Bandwidth-Sharing in Overloaded Networks[1]

Regina Egorova[*†], Sem Borst[*†‡], Bert Zwart[§]
E-mail: egorova@cwi.nl, sem@win.tue.nl, bertzwart@gatech.edu

[*]CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
[†]Dept. of Mathematics & Computer Science, Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands
[‡]Bell Labs, Alcatel-Lucent, P.O. Box 636, Murray Hill, NJ 07974, USA
[§] Georgia Institute of Technology, H. Milton Stewart School of ISyE, 765 Ferst Drive, Atlanta, GA 30332-0205, USA

***Abstract*— Bandwidth-sharing networks as considered by Massoulié & Roberts provide a natural modeling framework for describing the dynamic flow-level interaction among elastic data transfers. Under mild assumptions, it has been established that a wide family of so-called α-fair bandwidth-sharing strategies achieve stability in such networks provided that no individual link is overloaded.**

In the present paper we focus on α-fair bandwidth-sharing networks where the load on one or several of the links exceeds the capacity. Evidently, a well-engineered network should not experience overload, or even approach overload, in normal operating conditions. Yet, even in an adequately provisioned system with a low nominal load, the actual traffic volume may significantly fluctuate over time and exhibit temporary surges. Furthermore, gaining insight in the overload behavior is crucial in analyzing the performance in terms of long delays or low throughputs as caused by large queue build-ups. The way in which such rare events tend to occur, commonly involves a scenario where the system temporarily behaves as if it experiences overload.

In order to characterize the overload behavior, we examine the fluid limit, which emerges from a suitably scaled version of the number of flows of the various classes. Focusing on linear solutions to the fluid-limit equation, we derive a fixed-point equation for the corresponding asymptotic growth rates. It is proved that a fixed-point solution is also a solution to a related strictly concave optimization problem, and hence exists and is unique. The results are illustrated for linear topologies and star networks as two important special cases.

I. 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 [23], [26] provide a natural extension for modeling the dynamic interaction among competing elastic flows that traverse several links along their source-destination paths.

Assuming exponential flow size distributions and Poisson arrivals, De Veciana *et al.* [27], [28] proved that weighted max-min and proportional fair bandwidth-sharing strategies achieve stability in such networks (positive recurrence of the associated Markov process) under the nominal condition that no individual link is overloaded. Bonald & Massoulié [5] extended that result to a wide family of weighted α-fair bandwidth-sharing strategies as introduced by Mo & Walrand [24]. Massoulié [22] established that the nominal stability condition remains sufficient for the proportional fair strategy with an additional 'routing feature', thus further generalizing the result to phase-type flow size distributions. Bramson [9] showed that the max-min fair strategy guarantees stability under the nominal load condition for general flow size distributions and renewal arrival processes. Under similar general distributional assumptions and load conditions, Gromoll & Williams [15], [16] studied the fluid limit for weighted α-fair strategies, and established stability in some special cases, such as linear and tree topologies. Interesting stability results for α-fair strategies and general flow size distributions with bounded support were obtained by Chiang *et al.* [10]. Hansen *et al.* [17] examined the impact of rate allocation policies on the stability conditions for exponential flow sizes from the perspective of entrainment. The latter term is used by Kelly & Williams [20] to refer to the phenomenon that the simultaneous resource requirements may cause congested links to prevent other links from utilizing their full capacities.

In the present paper we focus on α-fair bandwidth-sharing networks where the load on one or several of the links exceeds the capacity. Obviously, with adequate provisioning, a network should not experience overload, or even approach overload, in normal operating conditions. However, even in a properly dimensioned system with a low typical load, the actual traffic volume may substantially fluctuate over time and exhibit transient surges, see also Bonald & Roberts [6]. Furthermore, an understanding of the overload behavior plays a crucial role in analyzing the performance in terms of long transfer delays or low flow throughputs as caused by large queue build-ups. The likely way for such rare events to occur, commonly entails a scenario where the system temporarily appears to deviate from the normal stochastic laws and behaves as if it experiences overload, see for instance Anantharam [2].

As alluded to above, an α-fair bandwidth-sharing network bears strong resemblance with a single-server processor-sharing system, since within each class the bandwidth is fairly shared among all competing flows. The overload behavior of a single-server processor-sharing system was first analyzed by Jean-Marie & Robert [19]. Their analysis was extended by Altman *et al.* [1] to the *discriminatory* processor-

[1]The work of the first author is financially supported by The Netherlands Organization for Scientific Research (NWO).

sharing discipline, which corresponds to a single-node network with a *weighted* α-fair strategy. Puha *et al.* [25] studied a single-server overloaded processor-sharing system in terms of measure-valued processes.

There are two key distinctions that arise in a network scenario: (i) the rate received by a class is no longer constant, but depends on the number of flows of all classes in some intricate fashion; and (ii) the network may show non-work-conserving behavior due to the entrainment phenomenon mentioned above. These two features not only render the flow-level performance largely intractable, but also complicate the analysis of the overload behavior. For example, even on links with excess capacity, the workloads may grow because of the non-work-conserving behavior mentioned above. In addition, while the total number of flows must grow in overload conditions, the exact nature of the growth patterns of the various classes is far from clear, and may even potentially involve oscillatory effects in certain cases as observed in Bramson [8] and Lu & Kumar [21] for example.

In order to characterize the growth dynamics, we examine the fluid limit, which emerges when the number of flows of the various classes is scaled in both space and time. Focusing on linear solutions to the fluid-limit equation, we obtain a fixed-point equation for the coefficients, which represent the corresponding asymptotic growth rates of the queue lengths. It is proved that a fixed-point solution is also a solution to a related strictly concave optimization problem, and therefore exists and is unique.

The results are illustrated for linear topologies and star networks as two important special cases. While admittedly simple, linear networks provide a useful model for flows that traverse several links and experience bandwidth competition from independent cross-traffic. Star networks offer a convenient abstraction for scenarios where the core is highly over-provisioned and congestion predominantly occurs at the edge with comparatively low-capacity access links, see also Fayolle *et al.* [12].

The remainder of the paper is organized as follows. In Section II we present a detailed model description and state some preliminaries. In Section III we examine linear solutions to the fluid-limit equation, derive a fixed-point equation for the corresponding asymptotic growth rates, and establish existence and uniqueness of the fixed point. We focus on the special case of a network with a linear topology in Section IV. In Section V we turn attention to the special case of a star network.

II. Model description and preliminaries

In this section we present a detailed model description and state some preliminaries.

Flow-level model

We consider a bandwidth-sharing network as in [23], [26] with a finite number of links labeled by $j = 1, \dots, J$. Denote by $C = (C_1, \dots, C_J)$ the vector of link capacities. The network is offered traffic from several classes indexed by $i = 1, \dots, I$. Each class is characterized by a route, i.e., a nonempty subset of $\{1, \dots, J\}$, which represents the links traversed by the traffic from that class. Let A be a $J \times I$ incidence matrix such that $A_{ji} = 1$ if link j belongs to the route of class i, and $A_{ji} = 0$ otherwise. For now, we do not make any specific assumptions on the topology of the network or the structure of the route sets.

The traffic of the various classes consists of elastic file transfers. Class-i flows arrive as a renewal process with rate λ_i, i.e., the mean interarrival time is $1/\lambda_i$, and have generally distributed sizes B_i with mean $1/\mu_i$. Denote by $\rho = (\rho_1, \dots, \rho_I)$ with $\rho_i := \lambda_i/\mu_i$ the vector of traffic intensities.

Bandwidth-sharing strategy

Denote by $\Lambda = (\Lambda_1, \dots, \Lambda_I)$ the vector of rates allocated to the various classes. Any rate allocation vector Λ must satisfy the capacity constraints $A\Lambda \leq C$. The bandwidth arbitration among competing flows is governed by a weighted α-fair strategy [24], which selects a rate allocation vector $\Lambda(z) = (\Lambda_1(z), \dots, \Lambda_I(z))$ based on the population $z = (z_1, \dots, z_I)$ of active flows and an optimization criterion as specified below. Within each class, the rate is fairly shared among all competing flows, i.e., if $z_i > 0$, then each class-i flow receives service at rate $\Lambda_i(z)/z_i$. Thus, a class-i flow that is continuously active throughout the time interval $[s, t]$, receives a cumulative amount of service

$$S_i(s,t) = \int_s^t \frac{\Lambda_i(Z(u))}{Z_i(u)} du,$$

with $Z(t) = (Z_1(t), \dots, Z_I(t))$ representing the population of active flows at time t.

A weighted α-fair strategy is parameterized by a fairness coefficient $\alpha \in (0, \infty)$ and a weight vector $w = (w_1, \dots, w_I) \in \mathbf{R}_I^+$. For a given population $z = (z_1, \dots, z_I) \neq (0, \dots, 0)$ of active flows, the weighted α-fair rate allocation $\Lambda(z)$ is determined by the solution to the optimization problem:

$$(P) \qquad \begin{array}{ll} \text{maximize} & G_z(\Lambda) \\ \text{subject to} & A\Lambda \leq C,\, \Lambda \geq 0, \end{array}$$

where the objective function $G_z(\cdot) : \mathbf{R}_+^I \to [-\infty, \infty]$ is defined by

$$G_z(\Lambda) = \begin{cases} \sum_{i=1}^I w_i\, z_i^\alpha\, \frac{\Lambda_i^{1-\alpha}}{1-\alpha}, & \alpha \in (0,\infty)/\{1\}, \\ \sum_{i=1}^I w_i\, z_i\, \log \Lambda_i, & \alpha = 1, \end{cases} \tag{1}$$

with the convention that $G_z(\Lambda) = -\infty$ if $\alpha \geq 1$ and $\Lambda_i = 0$, $z_i > 0$. With the additional convention that $\Lambda_i(z) = 0$ when $z_i = 0$, the rate allocation is uniquely determined since the above optimization problem is strictly concave.

The family of α-fair bandwidth-sharing strategies includes several common fairness concepts as special cases. In particular, the case $\alpha = 1$ and the limiting cases $\alpha \to 0$ and $\alpha \to \infty$ correspond to a rate allocation that is *proportional fair*, achieves *maximum throughput*, and is *max-min fair*, respectively.

Load conditions

Under Markovian assumptions, [5] established that α-fair strategies achieve stability provided no individual link is overloaded, i.e., $A\rho < C$. In the present paper we focus on an overload scenario where the above condition is violated for at least one of the links.

III. Main results

In this section we present our main results, which characterize the growth rates of the number of flows in an overloaded α-fair bandwidth-sharing network. Specifically, this section examines the fluid limit emerging from a suitably scaled version of the number of flows of the various classes, namely, the sequence of processes $(Z(rt)/r, t \geq 0)$ with $r \to \infty$. Gromoll & Williams [15] established that the sequence of these scaled processes is tight.

In preparation for the statement of the main result, we first introduce a slightly modified version of the rate allocation functions which may be interpreted as the service rates at the fluid scale. The service rate $R_i(z)$ received by class i is defined as follows: $R_i(z) \equiv \Lambda_i(z)$ if $z_i > 0$, where $\Lambda_i(\cdot)$ is the solution of the optimization problem (P), and $R_i(z) \equiv \rho_i$ if $z_i = 0$. The above distinction reflects the fact that on the fluid scale, $z_i = 0$ requires that class i receives service at rate ρ_i, rather than 0.

We now proceed to postulate a fluid-model solution $(z(t), t \geq 0)$. This definition is slightly different from the one in [15]: a nonnegative continuous function $(z(t), t \geq 0)$ is said to be a *fluid-model solution* if it satisfies

$$z_i(t) = \lambda_i \int_0^t \mathbf{P}(B_i > S_i(s,t))ds, \qquad (2)$$

where $S_i(s,t) = \int_s^t \frac{R_i(z(u))}{z_i(u)}du$ is the cumulative amount of service received by a class-i flow during the time interval $[s,t]$. As a matter of fact, we expect our definition to be equivalent to the one in [15] (extending the arguments in the proof of Lemma 4.3 in Gromoll *et al.* [14] from the single-server case to a network scenario).

The next proposition presents the main result of the paper.

Proposition 3.1: The fluid-limit Equation (2) admits a linear solution $z(t) \equiv mt$, where the vector $m = (m_1, \ldots, m_I)$ forms the unique solution to the fixed-point equation

$$R_i(m) = \rho_i \mathbf{E}\left[\mathrm{e}^{-\frac{m_i}{R_i(m)}B_i^*}\right], \qquad i = 1, \ldots, I, \qquad (3)$$

and B_i^* represents a residual flow size (i.e. a random variable with density $\mathbf{P}(B_i > x)/\mathbf{E}[B_i]$ and LST $\mathbf{E}\left[\mathrm{e}^{-xB_i^*}\right] = \mu_i(1 - \mathbf{E}\left[\mathrm{e}^{-xB_i}\right])/x)$.

The above proposition holds for arbitrary network topologies and arbitrary flow size distributions. In a single-link scenario, i.e., $J = 1$, it reduces to known results for single-server processor-sharing type systems. In particular, in the single-class case, i.e., $I = 1$, we have, dropping the class index, $R(m) = 1$, and Equation (3) specializes to

$$1 = \rho\, \mathbf{E}[\mathrm{e}^{-mB^*}],$$

which corresponds to the result in [19]. In the multi-class case, we have $R_i(m) = w_i m_i / \sum_{k=1}^I w_k m_k$, and Equation (3) takes the form

$$\frac{w_i m_i}{\sum_{k=1}^I w_k m_k} = \rho_i\, \mathbf{E}\left[\mathrm{e}^{-w_i^{-1}\sum_{k=1}^I w_k m_k B_i^*}\right], \qquad i = 1, \ldots, I,$$

which agrees with the result in [1] for overloaded discriminatory processor-sharing queues. We will give additional examples for specific network topologies in Sections IV and V.

A. Heuristic interpretation

The fixed-point Equation (3) may be heuristically derived in a similar way as explained by Jean-Marie [18]; we are not aware of an article where this derivation has been published.

Suppose that $\frac{Z(r)}{r} \to m$ a.s. as $r \to \infty$ for some vector $m = (m_1, \ldots, m_I)$. Then, for large t, a class-i flow will receive service at a rate of approximately $\frac{R_i(mt)}{m_i t}$. Let a_i^n be the arrival epoch of the n-th class-i flow. Then the size B_i of that flow and its sojourn time T_i may be related as:

$$B_i = \int_{a_i^n}^{a_i^n + T_i} \frac{R_i(mu)}{m_i u} du.$$

Since $R_i(mu) = R_i(m)$ (see [20]), it follows that

$$\frac{m_i}{R_i(m)} B_i = \int_{a_i^n}^{a_i^n + T_i} \frac{1}{u} du = \log(a_i^n + T_i) - \log a_i^n.$$

Taking the exponent on both sides, we obtain

$$T_i = a_i^n \left(\mathrm{e}^{\frac{m_i}{R_i(m)}B_i} - 1\right).$$

The number of active class-i flows at time t may then be expressed as $Z_i(t) = \#\{n : a_i^n + T_i^n \geq t, a_i^n \leq t\} = \#\{n : t \geq a_i^n \geq t\mathrm{e}^{-\frac{m_i}{R_i(m)}B_i}\}$. Because $a_i^n \approx n/\lambda_i$ for large n, we have

$$Z_i(t) = \#\{n : t \geq \frac{n}{\lambda_i} \geq t\mathrm{e}^{-\frac{m_i}{R_i(m)}B_i}\} \approx \lambda_i t \left(\mathbf{E}\left[1 - \mathrm{e}^{-\frac{m_i}{R_i(m)}B_i}\right]\right).$$

Dividing both sides by t and letting t tend to infinity, we deduce

$$m_i = \lambda_i \left(1 - \mathbf{E}\left[\mathrm{e}^{-\frac{m_i}{R_i(m)}B_i}\right]\right), \qquad (4)$$

which is equivalent to Equation (3).

Remark 3.1: In the case of exponential flow sizes, Equation (3) specializes to

$$m_i = \lambda_i - \mu_i R_i(m), \qquad i = 1, \ldots, I, \qquad (5)$$

which makes sense, since $\mu_i R_i(m)$ is indeed the departure rate of class-i flows. This is also consistent with the convention $R_i(m) = \rho_i$ when $m_i = 0$.

Remark 3.2: Note that we have *not* established convergence of $(Z(rt)/r, t \geq 0)$ to $(z(t), t \geq 0)$. In order to prove convergence, one would need to show that (i) any convergent subsequence of $(Z(rt)/r, t \geq 0)$ is a fluid-model solution in our sense; and (ii) that our fluid-model solution is unique. In [11] a proof is sketched of (i). To establish (ii), it remains to be shown that any solution to the fluid-limit equation is linear. While we strongly conjecture that to be the case for any finite initial state $Z(0)$, a rigorous proof appears quite challenging in general. In the special case of exponential flow sizes, any differentiable fluid-limit solution satisfies $z_i'(t) = \lambda_i - \mu_i R_i(z(t))$. Using a similar proof technique as in [5], it may then be shown that in case of the proportional fair policy the function $F(y(t))$, with $F(u) := \sum_{i=1}^{I} w_i u_i^2/\mu_i R_i(m)$ and $y(t) = z(t) - mt$, is non-increasing as function of t. In case $z(0) = 0$, so that $F(y(0)) = 0$, it then follows that $F(y(t)) = 0$ for all $t \geq 0$, i.e., $z(t) = mt$ for all $t \geq 0$. Thus, any differentiable fluid-limit solution is linear.

B. Proof of Proposition 3.1

The proof of Proposition 3.1 follows from the next two lemmas.

Lemma 3.1: At time t, for any class i, a solution of Equation (2) is given by

$$z_i(t) = m_i t,$$

where m_i is a solution of Equation (4).

Proof Suppose $z_i(t) = m_i t$, where m_i is some constant. Substituting this into Equation (2), we obtain that

$$\begin{aligned} z_i(t) = m_i t &= \lambda_i \int_0^t \mathbf{P}\left(B_i > \int_s^t \frac{R_i(m_i u)}{m_i u} du\right) ds \\ &= \lambda_i \int_0^t \mathbf{P}\left(B_i > \int_s^t \frac{R_i(m)}{m_i u} du\right) ds \\ &= \lambda_i \int_0^t \mathbf{P}\left(-B_i \frac{m_i}{R_i(m)} < \log \frac{s}{t}\right) ds \end{aligned}$$

$$= \lambda_i t \int_0^1 \mathbf{P}\left(\mathrm{e}^{-B_i \frac{m_i}{R_i(m)}} < u\right) du = \lambda_i t (1 - \mathbf{E}\left[\mathrm{e}^{-B_i \frac{m_i}{R_i(m)}}\right]),$$

which yields that m_i is a solution of Equation (4). $\square$

Lemma 3.2: Equation (3) has a unique solution $m = (m_1, \ldots, m_I)$.

Proof We will establish uniqueness of $R(m)$. The monotonicity of the LST's $\beta_i(\cdot)$ then implies uniqueness of $m = (m_1, \ldots, m_I)$.

A crucial role is played by a related optimization problem. To formulate this optimization problem, we rewrite the fixed-point Equation (3) in the equivalent form

$$\frac{m_i}{R_i(m)} = \beta_i^{-1}\left(\frac{R_i(m)}{\rho_i}\right), \qquad i = 1, \ldots, I, \tag{6}$$

where $\beta_i^{-1}(\cdot)$ is the inverse of the Laplace-Stieltjes Transform (LST) $\beta_i(y) = \mathbf{E}\left[\mathrm{e}^{-yB_i^*}\right]$, $\beta_i^{-1}(\beta_i(y)) = y$. Observe that the right-hand side only depends on m through $R_i(m)$. This motivates us to introduce the function $H_i(x)$, which has derivative $H_i'(x) = \left(\beta_i^{-1}\left(\frac{x}{\rho_i}\right)\right)^{\alpha}$. Since $\beta_i(x)$ is strictly decreasing in x, its inverse is strictly decreasing in x as well. Consequently, $H_i(x)$ is strictly concave in x.

Now, consider the optimization problem

$$(Q) \qquad \begin{array}{ll} \text{maximize} & H(R) = \sum_{i=1}^{I} w_i H_i(R_i) \\ \text{subject to} & AR \leq C,\ R \geq 0. \end{array}$$

This optimization problem is strictly concave, and hence has a unique solution $R = (R_1, \ldots, R_I)$ (see for instance [7]).

By using the fact that $R_i(m) \equiv \Lambda_i(m)$ when $m_i > 0$ and $R_i(m) \equiv \rho_i$ when $m_i = 0$ and observing that the rate allocation vector $\Lambda(z)$ satisfies the necessary Karush-Kuhn-Tucker (KKT) conditions [3] for problem (P), it may be shown that the vector $R(m)$ satisfying Equation (6) obeys the sufficient KKT conditions for problem (Q), and hence is a global optimum. We refer to [11] for details. $\square$

Besides proving uniqueness, the above equivalence also provides a way for actually computing the asymptotic growth rates m_i by solving the concave programming problem (Q).

IV. Linear networks

In this section we focus on the special case of linear networks The network consists of links $1, \ldots, L$, each of unit capacity, and is offered traffic from classes $0, 1, \ldots, L$. Class-j flows require service from link j only, $j = 1, \ldots, L$, while class-0 flows demand capacity on all links simultaneously. As mentioned earlier, linear networks provide a useful model for traffic that traverses several links and experiences bandwidth competition from independent cross-traffic. We assume that the load on at least one of the links exceeds the capacity, i.e., $\max_{j=1,\ldots,L} \rho_j > 1 - \rho_0$. The bandwidth arbitration is governed by the proportional fair policy with unit class weights, i.e., the objective function $G_z(\Lambda)$ is given by $\sum_{i=0}^{L} z_i \log(\Lambda_i)$. The capacity constraints take the form $\Lambda_0 + \Lambda_j \leq 1$ for all $j = 1, \ldots, L$.

We are interested in determining the asymptotic growth rates m_i of the various classes. According to Proposition 3.1, there exist nonnegative coefficients (Lagrange multipliers) p_j associated with the various links so that the m_i and the corresponding rate allocations R_i together with the p_j form a solution to the system of equations

$$\begin{cases} m_0 = R_0 \sum_{j=1}^{L} p_j, \\ m_j = R_j p_j, \qquad j = 1, \ldots, L, \\ p_j(R_0 + R_j - 1) = 0, \qquad j = 1, \ldots, L, \end{cases} \tag{7}$$

in conjunction with the set of fixed-point Equations (6). The latter equations in fact allow us to express the m_j's in terms of the R_j's, yielding

$$\begin{cases} \beta_0^{-1}\left(\frac{R_0}{\rho_0}\right) = \sum_{j=1}^{L} p_j, \\ \beta_j^{-1}\left(\frac{R_j}{\rho_j}\right) = p_j, \qquad j = 1, \ldots, L, \\ p_j(R_0 + R_j - 1) = 0, \qquad j = 1, \ldots, L. \end{cases} \tag{8}$$

In total the above system provides $2L+1$ equations for $2L+1$ unknown variables R_i, $i = 0, \ldots, L$, and p_j, $j = 1, \ldots, L$.

In order to solve the above system of equations, we consider the nonempty subset $\mathcal{J}_+ := \{j : p_j > 0\}$ of links with strictly positive Lagrange multipliers. (The subset $\mathcal{J}_+$ cannot be empty, since that would imply $\rho_0 + \rho_j \leq 1$ for all $j = 1, \ldots, J$, and contradict the overload assumption.) Observe that $p_j = 0$ means $m_j = 0$, and the growth rates of classes 0 and $j \in \mathcal{J}_+$ thus correspond to those in a scenario with classes $j \notin \mathcal{J}_+$ as well as links $j \notin \mathcal{J}_+$ removed. In particular, when $\mathcal{J}_+ = \{j_+\}$, the growth rates of classes 0 and j_+ are identical to those in a single-node processor-sharing system with classes 0 and j_+.

For compactness, denote $n = \sum_{j \in \mathcal{J}_+} p_j$. Then the solution to the system of Equations (7) may be represented as

$$R_0 = \frac{m_0}{n}; \;\; R_j \equiv S_{\mathcal{J}_+} = 1 - \frac{m_0}{n} \;\text{ if } j \in \mathcal{J}_+; \;\; R_j = \rho_j \;\text{ if } j \notin \mathcal{J}_+;$$

$$p_j = \frac{m_j}{R_j}, \; j = 1, \ldots, J. \tag{9}$$

Summing the last equality in Equation (9) over $j \in \mathcal{J}_+$, it follows that $n = m_0 + \sum_{j \in \mathcal{J}_+} m_j$.

What remains is to determine the subset $\mathcal{J}_+$ in terms of the system parameters. Note that $j \in \mathcal{J}_+$ implies $R_0 + R_j = 1$, and thus necessitates $\rho_0 + \rho_j \geq 1$. However, the latter inequality is not sufficient for $j \in \mathcal{J}_+$, since it is possible that $m_j = 0$ when other classes at other links sufficiently throttle the service rate of class 0. In order to characterize the subset $\mathcal{J}_+$, observe that $\rho_j \leq S_{\mathcal{J}_+}$ for all $j \notin \mathcal{J}_+$ and $\rho_j > S_{\mathcal{J}_+}$ for all $j \in \mathcal{J}_+$. In view of the inherent symmetry, we may assume without loss of generality that the links are indexed such that $\rho_1 \leq \rho_2 \leq \cdots \leq \rho_L$. Denote by σ_j the common service rate obtained by classes $j, \ldots, L$ in a system with links $j, \ldots, L$ and classes 0 and $j, \ldots, L$ only, $\sigma_j = 1 - \Lambda_0(m_j, \ldots, m_L) = 1 - \frac{m_0}{m_0 + \sum_{k=j}^{L} m_k}$. Then the subset $\mathcal{J}_+$ is of the form $\{j_+, \ldots, L\}$, with $j_+ := \max\{j : \rho_{j-1} \leq \sigma_j\}$. In case $B_0 \equiv B_L$, it is easily verified that $\sigma_L = \lambda_L / (\lambda_0 + \lambda_L)$.

The above characterization of the subset $\mathcal{J}_+$ may be interpreted as follows. If $\rho_{j-1} \leq \sigma_j$, then competition from classes $j, \ldots, L$ alone against class 0 is sufficient to throttle the rate of class 0 to an extent that what remains available for classes $1, \ldots, j-1$ exceeds their respective loads, and hence $m_1 = \cdots = m_{j-1} = 0$. This scenario occurs when the loads of classes $1, \ldots, j-1$ are relatively low and the loads of classes $j, \ldots, L$ are sufficiently high. Note that this may occur even when $\rho_0 + \rho_i > 1$ for some classes $i = 1, \ldots, j-1$. Although these classes rely on service from overloaded links, they remain stable thanks to the much stronger competition at other higher-loaded links. In contrast, if $\rho_{j-1} > \sigma_j$, then competition from classes $j, \ldots, L$ alone is not sufficient to provide stability to class $j-1$, and hence $m_{j-1} > 0$.

Exponential flow sizes

In order to determine the growth rates explicitly, we need to specify the flow size distributions of the various classes that occur in the set of fixed-point Equations (6). In the case of exponential flow sizes, substituting (6) into (9), the growth rates of the various classes may be represented in terms of the single variable n as

$$m_0 = \lambda_0 - \mu_0 \frac{m_0}{n} = \frac{\lambda_0 n}{\mu_0 + n}, \tag{10}$$

$$m_j = \lambda_j - \mu_j \frac{n - m_0}{n} = \lambda_j - \mu_j \left(1 - \frac{\lambda_0}{\mu_0 + n}\right), \qquad j \in \mathcal{J}_+. \tag{11}$$

Summing the above equations yields the quadratic equation $n^2 + \nu n + \kappa = 0$, with

$$\nu := \mu_0 + \sum_{j \in \mathcal{J}_+} \mu_j - \lambda_0 - \sum_{j \in \mathcal{J}_+} \lambda_j,$$

$$\kappa := \sum_{j \in \mathcal{J}_+} \mu_j (\mu_0 - \lambda_0) - \mu_0 \sum_{j \in \mathcal{J}_+} \lambda_j.$$

Substituting the positive solution in Equations (10)–(11) gives expressions for the asymptotic growth rates. To see that there is indeed a unique positive solution, recall that a quadratic equation has a unique positive solution when the zero-order constant is nonpositive, which may be written as

$$\rho_0 + \sum_{j \in \mathcal{J}_+} \rho_j \frac{\mu_j}{\sum_{j \in \mathcal{J}_+} \mu_j} \geq 1. \tag{12}$$

Noting that

$$\sum_{j \in \mathcal{J}_+} \rho_j \frac{\mu_j}{\sum_{j \in \mathcal{J}_+} \mu_j} \geq \sum_{j \in \mathcal{J}_+} \min_{k \in \mathcal{J}_+} \rho_k \frac{\mu_j}{\sum_{j \in \mathcal{J}_+} \mu_j} = \min_{k \in \mathcal{J}_+} \rho_k,$$

the inequality (12) is seen to hold by virtue of the fact that $\rho_0 + \rho_j \geq 1$, $j \in \mathcal{J}_+$.

V. Star networks

We now turn attention to the special case of a star network. As mentioned earlier, star networks offer a convenient abstraction for scenarios where the core is highly over-provisioned and congestion predominantly occurs at the edge with comparatively low-capacity access links. The network is composed of L links, each of unit capacity, and is offered traffic from $L(L-1)/2$ classes labelled as $\{i, j\}$, $i, j = 1, \ldots, L$, $i \neq j$. The route of class $\{i, j\}$ simply consists of the two links i and j. We assume that the load on at least one of the links exceeds the capacity, i.e., $\max_{j=1,\ldots,L} \sigma_j > 1$, with $\sigma_j := \sum_{k \neq j} \rho_{\{j,k\}}$. The bandwidth arbitration is governed by the proportional fair policy with unit class weights, i.e., the objective function $G_z(\Lambda)$ is given by $\sum_{j \neq k} z_{\{j,k\}} \log(\Lambda_{\{j,k\}})$. The capacity constraints take the form $\sum_{k \neq j} \Lambda_{\{j,k\}} \leq 1$ for all $j = 1, \ldots, L$.

For star networks, Proposition 3.1 and Equation (6) imply that the Lagrange multipliers p_j associated with the links in the network and the corresponding rate allocations $R_{\{j,k\}}$ satisfy the following system of equations

$$\begin{cases} \beta_{\{j,k\}}^{-1}\left(\frac{R_{\{j,k\}}}{\rho_{\{j,k\}}}\right) = R_{\{j,k\}}(p_j + p_k), & j \neq k, \\ p_j\left(\sum_{k \neq j} R_{\{j,k\}} - 1\right) = 0, & j = 1, \ldots, L. \end{cases} \tag{13}$$

In total the above system provides $L(L+1)/2$ equations for $L(L+1)/2$ unknown variables $R_{\{j,k\}}$, $j \neq k$, and p_j, $j = 1, \ldots, L$. In the case of exponential flow sizes, the set of fixed-point equations takes the explicit form in Equation (5). The above system of equations then simplifies to

$$\begin{cases} \lambda_{\{j,k\}} - \mu_{\{j,k\}} R_{\{j,k\}} = R_{\{j,k\}}(p_j + p_k), & j \neq k, \\ p_j(\sum_{k \neq j} R_{\{j,k\}} - 1) = 0, & j = 1, \ldots, L. \end{cases} \tag{14}$$

As before, we need to consider the subset of links with strictly positive Lagrange multipliers in order to solve the above system of equations.

Three-link network

As an illustrative example, we now focus on the case of a triangular network with three links and three classes. In that case we need to distinguish three scenarios, (I), (II) and (III), depending on whether one, two or all three of the Lagrange multipliers are strictly positive, respectively. It cannot occur that all three Lagrange multipliers are zero, since that would imply $\sum_{k \neq j} \rho_{\{j,k\}} < 1$, $j = 1, 2, 3$, and contradict the overload assumption.

With minor abuse of notation, we define $m_i := m_{\{1,2,3\} \setminus \{i\}}$, $R_i := R_{\{1,2,3\} \setminus \{i\}}$, and $\rho_i := \rho_{\{1,2,3\} \setminus \{i\}}$. The above system of Equations (13) may then be rewritten as

$$\begin{cases} m_i = R_i(p_j + p_k), & \{i,j,k\} = \{1,2,3\}, \\ p_j(R_i + R_k - 1) = 0, & \{i,j,k\} = \{1,2,3\}. \end{cases} \tag{15}$$

For compactness, denote $m := m_1 + m_2 + m_3$, and $n = m_i + m_j$. The solutions in the above three scenarios may be then represented as

$$\begin{aligned} (I) \qquad & (R_i, R_j, R_k) = \left(\tfrac{m_i}{n}, \tfrac{m_j}{n}, \rho_k\right), \\ & p_i = p_j = 0, \quad p_k = n, \\ (II) \qquad & (R_i, R_j, R_k) = \left(\tfrac{m_i}{m}, \tfrac{m_j + m_k}{m}, \tfrac{m_j + m_k}{m}\right), \\ & p_i = 0, \quad p_j = \tfrac{m_j}{m_j + m_k} m, \quad p_k = \tfrac{m_k}{m_j + m_k} m, \\ (III) \qquad & R_1 = R_2 = R_3 = \tfrac{1}{2}, \\ & p_i = \textstyle\sum_{j \neq i} m_j - m_i > 0, \quad i = 1, 2, 3. \end{aligned}$$

The above results reveal an interesting trichotomy in the behavior of the triangular network. In case (I) the network behaves as a single-node processor-sharing system with classes i and j only. The conditions for case (I) to occur in terms of the system parameters also coincide with the corresponding ones in the linear network with $|\mathcal{J}_+| = 1$. Case (II) corresponds to the case of the linear network with $|\mathcal{J}_+| = 2$. The conditions for this case to arise subsume the corresponding ones in the linear network, but include an additional condition that the loads of the three classes should be slightly unbalanced. If the latter condition is violated, i.e., the loads of the three classes are nearly equal, then case (III) arises, which has no counterpart in the linear network. In this case, each of the three classes behaves as in an isolated processor-sharing system with capacity $\frac{1}{2}$.

References

[1] Altman, E., Jimenez, T., Kofman, D. (2004). DPS queues with stationary ergodic service times and the performance of TCP in overload. In: *Proc. IEEE Infocom 2004*, Hong-Kong.

[2] Anantharam, V. (1989). How large delays build up in a GI/G/1 queue. *Queueing Systems* **5**, 345–367

[3] Bazaraa, M.S., Sherali, H.D., Shetty, C.M. (1993). *Nonlinear Programming: Theory and Algorithms.* John Wiley, New Jersey.

[4] Ben-Tal, A., Nemirovski, A. (2001). *Lectures on Modern Convex Optimization.* MPS-SIAM Series on Optimization, SIAM, Philadelphia.

[5] Bonald, T., Massoulié, L. (2001). Impact of fairness on Internet performance. In: *Proc. ACM Sigmetrics & Performance 2001 Conf.*, Boston MA, USA, 82–91.

[6] Bonald, T., Roberts, J.W. (2001). Performance modeling of elastic traffic in overload. In: *Proc. ACM Sigmetrics & Performance 2001 Conf.*, Boston MA, USA, 342–343.

[7] Boyd, S., Vanderberghe, L. (2004). *Convex Optimization.* University Press, Cambridge.

[8] Bramson, M. (1994). Instability of FIFO queueing networks. *Ann. Appl. Prob.* **4**, 414–431.

[9] Bramson, M. (2005). Stability of networks for max-min fair routing. Presentation at the 13th INFORMS Applied Probability Conference, Ottawa.

[10] Chiang, M., Shah, D., Tang, A. (2006). Stochastic stability of network utility maximization: general file size distribution. Submitted for publication in *IEEE Trans. Inf. Theory*.

[11] Egorova, R., Borst, S.C. Zwart, A.P. (2007). Bandwidth-sharing networks in overload. *Perf. Eval* **64** (Proc. Performance 2007 Conf.), 978–993.

[12] Fayolle, G., De la Fortelle, A., Lasgouttes, J.-M., Massoulié, L., Roberts, J.W. (2001). Best-effort networks: modeling and performance analysis via large network asymptotics. In: *Proc. Infocom 2001*.

[13] Gromoll, H.C., Robert, Ph., Zwart, A.P. (2006). Fluid limits for processor sharing queues with impatience. *Math. Oper. Res.*, to appear.

[14] Gromoll, H.C., Puha, A.L., Williams, R.J. (2002). Fluid limit of a heavily loaded processor sharing queue. *Ann. Appl. Prob.* **12**, 1–63.

[15] Gromoll, H.C., Williams, R.J. (2007). Fluid limit of a network with fair bandwidth sharing and general document size distribution. *Ann. Appl. Prob.*, to appear.

[16] Gromoll, H.C., Williams, R.J. (2007). Fluid model for a data network with alpha-fair bandwidth sharing and general document size distributions: two examples of stability. Preprint.

[17] Hansen, J., Reynolds, C., Zachary, S. (2007). Stability of processor-sharing networks with simultaneous resource requirements. *J. Appl. Prob.* **44**, 636–651.

[18] Jean-Marie, A. (2005). Presentation at the 30th Conference on the Mathematics of Operations Research, Lunteren, The Netherlands. http://www.lirmm.fr/ ajm/Talks/18Jan2005/transparents.pdf

[19] Jean-Marie, A., Robert, Ph. (1994). On the transient behavior of the processor sharing queue. *Queueing Systems* **17**, 129–136.

[20] Kelly, F.P., Williams, R.J. (2004). Fluid model for a network operating under a fair bandwidth-sharing policy. *Ann. Appl. Prob.* **14**, 1055–1083.

[21] Lu, S.H., Kumar, P.R. (1991). Distributed scheduling based on due dates and buffer priorities. *IEEE Trans. Aut. Control* **36**, 1406–1416.

[22] Massoulié, L. (2007). Structural properties of proportional fairness: stability and insensitivity. *Ann. Appl. Prob.* **17**, 809–839.

[23] Massoulié, L., Roberts, J.W. (1999). Bandwidth sharing: objectives & algorithms. In: *Proc. IEEE Infocom '99*, New York NY, USA, 1395–1403.

[24] Mo, J., Walrand, J. (2000). Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Netw.* **8**, 556–567.

[25] Puha, A.L., Stolyar, A.L., Williams, R.J. (2006). The fluid limit of an overloaded processor sharing queue. *Math. Oper. Res.* **31**, 316–350.

[26] Roberts, J.W., Massoulié, L. (1998). Bandwidth sharing and admission control for elastic traffic. In: *Proc. ITC Specialist Seminar*, Yokohama, Japan.

[27] De Veciana, G., Lee, T.-L., Konstantopoulos, T. (1999). Stability and performance analysis of networks supporting services with rate control – could the Internet be unstable? In: *Proc. Infocom '99*, New York NY, USA, 802–810.

[28] De Veciana, G., Lee, T.-L., Konstantopoulos, T. (2001). Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Trans. Netw.* **9**, 2–14.