^{-P^T t}$:

$$\begin{aligned} \bar{\Sigma}_{11}(t) &= 2\lambda_1 e^{-2p_{11}t}, \quad \bar{\Sigma}_{12}(t) = 2\lambda_1 \frac{p_{21}}{p_{22} - p_{11}} (e^{-p_{22}t - p_{11}t} - e^{-2p_{11}t}), \\ \bar{\Sigma}_{22}(t) &= 2\lambda_1 \left(\frac{p_{21}}{p_{22} - p_{11}} \right)^2 (e^{-p_{11}t} - 2e^{-p_{22}t - p_{11}t} + e^{-2p_{22}t}) + 2\lambda_2 e^{-2p_{22}t}. \end{aligned}$$

After integrating $\bar{\Sigma}(t)$ over t (componentwise), we eventually obtain:

$$\Sigma_{11} = \frac{\lambda_1}{p_{11}} = \frac{\lambda_1}{r_1\mu_1}, \quad \Sigma_{21} = \Sigma_{12} = -\frac{\lambda_1 p_{21}}{p_{11}(p_{11} + p_{22})}, \quad \Sigma_{22} = \frac{\lambda_1 p_{21}^2}{p_{11}p_{22}(p_{11} + p_{22})} + \frac{\lambda_2}{p_{22}}. \quad (12)$$

Observe that $\Sigma_{12} > 0$, as expected; we conclude that there is a positive correlation between the numbers of flows of class 1 and 2. It can be expected that for higher dimensions the correlation between two arbitrary classes will be positive as well.

Example 2.6 Return to the parameters of Example 2.5. We here specialize to $\mathcal{C} = 21/16 = 1.3125$, and consider the ‘ L -sped up model’ (i.e., arrival rate and service rate multiplied by L). Then $n_2^* = 3$, whereas n_1^* equals, as before, 3.75. Recall that $\vec{Y} := \lim_{t \rightarrow \infty} \vec{Y}(t)$, with

$$\vec{Y}(t) := \frac{1}{\sqrt{L}} (\vec{N}^{(L)}(t) - L \cdot \vec{n}^*),$$

converges to a zero-mean bivariate Normal random variable, with covariance matrix given by (12). Informally, one could say that $\vec{N}^{(L)} := \lim_{t \rightarrow \infty} \vec{N}^{(L)}(t)$ is approximately distributed as a bivariate Normal random variable with mean $L \cdot \vec{n}^*$ and covariance matrix $L \cdot \Sigma$. It can be verified that in this case

$$P = \begin{pmatrix} \frac{1}{5} & 0 \\ -\frac{2}{7} & \frac{5}{14} \end{pmatrix}; \quad \Sigma = \begin{pmatrix} \frac{15}{4} & \frac{25}{13} \\ \frac{25}{13} & \frac{2011}{455} \end{pmatrix} \approx \begin{pmatrix} 3.75 & 1.92 \\ 1.92 & 4.42 \end{pmatrix}.$$

As we have remarked in Example 2.5, in the fluid-limit regime class 1 is peak-rate constrained, and class 2 is not. Suppose that at some point in time there are n_i flows of type i present ($i = 1, 2$), then flows of type i are transmitting at peak rate if $g_1 n_1 + g_2 n_2 < \gamma_i LC$. It could be expected that most of the time the process is in a regime in which class 1 is peak-rate constrained, and class 2 is not, as

$$g_1 n_1 + g_2 n_2 < \gamma_1 LC \quad \text{and} \quad g_1 n_1 + g_2 n_2 > \gamma_2 LC.$$

To verify this, suppose the diffusion approximation is accurate. Observe that this implies that $\mathbb{V}\text{ar}(g_1 N_1^{(L)} + g_2 N_2^{(L)}) = L \cdot (g_1^2 \Sigma_{11} + 2g_1 g_2 \Sigma_{12} + g_2^2 \Sigma_{22})$. The fraction of time that both classes are transmitting at peak rate is roughly, with $\mathcal{N}(\mu, \sigma^2)$ denoting a Normally distributed random variable with mean μ and variance σ^2 ,

$$\begin{aligned} \mathbb{P}(g_1 N_1^{(L)} + g_2 N_2^{(L)} < \gamma_2 LC) &\approx \mathbb{P}\left(\mathcal{N}\left(\frac{21}{2} \cdot L, \frac{10586}{455} \cdot L\right) < \frac{105}{64} \cdot L\right) \\ &\approx \mathbb{P}(\mathcal{N}(10.5 \cdot L, 23.3 \cdot L) < 1.64 \cdot L). \end{aligned}$$

It is easily verified that this probability equals, with $\Phi(\cdot)$ denoting the standard Normal cumulative distribution function, $\Phi(-1.84 \cdot \sqrt{L})$. For instance, if we choose $L = 10$, then this gives a probability in the order of $4 \cdot 10^{-9}$. Conversely, one can estimate the fraction of time that both classes are transmitting at a rate higher than their access rate by computing $\mathbb{P}(g_1 N_1^{(L)} + g_2 N_2^{(L)} > \gamma_1 LC)$, which approximately equals $1 - \Phi(3.27 \cdot \sqrt{L})$, which is also extremely small already for moderate values of L . Hence we can safely conclude that the process is, virtually all the time, in a regime in which class 1 transmits at peak rate, while class 2 is transmitting at a rate higher than its peak rate.

3 Linear network

In this section we consider so-called linear networks. These consist of multiple nodes, say M , which we assume to have equal capacity \mathcal{C} . There are $M + 1$ classes sharing these nodes. There is one ‘common class’: class 0 goes through all nodes and gets served simultaneously at all of them (at the same rate). Then there are M ‘crossing classes’: class i , $i = 1, \dots, M$,

goes just through node i . This type of networks has been widely studied in the past; in the setting of Mo and Walrand (2000), one needs to solve

$$\max_{\vec{y}} \sum_{i=0}^M N_i(t) g_i \frac{y_i^{1-\alpha}}{1-\alpha}, \quad \text{under } N_0(t)y_0 + N_i(t)y_i \leq \mathcal{C}, \quad i = 1, \dots, M,$$

if at time t the network population is $\vec{N}(t)$. It is readily verified that the solution to the above optimization problem is

$$y_0 = \frac{g_0}{g_0 N_0(t) + S_\alpha(\vec{N}(t))}, \quad \text{where } S_\alpha(\vec{n}) := \left(\sum_{j=1}^M g_j n_j^\alpha \right)^{1/\alpha};$$

the value of y_i can be readily obtained from the capacity constraint.

The setting above does not involve limitations by access rates. When taking these into account, one could argue that the service rate attributed to a flow of class 0 at time t is approximately

$$C_0(t) = \min \left\{ \frac{g_0 \mathcal{C}}{g_0 N_0(t) + S_\alpha(\vec{N}(t))}, r_0 \right\},$$

whereas the allocation to a class $i, i = 1, \dots, M$ is (for $N_i(t) > 0$)

$$C_i(t) = \frac{1}{N_i(t)} \min \left\{ \frac{S_\alpha(\vec{N}(t)) \mathcal{C}}{g_0 N_0(t) + S_\alpha(\vec{N}(t))}, r_i N_i(t) \right\}. \quad (13)$$

These service rates give rise to a continuous-time Markov chain in a natural way if one assumes Poisson arrivals and exponentially distributed flows; we impose the stability condition $\varrho_0 + \varrho_i < 1$, for all $i = 1, \dots, M$, with $\varrho_i := \lambda_i / (\mu_i \mathcal{C})$, for $i = 0, \dots, M$. To identify the equilibrium point of the fluid limit, we have to solve the system $\lambda_i = \phi_i(\vec{n})$, $i = 0, \dots, M$, with

$$\begin{aligned} \phi_0(\vec{n}) &:= \mu_0 \min \left\{ \frac{g_0 n_0 \mathcal{C}}{g_0 n_0 + S_\alpha(\vec{n})}, r_0 n_0 \right\}; \\ \phi_i(\vec{n}) &:= \mu_i \min \left\{ \frac{S_\alpha(\vec{n}) \mathcal{C}}{g_0 n_0 + S_\alpha(\vec{n})}, r_i n_i \right\} \end{aligned}$$

for $i = 1, \dots, M$. As in the single link case, we can show that there is a unique fixed point \vec{n}^* . Interestingly, imposing the (weak) additional assumption that all ϱ_i are different, for $i = 1, \dots, M$, we can also see that the solution can have just three specific forms.

Proposition 3.1 *Assume that $\varrho_i \neq \varrho_j$ for all $i \neq j$ and $i, j = 1, \dots, M$. Then the solution to the system of equations $\lambda_i = \phi_i(\vec{n})$, $i = 0, \dots, M$, has three possible forms:*

- A ‘crossing class’ is binding: $\mathcal{S}^C = \{i^*\}$ for some $i^* \in \{1, \dots, M\}$ and $\mathcal{S} = \{0, 1, \dots, M\} \setminus i^*$. We have that for $i \in \mathcal{S}$

$$n_i^* = \lambda_i / (r_i \mu_i),$$

and for $i = i^*$

$$n_{i^*}^* = \left(\frac{\varrho_{i^*}}{(1 - \varrho_{i^*})g_{i^*}} \left(\frac{g_0 \lambda_0}{r_0 \mu_0} \right)^\alpha - \frac{1}{g_{i^*}} \sum_{i=1, i \neq i^*}^M \frac{g_i \lambda_i}{r_i \mu_i} \right)^{1/\alpha}.$$

- The ‘common class’ is binding: $\mathcal{S}^C = \{0\}$ and $\mathcal{S} = \{1, \dots, M\}$. Then for $i = 1, \dots, M$

$$n_i^* = \lambda_i / (r_i \mu_i),$$

and

$$n_0^* = \frac{\eta_0}{g_0} \left(\sum_{j=1}^M g_j \left(\frac{\lambda_j}{r_j \mu_j} \right)^\alpha \right)^{1/\alpha}.$$

- All classes are peak-rate constrained: $n_i^* = \lambda_i / (r_i \mu_i)$, for all $i = 0, \dots, M$.

Proof We first show that $i^* \in \mathcal{S}^C$ implies that the other classes are in \mathcal{S} . We have that

$$\varrho_{i^*} = \frac{S_\alpha(\vec{n})}{g_0 n_0 + S_\alpha(\vec{n})}.$$

It is now immediate from $\varrho_0 + \varrho_{i^*} < 1$ that $0 \in \mathcal{S}$. Also, the fact that $\varrho_i \neq \varrho_{i^*}$ for $i = 1, \dots, M$, $i \neq i^*$ implies that

$$\varrho_i \neq \frac{S_\alpha(\vec{n})}{g_0 n_0 + S_\alpha(\vec{n})}$$

for these i , and hence they are also in \mathcal{S} . The expressions for n_i^* , $i = 0, \dots, M$ are obtained directly from the system of equations $\lambda_i = \phi_i(\vec{n})$, $i = 0, \dots, M$.

Similarly, we now show that $0 \in \mathcal{S}^C$ implies that the other classes are in \mathcal{S} . Directly from $\varrho_0 + \varrho_i < 1$, we have that if $0 \in \mathcal{S}^C$, then $1, \dots, M \in \mathcal{S}$. The expressions for n_i^* , $i = 0, \dots, M$ are obtained from the system of equations $\lambda_i = \phi_i(\vec{n})$, $i = 0, \dots, M$.

The third case is straightforward. \square

In the next subsections we consider the first two possible solutions of Proposition 3.1 separately; in particular, as we did for the single-server case, we analyze the system under a diffusion scaling. More particularly, we explicitly compute the covariance matrix. In the third possibility, the classes behave essentially independently.

3.1 A ‘crossing class’ is binding

In this subsection we consider the case $\mathcal{S}^C = \{i^*\}$. As we did for the single-node case, we can consider the network population under a diffusion scaling: speed up time with a factor L , subtract L times the equilibrium point (as determined in the first part of this section), divide by \sqrt{L} , and let L go to ∞ . We thus obtain a random vector \vec{Y} , being Normally distributed with zero mean. To compute the corresponding covariance matrix Σ , we first consider the linearized system by constructing the matrix P , just as we did in the single-node case. Evidently, for $i \neq i^*$,

$$\frac{\partial \phi_i}{\partial n_i} = r_i \mu_i; \quad \frac{\partial \phi_i}{\partial n_j} = 0 \text{ for } j \neq i.$$

whereas for $i = i^*$,

$$\frac{\partial \phi_{i^*}}{\partial n_0} = -\mu_{i^*} \mathcal{C} \frac{S_\alpha(\vec{n})}{(g_0 n_0 + S_\alpha(\vec{n}))^2} g_0 < 0,$$

and for $i \neq 0$

$$\begin{aligned} \frac{\partial \phi_{i^*}}{\partial n_i} &= \mu_{i^*} \mathcal{C} \frac{g_0 n_0}{(g_0 n_0 + S_\alpha(\vec{n}))^2} \cdot \frac{\partial S_\alpha(\vec{n})}{\partial n_i} \\ &= \mu_{i^*} \mathcal{C} \frac{g_0 n_0}{(g_0 n_0 + S_\alpha(\vec{n}))^2} \cdot g_i \left(\frac{S_\alpha(\vec{n})}{n_i} \right)^{1-\alpha} > 0. \end{aligned}$$

It trivially follows that $P \equiv (p_{ij})_{i,j=0}^M$ has only positive eigenvalues (in fact, the eigenvalues are the diagonal elements). As a consequence the system will return in the direction of the equilibrium point \vec{n}^* after small perturbations.

Our goal is to study the correlation between the number of flows in two classes that do not share a path. The next proposition establishes that the correlation between the crossing class and any other class is negative.

Proposition 3.2 *In the case $\mathcal{S}^C = \{i^*\}$, the correlations between n_i^* and n_j^* for all $i \neq j$ and $i, j = 1, \dots, M$ are negative.*

Proof Let us calculate the covariance matrix Σ . Straightforward computations yield that $(e^{-Pt})_{ii} = e^{-p_{ii}t}$, and, for $i \neq i^*$,

$$(e^{-Pt})_{ii^*} = \frac{p_{ii^*}}{p_{i^*i^*} - p_{ii}} (e^{-p_{i^*i^*}t} - e^{-p_{ii}t});$$

all the other entries are 0. This leads to the following covariance matrix:

$$\Sigma_{i^*0} = -\frac{\lambda_0 p_{i^*0}}{p_{00}(p_{00} + p_{i^*i^*})} > 0, \quad \Sigma_{i^*i^*} = \frac{\lambda_{i^*}}{p_{i^*i^*}} + \sum_{i=0, i \neq i^*}^M \frac{\lambda_i p_{i^*i}^2}{p_{ii} p_{i^*i^*} (p_{ii} + p_{i^*i^*})},$$

$$\Sigma_{i^*i} = -\frac{\lambda_i p_{i^*i}}{p_{ii}(p_{ii} + p_{i^*i^*})} < 0, \text{ for } i \neq 0, i^*, \text{ and } \Sigma_{ii} = \frac{\lambda_i}{p_{ii}}, \text{ for } i \neq i^*;$$

the other covariances are 0. \square

Intuitively, the negative correlation can be explained by the fact that when there are many flows of class i^* , they ‘push down’ the bandwidth allocated to the common class, which turns out to be beneficial for the other crossing classes. This answers the question posed in the introduction: apparently it is beneficial for a specific crossing class if there is an unusually high number of flows of another crossing class. The other argumentation mentioned in the introduction (the number of flows of the crossing class is high, because the number of flows of the common class is high, and therefore also the numbers of flows of the other crossing classes will be high, which would suggest positive correlation) is apparently not valid.

3.2 The ‘common class’ is binding

Let us now consider the situation $\mathcal{S} = \{0\}$. Again we first construct the matrix P . For $i \neq 0$,

$$\frac{\partial \phi_i}{\partial n_i} = r_i \mu_i; \quad \frac{\partial \phi_i}{\partial n_j} = 0 \text{ for } j \neq i.$$

whereas for $i = 0$,

$$\begin{aligned} \frac{\partial \phi_0}{\partial n_0} &= \mu_0 \mathcal{C} \frac{S_\alpha(\vec{n})}{(g_0 n_0 + S_\alpha(\vec{n}))^2} g_0 > 0, \\ \frac{\partial \phi_0}{\partial n_i} &= -\mu_0 \mathcal{C} \frac{g_0 n_0}{(g_0 n_0 + S_\alpha(\vec{n}))^2} \cdot g_i \left(\frac{S_\alpha(\vec{n})}{n_i} \right)^{1-\alpha} < 0, \text{ for } i \neq 0. \end{aligned}$$

It is readily seen that $P \equiv (p_{ij})_{i,j=0}^M$ has only positive eigenvalues. The covariance matrix turns out to be, for $i \neq 0$,

$$\Sigma_{00} = \frac{\lambda_0}{p_{00}} + \sum_{i=1}^M \frac{\lambda_i p_{0i}^2}{p_{00} p_{ii} (p_{00} + p_{ii})}, \quad \Sigma_{0i} = -\frac{\lambda_i p_{0i}}{p_{ii} (p_{00} + p_{ii})} > 0, \quad \Sigma_{ii} = \frac{\lambda_i}{p_{ii}}.$$

The other covariances are 0. We obtain the following, plausible, result.

Proposition 3.3 *In the case $\mathcal{S}^C = \{0\}$, all correlations are positive.*

4 Concluding remarks

In this paper we have considered bandwidth-sharing networks under a diffusion scaling, with a focus on the single node and linear networks. As we have seen, imposing a (natural)

peak-rate restriction, an equilibrium point of the fluid limit exists. By following a diffusion approach, we have succeeded in characterizing the covariances between the numbers of flows of the various classes.

Future research may include the following subjects.

- (1) Extension to general networks operating under α -fair bandwidth sharing. So far, we have concentrated on single nodes and linear networks, as for these the rate allocation is explicitly known. As a first step we could focus on other topologies for which this is the case (cyclic networks, for instance), or on linear networks with *unequal* service rates.
- (2) Verification of the accuracy of the diffusion approximation: for which L does the diffusion reliably predict the steady-state distribution of our Markov chain?
- (3) Verification of the approximation for networks with TCP traffic. The bandwidth allocation obtained with TCP can be approximated by an α -fair allocation with $\alpha = 2$ [Kelly, Maulloo and Tan (1998)]. Thus, it would be interesting to see how well the results of this paper match estimates obtained through simulation of the actual system (i.e., of the IP network, not of the Markov model; this could be done for instance by relying NS2).
- (4) Analyze the impact of access rates, cf. also Ben Fredj *et al.* (2001). Our theory gives a handle on assessing the impact of the access rate. As we have seen in the single-node case, for classes that are not peak-rate constrained, the number of flows present of that class is not affected by its access rate, and hence upgrading the access rate is not very beneficial. For peak-rate constrained classes such an upgrade leads to performance improvements, but may also lead to a performance degradation for the other classes. It is interesting to characterize these sensitivities.
- (5) Approximations of the distribution away from the equilibrium point of the fluid limit. Results obtained by a diffusion scaling are typically accurate around the equilibrium point of the fluid limit $L \cdot \bar{n}^*$, but may be rather inaccurate away from the equilibrium point. In this region one may come up with other (large-deviations based) approximations; alternatively, one could devise importance-sampling simulation procedures for estimating the probabilities of such rare events.

Acknowledgments

We thank Laurent Massoulié (Thomson), Frank Kelly (University of Cambridge), and two anonymous referees for their useful remarks.

References

- [1] Abendroth, D., H. van den Berg, and M. Mandjes. (2006). "A versatile model for TCP bandwidth sharing in networks with heterogeneous users." *AEÜ International Journal on Electronics and Communications*, Vol. 60, pp. 267-278.
- [2] Altman, E., T. Jimenez, and D. Kofman. (2004). "DPS queues with stationary ergodic service times and the performance of TCP in overload." *Proc. Infocom*.
- [3] Asmussen, S. (2000). *Ruin Probabilities*. World Scientific.
- [4] Ben Fredj, S., T. Bonald, A. Proutire, G. Regnié, and J. Roberts. (2001). "Statistical bandwidth sharing: A study of congestion at flow level". *Proc. Sigcomm*.
- [5] Bonald, T. and L. Massoulié. (2001). "Impact of fairness on Internet performance." *Proc. Sigmetrics/Performance*, pp. 82-91.
- [6] Bonald, T. and A. Proutière. (2003). "Insensitive bandwidth sharing in data networks." *Queueing Systems*, Vol. 44, pp. 69-100.
- [7] Bonald, T. and A. Proutière. (2004) "On stochastic bounds for monotonic processor sharing networks." *Queueing Systems*, Vol. 47, pp. 81-106.
- [8] Dai, J. (1996). "A fluid limit model criterion for instability of multiclass queueing networks." *Annals of Applied Probability*, Vol. 6, pp. 751-757.
- [9] Fayolle, G., I. Mitrani, and R. Iasnogorodski. (1980). "Sharing a processor among many job classes." *Journal of the ACM*, Vol. 27, pp. 519-532.
- [10] Gibbens, R., S. Sargood, C. van Eijl, F. Kelly, H. Azmoodeh, R. Macfadyen, and N. Macfadyen. (2000). "Fixed-point models for the end-to-end performance analysis of IP networks." *Proc. 13th ITC Specialist Seminar: IP Traffic Measurement, Modeling and Management*.
- [11] Horn, R. and C. Johnson. (1985). *Matrix Analysis*. Cambridge University Press.
- [12] Hunt, P. and T. Kurtz. (1991). "Large loss networks." *Stochastic Processes and their Applications*, Vol. 53, pp. 363-378, 1991.
- [13] Kelly, F., A. Maulloo, and D. Tan. (1998). "Rate control in communication networks: shadow prices, proportional fairness and stability." *Journal of the Operational Research Society*, Vol. 49, pp. 237-252.

- [14] Karatzas, I. and S.E. Shreve. (1991). *Brownian Motion and Stochastic Calculus*. Springer-Verlag, New York.
- [15] Kelly, F. and R. Williams. (2004). "Fluid model for a network operating under a fair bandwidth-sharing policy." *Annals of Applied Probability*, Vol. 14, pp. 1055-1083.
- [16] Key, P. and L. Massoulié. (2003). "Probing strategies for distributed admission control in large and small scale systems." *Proc. Infocom*.
- [17] Khalil, H. (2001). *Nonlinear Systems*. Prentice Hall.
- [18] Kleinrock, L. (1967). "Time-shared systems: a theoretical treatment." *Journal of the ACM*, Vol. 14, pp. 242-261.
- [19] Kumar, S. and L. Massoulié. (2005). "Integrating streaming and file transfer internet traffic: Fluid and diffusion approximations." Available as *Microsoft Research Report MSR-TR-2005-160*.
- [20] Leskelä, L. (2006). "Stabilization of an overloaded queueing network using measurement based admission control." *Journal of Applied Probability*, Vol. 43, pp. 231-244.
- [21] Lieshout, P., S. Borst, and M. Mandjes. (2006). "Heavy-traffic approximations for linear networks operating under alpha-fair bandwidth-sharing policies." *Proc. ValueTools*, 2006.
- [22] Massoulié, L. and P. Key. (2006). "Fluid models of integrated traffic and multipath routing." *Queueing Systems*, Vol. 53, pp. 85-98.
- [23] Massoulié, L. and J. Roberts. (2002). "Bandwidth sharing: objectives and algorithms." *IEEE/ACM Transactions on Networking*, Vol. 10, pp. 320-328.
- [24] Mo, J. and J. Walrand. (2000). "Fair end-to-end window-based congestion control." *IEEE/ACM Transactions on Networking*, Vol. 8, pp. 556-567.
- [25] Meyn, S.P. and R.L. Tweedie. (1993). *Markov chains and stochastic stability*. Springer-Verlag, London.