

SRPT applied to bandwidth-sharing networks

Samuli Aalto^{a,1}, Urtzi Ayesta^{b,2}

^a Helsinki University of Technology, Networking Laboratory, P.O. Box 3000, FIN-02015 TKK, Finland

^b LAAS-CNRS, 7 Avenue du Colonel Roche, 31 077 Toulouse Cedex 4, France

Extended abstract

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 i , each flow being associated with a route r_i . Let \mathcal{R} denote the set of routes and $\mathcal{R}(l)$ the set of routes traversing through link l . The total bandwidth allocated to all flows with route r is denoted by ϕ_r . These *inter-route* bandwidth allocations are feasible if, for all links $l \in \mathcal{L}$,

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

Given the inter-route bandwidth allocations ϕ_r , the *intra-route* discipline determines how bandwidth is shared among the flows using the same route.

Regarding the static setting with a fixed number of flows, there has been various proposals how to allocate bandwidth to the flows in a fair way. Let n_r denote the number of flows on route r . The corresponding network state vector is denoted by \mathbf{n} . The classical fairness concept is max-min fairness. Kelly launched the proportional fairness concept [3]. These bandwidth allocation policies were combined in a parametric way by Mo and Walrand, whose concept is called α -fairness [4]. Even a more general class of bandwidth allocations, based on U -utility maximization, was introduced by Ye et al. [10]. Still another fairness concept was introduced by Bonald and Proutière, namely balanced fairness [2]. 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 routes r . In addition, the flows i with the same route $r_i = r$ get equal shares. In other words, the intra-route discipline is PS.

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 is varying randomly. Consider now the following dynamic setting (first introduced in [5]). 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. Necessary stability conditions require that, for all links $l \in \mathcal{L}$,

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

It has been shown that these conditions are also sufficient for the stability of any of the fair bandwidth allocation policies mentioned above, see [5, 6, 1, 10, 2].

On the other hand, it has been demonstrated that for some priority based bandwidth allocation policies the necessary conditions (2) are not sufficient for stability. For example, Bonald and Massoulié showed that this is the case for the maximum throughput bandwidth

¹email: samuli.aalto@tkk.fi

²email: urtzi@laas.fr

allocation policy in the linear network, in which the priority is given to the short routes at the cost of the long one [1]. Verloop et al. considered various size-based bandwidth allocation policies in the linear network, such as SRPT, which gives the full priority to the flow with the shortest remaining service requirement, and LAS, in which the full priority is given to the flow with the least amount of attained service [7]. They found that these policies may cause instability effects even at arbitrarily low traffic loads.

Within the family of stable allocation policies (in the sense that the necessary conditions are also sufficient for stability), it is interesting to find out an optimal policy. In the single link case, all work-conserving disciplines are stable, and SRPT is known to be optimal in many ways, for example, by minimizing the mean delay. If only non-anticipating disciplines are allowed, LAS minimizes the mean delay if the flow size distribution belongs to class DHR. The optimal bandwidth allocation in the network case seems to be a tricky problem, which has been studied only in very few papers. Yang and de Veciana found that an optimal policy in a linear network with equal link capacities gives (at a time) the flows on the long route either the full capacity or nothing [9]. Their optimality criterion was to minimize the average bit transmission delay (also known as slow-down). However, their sample path argument is valid, as well, for the mean delay criterion. Verloop et al. determined the optimal non-anticipating policy in a linear network with equal link capacities and exponential flow sizes [8].

Thus, 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 between the traffic classes (as in [7]) but only locally within each traffic class. More precisely, we show that any stable state-dependent 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. The proof together with an extensive numerical study is given in the complete version of this paper. In this short (extended abstract) version we only present the main result together with a numerical example.

Let Π denote the family of those feasible bandwidth allocation policies for which the necessary conditions (2) are also sufficient for stability. Let $Z_r^\pi(t)$ denote the total bandwidth allocated to traffic class r at time t for any $\pi \in \Pi$. In addition, let Π° denote the subfamily of those bandwidth allocation policies that belong to Π and for which $Z_r^\pi(t)$ is given as follows:

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

This subfamily includes clearly all fair bandwidth allocation policies mentioned previously. Now let $\pi \in \Pi^\circ$. Denote by $\tilde{\pi}$ a modified policy for which the inter-route bandwidth allocation process is the same as for π ,

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

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. Note that π' differs from π^* , which also applies SRPT as the intra-route discipline but for which the inter-route bandwidth allocation policy is defined by

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

The main theorem tells that π' is the optimal modification in a very strong sense, minimizing the number of flows $N_r(t)$ in any class r at any time t in each sample path.

Theorem 1 *Let $\pi \in \Pi^\circ$, $r \in \mathcal{R}$ and $t \geq 0$. Then $N_r^{\pi'}(t) \leq N_r^{\tilde{\pi}}(t)$ for any modification $\tilde{\pi}$.*

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, of course, the mean delay as well by Little's result.

Corollary 2 *Let $\pi \in \Pi^\circ$. Then $\bar{N}^{\pi'} \leq \bar{N}^{\tilde{\pi}}$ for any modification $\tilde{\pi}$.*

However, if we want to compare policies π^* and π , a similar pathwise result as given in Theorem 1 is not valid in general. It is easy to construct a counter-example for that.

We made simulations in which the three policies (π , π' and π^*) mentioned above are compared. As an example, we considered a linear bandwidth sharing network consisting of two links with unit capacities. 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. We assumed Poisson arrivals for class r with intensity λ_r and, for all classes, the mean flow sizes of $\beta = 0.8$. Below we present the results of such simulations with hyperexponentially distributed flow sizes for two basic policies. In this numerical example, $\lambda_0 = 0.4$ and $\lambda_1 = \lambda_2 = 0.3$. For the proportionally fair basic policy π , the results were as follows:

$$\bar{N}^\pi = 2.07, \quad \bar{N}^{\pi'} = 1.58, \quad \bar{N}^{\pi^*} = 1.56.$$

For the basic policy π which is known to be the optimal non-anticipating discipline in the corresponding case with exponential flow sizes (giving full priority to the short routes if there are flows on all of them [8]), the results were as follows:

$$\bar{N}^\pi = 2.72, \quad \bar{N}^{\pi'} = 2.14, \quad \bar{N}^{\pi^*} = 2.14.$$

Three interesting notes can be made even from this single example. 1) The policy which is optimal for the exponential case is far from that when flow sizes are non-exponential. 2) If anticipating policies are allowed, the mean delay performance can be reduced significantly by applying SRPT in a controlled way. 3) Numerically the difference between the policies π' and π^* is very small.

Bibliography

- [1] T. Bonald and L. Massoulié. Impact of fairness on Internet performance. In: *ACM Sigmetrics 2001*, Cambridge, MA, pp. 82-91.
- [2] T. Bonald and A. Proutière. Insensitive bandwidth sharing in data networks. *Queueing Systems* **44** (2003) 69-100.
- [3] F.P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications* **8** (1997) 33-37.
- [4] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE Transactions on Networking* **8** (2000) 556-567.

- [5] J. Roberts and L. Massoulié. Bandwidth sharing and admission control for elastic traffic. In: *ITC Specialist Seminar*, Yokohama, Japan, 1998.
- [6] G. de Veciana, T.-J. Lee and T. Konstantopoulos. Stability and performance analysis of networks supporting services with rate control — Could the Internet be unstable? In: *IEEE Infocom 1999*, New York, NY, pp. 802-810.
- [7] I.M. Verloop, S.C. Borst and R. Núñez-Queija. Stability of size-based scheduling disciplines in resource-sharing networks. *Performance Evaluation* **62** (2005) 247-262.
- [8] I.M. Verloop, S.C. Borst and R. Núñez-Queija. Delay optimization in bandwidth-sharing networks. In: *CISS 2006*, Princeton, NJ.
- [9] S. Yang and G. de Veciana. Size-based adaptive bandwidth allocation: Optimizing the average QoS for elastic flows. In: *IEEE Infocom 2002*, New York, NY, pp. 657-666.
- [10] H.-Q. Ye. Stability of data networks under an optimization-based bandwidth allocation. *IEEE Transactions on Automatic Control* **48**, 7 (2003) 1238-1242.