

SRPT applied to bandwidth-sharing networks

Samuli Aalto*

TKK Helsinki University of Technology, Networking Laboratory
P.O. Box 3000, FI-02015 TKK, Finland
E-mail: samuli.aalto@tkk.fi

Urtzi Ayesta

LAAS-CNRS

7 Avenue du Colonel Roche, 31 077 Toulouse Cedex 4, France
E-mail: urtzi@laas.fr

August 28, 2007

Abstract

We consider bandwidth-sharing networks, and show how the SRPT (Shortest Remaining Processing Time) discipline can be used in order to improve the delay performance of the system. Our main idea is not to use SRPT globally between the traffic classes, which has been shown to induce instability, but rather deploy SRPT only locally within each traffic class. We show that with this approach, the performance of any stable bandwidth allocation policy can be improved. Importantly, our result is valid for any network topology and any flow size distribution. A numerical study is included to illustrate the results.

Keywords: Bandwidth-sharing network, bandwidth allocation, scheduling, SRPT

*Corresponding author: S. Aalto

A bandwidth-sharing network is a flow-level model of a data network. It consists of a set of links $l \in \mathcal{L}$ with capacities (bandwidths) c_l . The network is loaded with elastic flows (such as file transfers), each flow being associated with a route. Let \mathcal{R} denote the set of routes and $\mathcal{R}(l)$ the set of routes traversing through link l . Let n_r denote the number of flows on route $r \in \mathcal{R}$. The corresponding network state vector is denoted by $\mathbf{n} = (n_r; r \in \mathcal{R})$. The total bandwidth allocated to the n_r flows on route r is denoted by ϕ_r . These *inter-route* bandwidth allocations are feasible if, for all $l \in \mathcal{L}$,

$$\sum_{r \in \mathcal{R}(l)} \phi_r \leq c_l. \quad (1)$$

Given the inter-route bandwidth allocations ϕ_r , the *intra-route* discipline determines how bandwidth is shared among the flows using the same route. Note that in this model a flow requires the same amount of bandwidth simultaneously in each link along its route.

Within this framework, researchers have looked at both the static and dynamic settings. The static setting with a fixed number of flows is valid for a small time scale. However, in a longer time scale, the number of flows varies over time randomly, leading to a dynamic setting.

Regarding the static setting, there have been various proposals how to allocate bandwidth to the flows in a fair way. The classical fairness concept is max-min fairness [Bertsekas and Gallager (1992)]. Ever since there have been many other proposals such as proportional fairness [Kelly (1997), Kelly, Maulloo, and Tan (1998)], potential delay minimization [Massoulié and Roberts (1999)], and balanced fairness [Bonald and Proutière (2003)]. Interestingly, Mo and Walrand (2000) showed that many bandwidth allocation policies could be combined in a parametric way with the so-called α -fair allocations. It is common to all these (fair) bandwidth allocation policies that the inter-route bandwidth allocations depend on the whole network state, $\phi_r = \phi_r(\mathbf{n})$ for all r . In addition, the flows with the same route get equal shares. In other words, the intra-route discipline is typically PS (Processor Sharing).

Consider now the following dynamic setting first introduced in Massoulié and Roberts (2000). The flows on route r constitute traffic class r . New flows of class r arrive according to a Poisson process with intensity λ_r . The size of a flow of class r has a general distribution with mean β_r . Let $\rho_r = \lambda_r \beta_r$ denote the load of class r . In this dynamic setting, the primary concern is stability of the bandwidth allocation policy, that is, given a certain allocation policy, what are the conditions that ρ_r must satisfy such that the number of flows does not explode. Necessary stability conditions are that for all $l \in \mathcal{L}$,

$$\sum_{r \in \mathcal{R}(l)} \rho_r < c_l. \quad (2)$$

It has been shown that with exponentially distributed flow sizes, these conditions are also sufficient for α -fair allocations [Massoulié and Roberts (2000), de Veciana, Lee, and Konstantopoulos (2001), Bonald and Massoulié (2001), Ye (2003), Ye, Ou, and Yuan (2005)]. In the case of generally distributed flow sizes, the same result has been proven for the balanced fair [Bonald and Proutière (2003)], the max-min fair [Bramson (2005)], and the proportionally fair [Massoulié (2005)] allocations. Recently, Gromoll and Williams (2006) proved the result for any α -fair policy, however, assuming a tree topology. On the other hand, no counter-example has been given in which these fair allocation policies fail to have this property.

As a next logical step beyond stability, researchers have tried to determine how to improve the performance of bandwidth-sharing networks, by, for example, minimizing the mean number of flows in the system or the mean flow delay (i.e. sojourn time), as well. There exists a vast literature on this problem for the single-link (single-server) case. It has been shown that SRPT (Shortest Remaining Processing Time) is the optimal anticipating discipline [Schrage (1968)], while LAS (Least Attained Service) is optimal among the nonanticipating disciplines, for which the remaining service times are unknown, if the service times are of type DHR (Decreasing Hazard Rate) [Yashkov (1987)]. However, extending these policies to the network case with multiple links is not that straightforward. For

example, Verloop, Borst, and Núñez-Queija (2005) showed that SRPT and LAS may render a network unstable at arbitrarily low traffic loads. The main reason for this difference is that there may be a trade-off between the number of links used and the instantaneous departure rate of flows. In some particular cases these two criteria are non-conflicting, and the optimal policy can be determined. For example, Verloop, Borst, and Núñez-Queija (2006) considered a linear network with exponentially distributed flow sizes, and showed that, under certain assumptions on the mean flow sizes, static priority disciplines are optimal.

It is a widely open problem how to improve the delay performance in a bandwidth-sharing network with a general topology, as well as in a linear network with non-exponential flow sizes. In this paper we take a step towards the solution of this problem by using SRPT in a controlled way. To avoid possible instabilities, we do not apply SRPT globally across all traffic classes (as in Verloop, Borst, and Núñez-Queija (2005)) but only locally within each traffic class. More precisely, we show that any stable bandwidth allocation policy (such as the fair policies mentioned above) can be improved by replacing the original intra-route discipline with SRPT. Importantly, our result is valid for any network topology and any flow size distribution. A similar idea to locally apply SRPT to improve the performance of a complex system has recently been presented in the multiprocessor setting with immediate dispatching, see Avrahami and Azar (2007).

The rest of the paper is organized as follows. Section 1 introduces the notation used in this paper and the bandwidth allocation policies used in the examples. It also includes an example illustrating the possible stability problems when SRPT is applied globally. In Section 2, we present how to apply SRPT locally to improve the delay performance of any stable policy. In Section 3, we consider the allocation policies that only depend on the workload, and show how a more obvious application of SRPT performs better than the original policy. A numerical study comparing different allocation policies is given in Section 4.

1 Bandwidth-sharing networks: notation and examples

Consider a bandwidth-sharing network with a general topology, Poisson flow arrivals and generally distributed and independent flow sizes. Let A_i and S_i denote the arrival time and the size of flow i , respectively. The flow size distributions may be different for different traffic classes.

Let Π denote the family of those feasible bandwidth allocation policies for which the necessary conditions (2) are also sufficient for stability and the intra-route disciplines are work-conserving.

For any $\pi \in \Pi$ we use the following notation. Denote by $\sigma_i^\pi(t)$ the bandwidth allocated to flow i at time t under policy π . The attained service and the remaining service of flow i at time t are, respectively, defined by

$$X_i^\pi(t) = \int_0^t \sigma_i^\pi(u) du \quad \text{and} \quad Y_i^\pi(t) = S_i - X_i^\pi(t).$$

Note that $X_i^\pi(t) = 0$ and $Y_i^\pi(t) = S_i$ for all $t \leq A_i$. Flow i departs as soon as $X_i^\pi(t)$ reaches level S_i (or, equivalently, $Y_i^\pi(t)$ becomes 0). Let $\mathcal{N}_r^\pi(t)$ denote the set of flows on route r at time t . In addition, let

$$N_r^\pi(t) = |\mathcal{N}_r^\pi(t)|.$$

The total bandwidth allocated to traffic class r at time t is thus

$$Z_r^\pi(t) = \sum_{i \in \mathcal{N}_r^\pi(t)} \sigma_i^\pi(t).$$

Furthermore, let $\mathbf{R}_r^\pi(t) = (R_{rn}^\pi(t); n \in \{1, 2, \dots\})$ denote the *ordered* vector of remaining services at time t related to the class r flows with the first element $R_{r1}^\pi(t)$ referring to the maximum remaining service. Note that the first $N_r^\pi(t)$ elements of this vector are non-zero, while the rest equal zero.

Let Π^{a} denote the subfamily of those bandwidth allocation policies π that belong to Π and depend only on the number of flows, i.e., for which $Z_r^\pi(t)$ is given as follows:

$$Z_r^\pi(t) = \phi_r^\pi(\mathbf{N}^\pi(t)) \quad \text{for all } r, \tag{3}$$

where $\mathbf{N}^\pi(t) = (N_r^\pi(t); r \in \mathcal{R})$. This subfamily includes, for example, the balanced fair, the max-min fair, and the proportionally fair bandwidth allocation policies.

In addition, let Π^w denote the subfamily of those bandwidth allocation policies π that belong to Π and depend only on the workload, i.e., for which $Z_r^\pi(t)$ is given as follows:

$$Z_r^\pi(t) = \xi_r^\pi(\mathbf{W}^\pi(t)) \quad \text{for all } r, \quad (4)$$

where $\mathbf{W}^\pi(t) = (W_r^\pi(t); r \in \mathcal{R})$ and $W_r^\pi(t)$ refers to the total workload for route r at time t ,

$$W_r^\pi(t) = \sum_{i \in \mathcal{N}_r^\pi(t)} Y_i^\pi(t) = \sum_{n=1}^{N_r^\pi(t)} R_{rn}^\pi(t).$$

Finally, let Π^b denote the subfamily of those bandwidth allocation policies π that belong to Π and depend only on the set of busy flows, i.e., for which $Z_r^\pi(t)$ is given as follows:

$$Z_r^\pi(t) = \psi_r^\pi(\mathbf{B}^\pi(t)) \quad \text{for all } r, \quad (5)$$

where $\mathbf{B}^\pi(t) = (B_r^\pi(t); r \in \mathcal{R})$ and

$$B_r^\pi(t) = 1_{\{N_r^\pi(t) > 0\}} = 1_{\{W_r^\pi(t) > 0\}},$$

where $1_{\{\cdot\}}$ refers to the indicator function. Note that $\Pi^b = \Pi^n \cap \Pi^w$. Examples of resource allocation policies belonging to Π^b are given below.

1.1 Linear Network

Due to the difficulty of the problem, researchers have often considered the particular case of a *linear network*, see e.g. Massoulié and Roberts (1999, 2000), Bonald and Massoulié (2001). Figure 1 depicts a linear network consisting of two unit-capacity links, $c_1 = c_2 = 1$. In such a network there are three traffic classes, class 0 corresponding to the long route and classes 1 and 2 corresponding to the two short routes.




Figure 1: Linear network with two links.

While the theoretical results of this paper are valid for any topology, we use the linear network with two unit-capacity links as an illustrative example in our numerical studies. We will use the following three bandwidth allocation policies (with intra-route discipline PS for all classes) in our examples.

BF: This is the balanced fair bandwidth allocation policy belonging to Π^n with allocations

$$\begin{aligned} \phi_0(\mathbf{n}) &= \frac{n_0}{n_0 + n_1 + n_2}, \\ \phi_1(\mathbf{n}) &= \phi_2(\mathbf{n}) = \frac{n_1 + n_2}{n_0 + n_1 + n_2}. \end{aligned}$$

We note that for the linear network, the balanced fair allocation coincides with the proportionally fair allocation.

PR0: This is the priority allocation policy denoted by π^* in Verloop, Borst, and Núñez-Queija (2006) and belonging to Π^b . It gives preemptive priority to class 0 whenever there are flows in this class, and otherwise serves any other class with at least one flow. Thus,

$$\begin{aligned}\psi_0(\mathbf{b}) &= b_0, \\ \psi_1(\mathbf{b}) &= (1 - b_0)b_1, \\ \psi_2(\mathbf{b}) &= (1 - b_0)b_2.\end{aligned}$$

PR12: This is the priority allocation policy denoted by π^{**} in Verloop, Borst, and Núñez-Queija (2006) and also belonging to Π^b . It serves simultaneously both classes 1 and 2 whenever there are flows in both classes. Otherwise class 0 is served if non-empty. When there are no flows in class 0, any other class with at least one flow is served. Thus,

$$\begin{aligned}\psi_0(\mathbf{b}) &= (1 - b_1b_2)b_0, \\ \psi_1(\mathbf{b}) &= b_1b_2 + (1 - b_1b_2)(1 - b_0)b_1, \\ \psi_2(\mathbf{b}) &= b_1b_2 + (1 - b_1b_2)(1 - b_0)b_2.\end{aligned}$$

Recall that stability of the balanced fair allocation policy BF has been proven in Bonald and Proutière (2003). Stability of policies PR0 and PR12 is guaranteed by the fact that the whole capacity of each link is used whenever there are flows loading the link.

1.2 Stability problems with SRPT when applied globally

Below we give a simple example that illustrates the possible stability problems when SRPT is applied globally.

Consider the linear network depicted in Figure 1. Assume that flow sizes on the long route (class 0) are large compared to those on the short routes (classes 1 and 2). As an extreme example, suppose that all flow sizes are deterministic, on the short routes of size 1 and on the long route equal to $M \gg 1$. If SRPT were deployed globally, at any time the job with the smallest residual service time in the whole network would be given preference. Thus, from the point of view of classes 1 and 2 the system behaves approximately like two independent M/D/1 queues with loads ρ_1 and ρ_2 , respectively. Note that for very large M , it becomes a rare event that there is a flow of class 0 with the remaining service less than 1. Thus, from the point of view of class 0, the system is stable only if

$$\rho_0 < (1 - \rho_1)(1 - \rho_2) + \epsilon,$$

where $\epsilon \rightarrow 0$ as $M \rightarrow \infty$. For sufficiently small ϵ , this is clearly a more stringent requirement than the necessary stability condition given in (2):

$$\rho_0 < \min\{1 - \rho_1, 1 - \rho_2\}.$$

2 Delay improvement by applying SRPT locally

In this section we present the main theoretical results, which reveal how SRPT can be applied to bandwidth-sharing networks to improve the delay performance while keeping the network stable.

Let $\pi \in \Pi$ be fixed. Below it is called the *basic policy*. Denote by $\tilde{\pi}$ a modified policy for which the inter-route bandwidth allocation process is the same as for π ,

$$Z_r^{\tilde{\pi}}(t) = Z_r^{\pi}(t) \quad \text{for all } r,$$

but the intra-route disciplines may be different from the original ones. Among these modified policies, let π' denote the one that applies SRPT as the intra-route discipline for all traffic classes. More precisely, for any flow i and time t , we have

$$\sigma_i^{\pi'}(t) = \begin{cases} Z_r^{\pi}(t), & \text{if } i \in \mathcal{N}_r^{\pi'}(t) \text{ and } i = \arg \min_{j \in \mathcal{N}_r^{\pi'}(t)} Y_j^{\pi'}(t) \\ 0, & \text{otherwise.} \end{cases}$$

Note that this kind of a local application of SRPT avoids the stability problems of the global SRPT in the example given above in Section 1.2.

Our main result presented in Theorem 1 says that π' is the optimal modification in a very strong sense, minimizing the number of flows in any class r at any time t in each sample path. This is an extension of the optimality property that SRPT has in the single-link case. The formal proof is adapted from Smith (1978). We start with an intermediate result presented below in Proposition 1, from which the main result easily follows.

Proposition 1 *Let $\pi \in \Pi$, $r \in \mathcal{R}$, $t \geq 0$, and $k \in \{1, 2, \dots\}$. Then*

$$\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t) \leq \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t) \quad (6)$$

for any modification $\tilde{\pi}$ of the basic policy π .

Proof Consider a fixed modification $\tilde{\pi}$. Let r and t be fixed. Assume that (6) is true in the interval $[0, t)$ for any k .

1° Assume first that time t is an arrival epoch of, say, the j th flow with size S_j . Let

$$n' = \arg \min\{n : R_{rn}^{\pi'}(t) = S_j\}, \quad \tilde{n} = \arg \min\{n : R_{rn}^{\tilde{\pi}}(t) = S_j\}.$$

If $n' < k$ and $\tilde{n} < k$, then

$$\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t) = \sum_{n=k-1}^{\infty} R_{rn}^{\pi'}(t^-) \leq \sum_{n=k-1}^{\infty} R_{rn}^{\tilde{\pi}}(t^-) = \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t).$$

If $n' < k$ and $\tilde{n} \geq k$, then

$$\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t) \leq S_j + \sum_{n=k}^{\infty} R_{rn}^{\pi'}(t^-) \leq S_j + \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t^-) = \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t).$$

If $n' \geq k$ and $\tilde{n} < k$, then

$$\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t) \leq \sum_{n=k-1}^{\infty} R_{rn}^{\pi'}(t^-) \leq \sum_{n=k-1}^{\infty} R_{rn}^{\tilde{\pi}}(t^-) = \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t).$$

If $n' \geq k$ and $\tilde{n} \geq k$, then

$$\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t) = S_j + \sum_{n=k}^{\infty} R_{rn}^{\pi'}(t^-) \leq S_j + \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t^-) = \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t).$$

Thus, (6) is true in the closed interval $[0, t]$ for any k .

On the other hand, if time t is not an arrival epoch, the claim is justified by a continuity argument.

2° Let $h > 0$ and suppose that no flows belonging to class r arrive in the interval $(t, t+h]$. If $\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t+h) = 0$, then surely

$$\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t+h) = 0 \leq \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t+h).$$

Assume now that $\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t+h) > 0$. Due to the intra-route SRPT discipline, we have

$$\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t+h) = \sum_{n=k}^{\infty} R_{rn}^{\pi'}(t) - \int_t^{t+h} Z_r^{\pi}(u) du.$$

By 1°, we have

$$\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t+h) \leq \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t) - \int_t^{t+h} Z_r^{\pi}(u) du.$$

On the other hand, it is easy to see that

$$\sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t+h) \geq \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t) - \int_t^{t+h} Z_r^{\pi}(u) du.$$

Thus,

$$\sum_{n=k}^{\infty} R_{rn}^{\pi'}(t+h) \leq \sum_{n=k}^{\infty} R_{rn}^{\tilde{\pi}}(t+h),$$

which completes the proof. \square

Let $r \in \mathcal{R}$ be fixed for a short while. Recall that the intra-route disciplines are work-conserving for all $\pi \in \Pi$, and the same is true for the modification π' by its definition. Thus, if we restrict ourselves to the modifications $\tilde{\pi}$ (including π and π') that apply only work conserving intra-route disciplines, then the total workload related to route r remains the same implying that

$$\sum_{n=1}^{\max\{N_r^{\tilde{\pi}}(t), N_r^{\pi'}(t)\}} R_{rn}^{\tilde{\pi}}(t) = W_r^{\tilde{\pi}}(t) = W_r^{\pi'}(t) = \sum_{n=1}^{\max\{N_r^{\tilde{\pi}}(t), N_r^{\pi'}(t)\}} R_{rn}^{\pi'}(t).$$

Combining this with (6), we can say that the (truncated) vector

$$(R_{rn}^{\tilde{\pi}}(t); n = 1, \dots, \max\{N_r^{\tilde{\pi}}(t), N_r^{\pi'}(t)\})$$

is *majorized* by the (truncated) vector

$$(R_{rn}^{\pi'}(t); n = 1, \dots, \max\{N_r^{\tilde{\pi}}(t), N_r^{\pi'}(t)\}),$$

cf. Chang and Yao (1993), Hirayama and Kijima (1992). As a consequence, π' minimizes [maximizes] any *Schur-concave* [*Schur-convex*] function of this truncated vector among such modifications. This could be applied to prove that π' minimizes the number of flows in each route at any time among such modifications. However, below we use a direct argument to prove that π' has this property among *all* possible modifications.

Theorem 1 *Let $\pi \in \Pi$, $r \in \mathcal{R}$ and $t \geq 0$. Then*

$$N_r^{\pi'}(t) \leq N_r^{\tilde{\pi}}(t)$$

for any modification $\tilde{\pi}$ of the basic policy π .

Proof Assume that $N_r^{\tilde{\pi}}(t) = n$. Then, by Proposition 1,

$$\sum_{j=n+1}^{\infty} R_{rj}^{\pi'}(t) \leq \sum_{j=n+1}^{\infty} R_{rj}^{\tilde{\pi}}(t) = 0,$$

implying that $N_r^{\pi'}(t) \leq n$. \square

Theorem 1 expresses a stochastic ordering between the flow number processes of the two systems. This pathwise result implies a corresponding result for mean values, and furthermore, by taking the sum over all traffic classes, we find that π' minimizes the mean total number of flows \bar{N} and the mean delay as well by Little's result. As a by-product we also see that π' is stable.

Corollary 1 *Let $\pi \in \Pi$. Then $\bar{N}^{\pi'} \leq \bar{N}^{\tilde{\pi}}$ for any modification $\tilde{\pi}$ of the basic policy π .*

In the following section we consider another SRPT-related modification, denoted by π^* , and defined for any basic policy $\pi \in \Pi^n \cup \Pi^w$. If $\pi \in \Pi^n$, there are functions ϕ_r^π such that

$$Z_r^\pi(t) = \phi_r^\pi(\mathbf{N}^\pi(t)).$$

Let $\pi^* \in \Pi^n$ be the policy such that

$$Z_r^{\pi^*}(t) = \phi_r^\pi(\mathbf{N}^{\pi^*}(t))$$

and which deploys SRPT as intra-route discipline in every route r . If $\pi \in \Pi^w$, there are functions ξ_r^π such that

$$Z_r^\pi(t) = \xi_r^\pi(\mathbf{W}^\pi(t)).$$

In this case, let $\pi^* \in \Pi^w$ denote the policy such that

$$Z_r^{\pi^*}(t) = \xi_r^\pi(\mathbf{W}^{\pi^*}(t))$$

and which deploys SRPT as intra-route discipline in every route r .

Note the difference between π^* and π' : for π' we require that the inter-route bandwidth allocation process be the same as for the basic policy π , while for π^* we only require that the corresponding allocation function be the same.

3 On the optimality of π^*

Below we first demonstrate by a counter-example that a similar pathwise result as given in Theorem 1 for π' is not valid for π^* in general. Thereafter, in Theorem 2, we give a sufficient condition under which π^* beats the basic policy π and all its modifications $\tilde{\pi}$ pathwise.

Example 1 Consider the linear network depicted in Figure 1 with unit-capacity links. Assume that $A_1 = 0$, $A_2 = \epsilon$, and $A_3 = 2\epsilon$ with the first arriving flow belonging to class 0 and the other two to class 1. In addition, assume that $S_1 = 3$ and $S_2 = S_3 = 2$. The basic policy π is chosen to be the balanced fair bandwidth allocation $\text{BF} \in \Pi^n$ defined in Section 1.1 so, as $\epsilon \rightarrow 0$, $Z_0^\pi(t) = 1 - Z_1^\pi(t) = 1/3$ for $0 \leq t \leq 6$, and $Z_0^\pi(t) = 1$ for $6 < t \leq 7$. The total number of flows (for policies π, π', π^*) develops in time as given in Figure 2, in which $\epsilon \rightarrow 0$. After the departure of a flow (approximately) at time 3 under π^* , class 1 gets less bandwidth than under π and π' . As a result, there is an interval (approximately from 6 to 7) during which we have

$$N^\pi(t) = N^{\pi'}(t) = 1 < 2 = N^{\pi^*}(t),$$

Thus, π^* does not beat the basic policy π pathwise. □

The following theorem is the main result of this section.

Theorem 2 *Let $\pi \in \Pi^w$, $r \in \mathcal{R}$ and $t \geq 0$. Then*

$$N_r^{\pi^*}(t) = N_r^{\pi'}(t) \leq N_r^{\tilde{\pi}}(t)$$

for any modification $\tilde{\pi}$ of the basic policy π .




Figure 2: The total number of flows $N(t)$ as a function of time t for policies π, π', π^* (from the left to the right).

Proof Since $\Pi^w \subset \Pi$, the latter part ($N_r^{\pi'}(t) \leq N_r^{\tilde{\pi}}(t)$) is due to Theorem 1. For the former part ($N_r^{\pi^*}(t) = N_r^{\pi'}(t)$), we consider the remaining services and show that for any i and t

$$Y_i^{\pi^*}(t) = Y_i^{\pi'}(t), \quad (7)$$

from which the claim follows.

Let t be fixed. Assume that (7) is true in the interval $[0, t)$ for any i .

1° If time t is an arrival epoch of, say, the j th customer, then clearly

$$Y_j^{\pi^*}(t) = Y_j^{\pi'}(t) = S_j$$

and

$$Y_i^{\pi^*}(t) = Y_i^{\pi^*}(t^-) = Y_i^{\pi'}(t^-) = Y_i^{\pi'}(t)$$

for any $i \neq j$. On the hand, if time t is not an arrival epoch, the previous equation is true for all i .

2° Moreover, for any i the remaining services are decreasing with the same rate

$$\sigma_i^{\pi^*}(t) = \sigma_i^{\pi'}(t).$$

This is due to the same intra-route discipline (SRPT) applied in π^* and π' and the fact that the inter-route bandwidth allocations are the same for any r ,

$$Z_r^{\pi^*}(t) = Z_r^{\pi'}(t). \quad (8)$$

Due to (4), equation (8) follows from

$$W_r^{\pi^*}(t) = \sum_{i \in N_r^{\pi^*}(t)} Y_i^{\pi^*}(t) = \sum_{i \in N_r^{\pi'}(t)} Y_i^{\pi'}(t) = W_r^{\pi'}(t).$$

This completes the proof. \square

Theorem 2 says that policies π^* and π' indeed result in the same system in the stochastic sense. Again, this pathwise result implies a corresponding result for mean values, and furthermore, by taking the sum over all traffic classes we get the following corollary.

Corollary 2 *Let $\pi \in \Pi^w$. Then $\bar{N}^{\pi^*} = \bar{N}^{\pi'} \leq \bar{N}^{\tilde{\pi}}$ for any modification $\tilde{\pi}$ of the basic policy π .*

4 Numerical results

We ran simulations assuming the linear network with two unit-capacity links and three classes depicted in Figure 1. In these simulations we compared a basic policy π with its modifications π' and π^* . To improve the accuracy of the comparison, the same realizations of the arrival process were used for all these policies. The three basic policies applied (BF, PR0, PR12) were defined in Section 1.1. Since flow sizes in the Internet have been modelled by hyperexponential and Pareto distributions [Crovella and Bestavros (1996), Feldmann and Whitt (1997)], we included these distributions in our simulations. Two aspects were considered, the variation in the flow size distribution (Section 4.1) and the heterogeneity in the classwise mean flow sizes (Section 4.2).

In fact, we additionally considered two α -fair policies (corresponding to $\alpha = 2$ and $\alpha = 10$). However, the results of these simulations were virtually identical to those for the BF policy, which in our example network coincides with the α -fair policy with $\alpha = 1$, and are therefore not presented here. We just conclude that the performance of the α -fair basic policies π and their SRPT-modifications π' and π^* seems to be largely insensitive to the parameter α .

4.1 Effect of the variability in the flow size distribution

To get an idea about the effect of flow-size variability, we considered three very different flow size distributions: deterministic, exponential, and hyperexponential. We assumed the mean flow sizes of $\beta = 4/5$ for all classes. The hyperexponential distribution with this mean was chosen to be

$$P\{S > x\} = p_1 e^{-\mu_1 x} + p_2 e^{-\mu_2 x}$$

with parameters $p_1 = 9/10$, $\mu_1 = 45/4$, $p_2 = 1/10$, and $\mu_2 = 5/36$. In this simulation set, we let the flow arrival rates λ_r vary in such a way that $\lambda_0 + \lambda_1 + \lambda_2 = 1$ and $\lambda_1 = \lambda_2 = (1 - \lambda_0)/2$. We note that, in the case of the exponential distribution, the policy PR12 is optimal for all the parameter combinations used in these simulations, cf. Verloop, Borst, and Núñez-Queija (2006).




Figure 3: Mean total number of flows as a function of λ_0 for deterministic (left), exponential (middle) and hyperexponential (right) flow sizes. BF = solid black, BF' = dashed black, PR12 = solid gray, PR12' = dashed gray.

The mean total number of flows (\bar{N}^π and $\bar{N}^{\pi'}$) is plotted in Figure 3 for basic policies $\pi =$ BF, PR12 and their corresponding modifications $\pi' =$ BF', PR12'. The more detailed results related to different flow size distributions are given in Tables 1-3. Since classes 1 and 2 are symmetric, only one of them (class 1) appears in the tables.

We make the following observations based on these numerical results:

- As Theorem 1 states, we have $\bar{N}_r^{\pi'} \leq \bar{N}_r^\pi$ for all basic policies π , all flow size distributions, and all classes r .
- As Theorem 2 states, we have $\bar{N}_r^{\pi^*} = \bar{N}_r^{\pi'} \leq \bar{N}_r^\pi$ for the basic policies PR0 and PR12, all flow size distributions, and all classes r .
- Interestingly, in all our numerical cases we also have $\bar{N}_r^{\pi^*} \leq \bar{N}_r^\pi$ for the basic policy BF, all flow size distributions, and all classes r . In addition, $\bar{N}^{\pi^*} \leq \bar{N}^{\pi'} \leq \bar{N}^\pi$ in all our experiments. Numerically, the difference between \bar{N}^{π^*} and $\bar{N}^{\pi'}$ is, however, very small.

Table 1: Mean number of flows for deterministic flow sizes.

Basic policy BF										
λ_0	λ_1	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.0	4.00	2.40	2.40	0.00	0.00	0.00	4.00	2.40	2.40
0.8	0.1	2.79	1.77	1.84	0.29	0.25	0.22	3.37	2.28	2.28
0.6	0.2	1.74	1.20	1.24	0.44	0.38	0.35	2.63	1.95	1.94
0.4	0.3	0.98	0.75	0.75	0.54	0.45	0.44	2.07	1.65	1.64
0.2	0.4	0.42	0.36	0.35	0.61	0.50	0.50	1.65	1.36	1.35
0.0	0.5	0.00	0.00	0.00	0.67	0.53	0.53	1.33	1.07	1.07

Basic policy PR0										
λ_0	λ_1	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.0	4.00	2.40	2.40	0.00	0.00	0.00	4.00	2.40	2.40
0.8	0.1	1.77	1.21	1.21	0.59	0.51	0.51	2.96	2.22	2.22
0.6	0.2	0.92	0.70	0.70	0.71	0.58	0.58	2.34	1.86	1.86
0.4	0.3	0.47	0.40	0.40	0.72	0.58	0.58	1.91	1.55	1.55
0.2	0.4	0.19	0.17	0.17	0.70	0.56	0.56	1.58	1.29	1.29
0.0	0.5	0.00	0.00	0.00	0.67	0.53	0.53	1.33	1.07	1.07

Basic policy PR12										
λ_0	λ_1	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.0	4.00	2.40	2.40	0.00	0.00	0.00	4.00	2.40	2.40
0.8	0.1	2.05	1.36	1.36	0.43	0.36	0.36	2.93	2.07	2.07
0.6	0.2	1.16	0.85	0.85	0.55	0.44	0.44	2.26	1.73	1.73
0.4	0.3	0.63	0.50	0.50	0.60	0.48	0.48	1.83	1.47	1.47
0.2	0.4	0.26	0.23	0.23	0.63	0.51	0.51	1.53	1.25	1.25
0.0	0.5	0.00	0.00	0.00	0.67	0.53	0.53	1.33	1.07	1.07

- As is well known, the results for the basic policy BF are insensitive to the flow size distribution. The improvement achieved with modifications π' and π^* is larger for the (more variable) exponential and hyperexponential flow size distributions. However, the improvement when replacing the exponential distribution with the hyperexponential one is quite small. To summarize, while the performance of BF is insensitive to the flow size distribution, the performance of its SRPT-modifications increases with flow-size variability.
- Basic policy PR0 is better than BF only for the deterministic flow size distribution, but much worse for the more variable hyperexponential distribution. The SRPT-modifications π' and π^* of PR0 are worse than the corresponding modifications of BF even for the exponential distribution. Indeed, the performance of basic policy PR0 (as well as its modifications) decreases as the flow-size variability increases.
- Basic policy PR12 is better than BF for deterministic and exponential flow size distributions, but worse for the more variable hyperexponential distribution. As for PR0, the performance of basic policy PR12 (as well as its modifications) decreases as the flow-size variability increases.

To get a more systematic view about the effect of flow-size variability, we performed another series of simulations where we kept the same mean flow size $\beta = E[S] = 4/5$ as above, but let the coefficient of variation, $C[S] = \sqrt{\text{Var}[S]}/E[S]$, vary between 0 and 4, which we implemented with deterministic, gamma, exponential, and hyperexponential distributions. The arrival rates were fixed to $\lambda_0 = 0.8, \lambda_1 = \lambda_2 = 0.1$.

Table 2: Mean number of flows for exponential flow sizes.

Basic policy BF										
λ_0	λ_1	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.0	4.00	1.88	1.88	0.00	0.00	0.00	4.00	1.88	1.88
0.8	0.1	2.79	1.50	1.57	0.29	0.25	0.20	3.37	2.01	1.97
0.6	0.2	1.74	1.08	1.12	0.44	0.37	0.34	2.63	1.83	1.79
0.4	0.3	0.98	0.71	0.72	0.54	0.44	0.43	2.07	1.60	1.57
0.2	0.4	0.42	0.36	0.35	0.61	0.48	0.48	1.65	1.33	1.32
0.0	0.5	0.00	0.00	0.00	0.67	0.52	0.52	1.33	1.04	1.04

Basic policy PR0										
λ_0	λ_1	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.0	4.00	1.88	1.88	0.00	0.00	0.00	4.00	1.88	1.88
0.8	0.1	1.78	1.08	1.08	0.80	0.70	0.70	3.37	2.49	2.49
0.6	0.2	0.92	0.67	0.67	0.86	0.72	0.72	2.64	2.10	2.10
0.4	0.3	0.47	0.39	0.39	0.81	0.65	0.65	2.09	1.70	1.70
0.2	0.4	0.19	0.17	0.17	0.73	0.58	0.58	1.65	1.33	1.33
0.0	0.5	0.00	0.00	0.00	0.67	0.52	0.52	1.33	1.04	1.04

Basic policy PR12										
λ_0	λ_1	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.0	4.00	1.88	1.88	0.00	0.00	0.00	4.00	1.88	1.88
0.8	0.1	2.06	1.20	1.20	0.52	0.45	0.45	3.09	2.09	2.09
0.6	0.2	1.19	0.81	0.81	0.63	0.51	0.51	2.41	1.82	1.82
0.4	0.3	0.63	0.50	0.50	0.63	0.51	0.51	1.91	1.52	1.52
0.2	0.4	0.27	0.24	0.24	0.65	0.51	0.51	1.57	1.27	1.27
0.0	0.5	0.00	0.00	0.00	0.67	0.52	0.52	1.33	1.04	1.04

The mean total number of flows (\bar{N}^π , $\bar{N}^{\pi'}$ and \bar{N}^{π^*}) is plotted in Figure 4 for basic policies $\pi = \text{BF, PR0, PR12}$ and their corresponding modifications π' and π^* . In fact the results for the modifications π' and π^* are so close to each other that the difference is almost indistinguishable from the figure. Based on this simulation set, we make the following further observations:

- While the performance of BF is insensitive to the flow size distribution, the performance of basic policies PR0 and PR12 decreases as the flow-size variability increases.
- The performance gain achieved by the SRPT modifications ($\bar{N}^\pi - \bar{N}^{\pi'}$ and $\bar{N}^\pi - \bar{N}^{\pi^*}$) is rather insensitive to the flow size distribution. It seems to be increasing, but only very slowly, as a function of $C[S]$.

4.2 Effect of mean flow-size heterogeneity across routes

The effect of the heterogeneity in the classwise mean flow sizes was studied with the following classwise Pareto flow size distributions:

$$P\{S_r > x\} = \left(\frac{k_r}{x}\right)^\alpha, \quad x \geq k_r,$$

where we used a joint value $\alpha = 2.2$ for the shape parameter but let the location parameter k_r vary for different classes to achieve the desired classwise mean flow sizes $E[S_r] = \beta_r$. For all these Pareto

Table 3: Mean number of flows for hyperexponential flow sizes.

Basic policy BF										
λ_0	λ_1	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.0	4.00	1.84	1.84	0.00	0.00	0.00	4.00	1.84	1.84
0.8	0.1	2.79	1.47	1.54	0.29	0.26	0.20	3.37	1.98	1.94
0.6	0.2	1.74	1.08	1.12	0.44	0.37	0.33	2.63	1.82	1.78
0.4	0.3	0.98	0.70	0.71	0.54	0.44	0.43	2.07	1.58	1.56
0.2	0.4	0.42	0.36	0.34	0.61	0.48	0.48	1.65	1.33	1.31
0.0	0.5	0.00	0.00	0.00	0.67	0.52	0.52	1.33	1.03	1.03

Basic policy PR0										
λ_0	λ_1	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.0	4.00	1.84	1.84	0.00	0.00	0.00	4.00	1.84	1.84
0.8	0.1	1.77	1.07	1.07	3.64	3.49	3.49	9.10	8.09	8.09
0.6	0.2	0.94	0.68	0.68	3.07	2.87	2.87	7.10	6.44	6.44
0.4	0.3	0.48	0.39	0.39	2.07	1.88	1.88	4.62	4.16	4.16
0.2	0.4	0.19	0.18	0.18	1.31	1.13	1.13	2.80	2.44	2.44
0.0	0.5	0.00	0.00	0.00	0.67	0.52	0.52	1.33	1.03	1.03

Basic policy PR12										
λ_0	λ_1	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.0	4.00	1.84	1.84	0.00	0.00	0.00	4.00	1.84	1.84
0.8	0.1	2.18	1.28	1.28	1.11	0.86	0.86	4.40	3.01	3.01
0.6	0.2	1.34	0.95	0.95	1.08	0.82	0.82	3.46	2.58	2.58
0.4	0.3	0.84	0.70	0.70	0.94	0.73	0.73	2.72	2.14	2.14
0.2	0.4	0.41	0.38	0.38	0.80	0.62	0.62	2.01	1.61	1.61
0.0	0.5	0.00	0.00	0.00	0.67	0.52	0.52	1.33	1.03	1.03

distributions, the coefficient of variation is the same, depending only on the shape parameter α ,

$$C[S_r] = \frac{1}{\sqrt{\alpha}\sqrt{\alpha-2}} \approx 1.5.$$

In this simulation set we fixed the arrival rates to be $\lambda_0 = \lambda_1 = \lambda_2 = 0.3$, and let the classwise means β_r vary in such a way that $\beta_0 + \beta_1 + \beta_2 = 2.2$.

The results related to the Pareto flow sizes are given in Tables 4 and 5. The former one includes the classwise mean values, while the latter gives us the totals.

In fact, we also ran the same simulation set by replacing each Pareto distribution with a corresponding hyperexponential distribution (with the same mean and the same coefficient of variation). Although the results did not change much, they are given below in Tables 6 and 7 for completeness.

We make the following observations based on these numerical results:

- As Theorem 1 states, we have $\bar{N}_r^{\pi^*} \leq \bar{N}_r^\pi$ for all basic policies π and all classes r .
- In all our numerical cases we also have $\bar{N}_r^{\pi^*} \leq \bar{N}_r^{\pi'}$ for all basic policies π and all classes r . In addition, $\bar{N}^{\pi^*} \leq \bar{N}^{\pi'} \leq \bar{N}^\pi$ in all our experiments. Numerically, the difference between \bar{N}^{π^*} and $\bar{N}^{\pi'}$ is, however, very small.
- While the performance of BF is insensitive to the flow size distribution, the performance of basic policies PR0 and PR12, when considering the total means, is slightly better for the Pareto distributions (than for the corresponding hyperexponential distributions). However, for the classwise means, the situation is no longer that clear.




Figure 4: Mean total number of flows as a function of $C[S]$ for basic policies $\pi = \text{BF}$ (left), $\pi = \text{PR0}$ (middle) and $\pi = \text{PR12}$ (right), together with their modifications π' and π^* . $\pi =$ solid black, $\pi' =$ dashed black, $\pi^* =$ dashed gray. The solid gray line is the analytically calculated value for the basic policy BF, which is insensitive to the variation in the distribution.

- The performance gain obtained by π' and π^* with the Pareto distributions is of the same magnitude as that obtained with the corresponding hyperexponential distributions.

5 Summary

The problem of efficient scheduling of elastic flows in bandwidth-sharing networks is largely open. Until now, most of the research has focused on stability issues, but it seemed hard to combine the stability of the network together with size-based scheduling disciplines, such as SRPT. In this paper, we have shown that with a simple mechanism, it is easy to modify any stable bandwidth allocation policy in such a way that flows within each class are served according to SRPT. The obtained discipline preserves the stability property of the original policy, but it comes with the benefit of reducing the mean number of flows in each of the classes. Numerical results show that the reduction can be significant. The reduction seems to be of the same magnitude for flow size distributions with the same first and second order characteristics (i.e., mean and variance). Our numerical experiments further propose that the performance of the α -fair policies is largely insensitive to the parameter α .

Table 4: Classwise mean number of flows for Pareto flow sizes.

Basic policy BF											
β_0	β_1	β_2	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}_2^π	$\bar{N}_2^{\pi'}$	$\bar{N}_2^{\pi^*}$
1.0	0.8	0.4	0.74	0.57	0.57	0.52	0.42	0.41	0.21	0.19	0.18
0.8	0.8	0.6	0.56	0.45	0.45	0.46	0.38	0.38	0.31	0.27	0.26
0.6	0.8	0.8	0.40	0.34	0.33	0.41	0.35	0.35	0.41	0.35	0.34
0.4	0.8	1.0	0.26	0.23	0.22	0.38	0.32	0.32	0.52	0.42	0.42

Basic policy PR0											
β_0	β_1	β_2	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}_2^π	$\bar{N}_2^{\pi'}$	$\bar{N}_2^{\pi^*}$
1.0	0.8	0.4	0.43	0.36	0.36	0.92	0.74	0.74	0.51	0.46	0.46
0.8	0.8	0.6	0.32	0.28	0.28	0.64	0.52	0.52	0.47	0.40	0.40
0.6	0.8	0.8	0.22	0.20	0.20	0.50	0.42	0.42	0.50	0.42	0.42
0.4	0.8	1.0	0.14	0.13	0.13	0.40	0.34	0.34	0.55	0.45	0.45

Basic policy PR12											
β_0	β_1	β_2	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}_2^π	$\bar{N}_2^{\pi'}$	$\bar{N}_2^{\pi^*}$
1.0	0.8	0.4	0.50	0.41	0.41	0.79	0.60	0.60	0.25	0.22	0.22
0.8	0.8	0.6	0.38	0.32	0.32	0.57	0.46	0.46	0.35	0.30	0.30
0.6	0.8	0.8	0.27	0.24	0.24	0.45	0.37	0.37	0.44	0.37	0.37
0.4	0.8	1.0	0.17	0.16	0.16	0.38	0.32	0.32	0.53	0.43	0.43

Table 5: Total mean number of flows for Pareto flow sizes.

			BF			PR0			PR12		
β_0	β_1	β_2	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.8	0.4	1.47	1.18	1.16	1.87	1.56	1.56	1.53	1.23	1.23
0.8	0.8	0.6	1.33	1.10	1.09	1.42	1.20	1.20	1.31	1.09	1.09
0.6	0.8	0.8	1.23	1.03	1.02	1.22	1.04	1.04	1.16	0.98	0.98
0.4	0.8	1.0	1.15	0.98	0.97	1.09	0.92	0.92	1.08	0.91	0.91

Table 6: Classwise mean number of flows for hyperexponential flow sizes.

Basic policy BF											
β_0	β_1	β_2	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}_2^π	$\bar{N}_2^{\pi'}$	$\bar{N}_2^{\pi^*}$
1.0	0.8	0.4	0.74	0.57	0.57	0.52	0.43	0.42	0.21	0.19	0.18
0.8	0.8	0.6	0.56	0.45	0.45	0.46	0.38	0.38	0.31	0.27	0.27
0.6	0.8	0.8	0.40	0.34	0.33	0.41	0.35	0.34	0.41	0.35	0.35
0.4	0.8	1.0	0.26	0.23	0.22	0.38	0.32	0.32	0.52	0.42	0.42

Basic policy PR0											
β_0	β_1	β_2	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}_2^π	$\bar{N}_2^{\pi'}$	$\bar{N}_2^{\pi^*}$
1.0	0.8	0.4	0.43	0.36	0.36	0.91	0.76	0.76	0.54	0.50	0.50
0.8	0.8	0.6	0.32	0.28	0.28	0.67	0.57	0.57	0.50	0.45	0.45
0.6	0.8	0.8	0.22	0.20	0.20	0.52	0.44	0.44	0.51	0.44	0.44
0.4	0.8	1.0	0.14	0.13	0.13	0.41	0.36	0.36	0.56	0.46	0.46

Basic policy PR12											
β_0	β_1	β_2	\bar{N}_0^π	$\bar{N}_0^{\pi'}$	$\bar{N}_0^{\pi^*}$	\bar{N}_1^π	$\bar{N}_1^{\pi'}$	$\bar{N}_1^{\pi^*}$	\bar{N}_2^π	$\bar{N}_2^{\pi'}$	$\bar{N}_2^{\pi^*}$
1.0	0.8	0.4	0.50	0.41	0.41	0.78	0.63	0.63	0.29	0.27	0.27
0.8	0.8	0.6	0.38	0.33	0.33	0.57	0.48	0.48	0.38	0.33	0.33
0.6	0.8	0.8	0.28	0.25	0.25	0.45	0.38	0.38	0.46	0.39	0.39
0.4	0.8	1.0	0.19	0.18	0.18	0.39	0.33	0.33	0.53	0.43	0.43

Table 7: Total mean number of flows for hyperexponential flow sizes.

			BF			PR0			PR12		
β_0	β_1	β_2	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}	\bar{N}^π	$\bar{N}^{\pi'}$	\bar{N}^{π^*}
1.0	0.8	0.4	1.47	1.18	1.16	1.87	1.62	1.62	1.57	1.30	1.30
0.8	0.8	0.6	1.33	1.10	1.09	1.49	1.29	1.29	1.34	1.14	1.14
0.6	0.8	0.8	1.23	1.03	1.02	1.25	1.07	1.07	1.20	1.02	1.02
0.4	0.8	1.0	1.15	0.97	0.96	1.11	0.94	0.94	1.11	0.94	0.94

References

- [1] Avrahami, N. and Y. Azar. (2007). “Minimizing total low time and total completion time with immediate dispatching.” *Algorithmica* 47, 253–268.
- [2] Bertsekas, D. and R. Gallager. (1992). *Data Networks*. Second Edition. Prentice-Hall.
- [3] Bonald, T. and L. Massoulié. (2001). “Impact of fairness on Internet performance.” In *ACM Sigmetrics 2001*, Cambridge, MA, pp. 82–91.
- [4] Bonald, T. and A. Proutière. (2003). “Insensitive bandwidth sharing in data networks.” *Queueing Systems* 44, 69–100.
- [5] Bramson, M. (2005). “Stability of networks for max-min fair routing.” In *13th INFORMS Applied Probability Conference*, Ottawa, Canada.
- [6] Chang, C.S. and D.D. Yao. (1993). “Rearrangement, majorization, and stochastic scheduling.” *Mathematics of Operations Research* 18, 658–684.
- [7] Crovella, M. and A. Bestavros. (1996) “Self-similarity in world wide web traffic: evidence and possible causes.” In *ACM SIGMETRICS 1996*, Philadelphia, PA, pp. 160–169.
- [8] Feldmann, A. and W. Whitt. (1997) “Fitting mixtures of exponentials to long-tail distributions to analyze network performance models.” In *IEEE Infocom 1997*, Kobe, Japan, pp. 1096–1104.
- [9] Gromoll, H.C. and R.J. Williams. (2006). “Fluid limit of a network with fair bandwidth sharing and general document size distribution.” Preprint.
- [10] Hirayama, T. and M. Kijima. (1992) “Single machine scheduling problem when the machine capacity varies stochastically.” *Operations Research* 40, 376–383
- [11] Kelly, F.P. (1997). “Charging and rate control for elastic traffic.” *European Transactions on Telecommunications* 8, 33–37.
- [12] Kelly, F.P., A.K. Maulloo, and D.K.H. Tan. (1998). “Rate control in communication networks: shadow prices, proportional fairness and stability.” *Journal of the Operational Research Society* 49, 237–252.
- [13] Massoulié, L. and J. Roberts. (1999). “Bandwidth sharing: objectives and algorithms.” In *IEEE Infocom 1999*, New York, NY, pp. 1395–1403.
- [14] Massoulié, L. and J. Roberts. (2000). “Bandwidth sharing and admission control for elastic traffic.” *Telecommunication Systems* 15, 185–201.
- [15] Massoulié, L. (2005). “Structural properties of proportional fairness: stability and insensitivity.” Preprint.
- [16] Mo, J. and J. Walrand. (2000). “Fair end-to-end window-based congestion control.” *IEEE Transactions on Networking* 8, 556–567.
- [17] Schrage, L.E. (1968). “A proof of the optimality of the shortest remaining processing time discipline.” *Operations Research* 16, 687–690.
- [18] Smith, D.R. (1978). “A new proof of the optimality of the shortest remaining processing time discipline.” *Operations Research* 26, 197–199.
- [19] de Veciana, G., T.-J. Lee, and T. Konstantopoulos. (2001). “Stability and performance analysis of networks supporting elastic services.” In *IEEE Transactions on Networking* 8, 556–567.

- [20] Verloop, I.M., S.C. Borst, and R. Núñez-Queija. (2005). “Stability of size-based scheduling disciplines in resource-sharing networks.” *Performance Evaluation* 62, 247–262.
- [21] Verloop, I.M., S.C. Borst, and R. Núñez-Queija. (2006). “Delay optimization in bandwidth-sharing networks.” In *CISS 2006*, Princeton, NJ.
- [22] Yashkov, S.F. (1987). “Processor-sharing queues: Some progress in analysis.” *Queueing Systems* 2, 1–17.
- [23] Ye, H.-Q. (2003). “Stability of data networks under an optimization-based bandwidth allocation.” *IEEE Transactions on Automatic Control*, 48, 1238–1242.
- [24] Ye, H.-Q., J. Ou, and X.-M. Yuan. (2005). “Stability of data networks: stationary and bursty models.” *Operations Research*, 53, 107–125.