

Short-Term Fairness and Long-Term QoS

Bo Tan

Department of Electrical
and Computer Engineering
University of Illinois at Urbana - Champaign
Urbana, IL 61801, USA
Email: botan2@uiuc.edu

Lei Ying

Department of Electrical
and Computer Engineering
Iowa State University
Ames, IA 50014, USA
Email: leiying@iastate.edu

R. Srikant

Department of Electrical
and Computer Engineering
University of Illinois at Urbana - Champaign
Urbana, IL 61801, USA
Email: rsrikant@uiuc.edu

Abstract—We consider connection-level models of resource allocation in the Internet, where files arrive into the network according to a Poisson process and the size of each file is exponentially distributed. The link capacities in the network are assumed to be shared among the files in the network using a proportionally-fair resource allocation scheme. In this paper, we are interested in the impact of this resource allocation on the file-transfer delay experienced by the users of the network. Towards this end, we derive a simple upper bound on the expected number of files in the network using Lyapunov techniques.

I. INTRODUCTION

Resource allocation algorithms for the Internet are designed to ensure fairness of allocation among the many users who are present in the network at any one instant of time. Such algorithms are often designed under the assumption that the number of users is fixed. We will refer to such a model as a *static model*. The study of distributed algorithms for static models was initiated in [6] and has been studied by a number of authors (see [10] for references). In reality, the users in the network arrive, bringing in a certain amount of work in the form of a file to be transferred, and depart when the work is completed, i.e., when the file has been transferred. The stability of the network when there are file arrivals and departures has been studied in a number of papers [9], [3], [1], [8]. Such models are often called *connection-level models*. In this paper, our goal is to go one step beyond stability, and study the impact of resource-allocation on the connection-level performance of the network under a connection-level model.

Specifically, we consider the proportionally-fair resource allocation policy introduced in [5]. The proportionally-fair resource allocation policy, which can be implemented in a distributed fashion using congestion control algorithm, has been widely-studied. It has been shown that the proportionally-fair scheme can

- (1) provide fair resource allocation at each time instant;
- (2) support the maximum connection-level throughput.

While this proportionally-fair policy guarantees short-term fairness, it does not maximize the departure rate at each time instant (an example will be presented in Section II). Thus, the connection-level delay performance may not be optimized under this proportionally-fair policy. Using the Foster-Lyapunov criterion, we derive an upper-bound on the expected number of files in the network under the proportionally-fair policy, and

the upper bound increases linearly with the number of flows in the network. This result complements the work of Kang et al [4] who have obtained a conjecture on the distribution on the number of files in the network in a heavy-traffic regime using Brownian motion models. Our upper bound matches the result in [4] in certain symmetric networks. It should be noted that the result in [4] is a generalization of earlier results on product-form solutions obtained by [9] and [2].

We would like to mention that the delay performance of proportional fairness has been compared to the performance of delay-optimal policies in a linear network in [12], [13] and in a symmetric star network in [14].

II. MODEL AND RESULTS

In this paper, we consider a network with L links and R routes. Each file transfer in the network will be associated with a route. Thus, there are R file classes in the network. We assume that files of each route arrive according to independent Poisson processes with rate λ_r files/sec, the file sizes are independent and exponentially distributed, with mean file size $1/\mu_r$ bits, and the capacity of each link is c_l bit/sec. Besides, we let $n_r(t)$ denote the number of files using route r in the network at time t , and $x_r(t)$ denote the rate allocated to serve files using r at time t . Additionally, we assume that the files using the same route equally share the allocated rate, i.e., at time t , a file on route r will be served at rate $x_r(t)/n_r(t)$ if $n_r(t) > 0$. We further define a routing matrix \mathbf{A} such that $A_{lr} = 1$ if route r uses link l , so so a feasible rate allocation \mathbf{x} has to lie in the set \mathcal{F} defined as

$$\mathcal{F} = \{\mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{c}, \mathbf{x} \geq 0\}.$$

The evolution of $\mathbf{n}(t) = (n_1(t), \dots, n_R(t))$ can be described by a continuous-Markov chain. Furthermore, using uniformization [7], we can construct a discrete-time Markov chain $\mathbf{n}[t]$ by sampling the network at certain random times as we will describe now [10], [11]. Let ν denote the maximum rate at which events occur in this network, where

$$\nu = \max_{\mathbf{x} \in \mathcal{F}} \left\{ \sum_{r=1}^R (\lambda_r + x_r \mu_r) \right\}.$$

We sample the system according to a Poisson process of this rate. Specifically, after each sample which we call an

event, the next event is an arrival of a file using route r with probability λ_r/ν , a departure of a file using route r with probability $x_r\mu_r/\nu$, or the system state is left unchanged with the remaining probability. These events are called arrivals, departures and fictitious departures, respectively. From now on, without loss of generality we assume $\nu = 1$.

For the discrete-time system obtained through uniformization, let $n_r[t]$ denote the number of files of type r during the time from the t^{th} event to $t + 1^{\text{th}}$ event, named as *time slot* t . We further let $x_r[t]$ denote the rate allocated to serve files at route r at time slot t . It is obvious that the network dynamics at the connection level depends on the resource allocation algorithm that decides $\mathbf{x}[t]$ at each time slot t . One particular resource allocation scheme that has been widely studied is the proportionally-fair resource allocation algorithm, which is defined as follows:

Proportionally-Fair (PF) Resource Allocation Algorithm: At time slot t , files on route r are served at rate $x_r[t]$, where $\mathbf{x}[t]$ is the solution of the following optimization problem:

$$\mathbf{x}[t] = \arg \max_{\mathbf{x} \in \mathcal{F}} \sum_r n_r[t] \log x_r. \quad (1)$$

The following result has been proved in [1], [10], [11].

Lemma 1: The network is stochastically stable under the proportionally-fair algorithm if

$$\sum_{r:l \in r} \frac{\lambda_r}{\mu_r} < c_l.$$

□

From Lemma 1, we can see that the proportionally-fair resource allocation algorithm supports the maximum connection-level throughput. However, the links could be occasionally under-utilized under the proportionally-fair resource allocation algorithm.

Consider a line network as in Figure 1, where we assume $\lambda_r = \lambda$, $\mu_r = \mu$ and $c_l = 1$. Given $n_1[t] = n_2[t] = n_3[t] \neq 0$, the proportionally-fair resource allocation yields $x_1[t] = x_2[t] = x_3[t] = 0.5$. So the overall departure rate is 1.5μ . However, if we choose $x_2[t] = x_3[t] = 1$, the departure rate becomes 2μ . Thus, it is quite possible that the number of users




Fig. 1. A Line Network

under proportional fairness can be larger than under an optimal scheme.

The example above motivates us to study the delay-performance of the proportionally-fair resource allocation. We use the approach in [10], [11] to study the following objective

$$\lim_{h \rightarrow \infty} \frac{1}{h+1} \sum_{t=0}^h E \left[\sum_{r=1}^R n_r[t] \mid \mathbf{n}[0] \right]. \quad (2)$$

For brevity, we will call this objective “the average number of files in the network” in the rest of the paper. Note that $\mathbf{n}[0]$ is the initial condition. Since the proportionally-fair controller is stable-in-the-mean [10], [11], it follows that the above objective is also equal to

$$\lim_{t \rightarrow \infty} E \left[\sum_{r=1}^R n_r[t] \mid \mathbf{n}[0] \right]$$

under proportional fairness.

Next we give an upper bound on the expected number of files in the network under the proportionally-fair algorithm.

Theorem 2: Define $\rho_r = \lambda_r/\mu_r$. We assume there exists $\epsilon > 0$ such that $(1 + \epsilon)\boldsymbol{\rho} \in \mathcal{F}$. Then under the proportionally-fair algorithm, we have

$$\lim_{t \rightarrow \infty} E \left[\sum_{r=1}^R n_r^{\text{PF}}[t] \mid \mathbf{n}[0] \right] \leq \frac{R}{\epsilon}, \quad \forall \mathbf{n}[0],$$

where the superscript *PF* indicates the proportionally-fair algorithm, and we recall R is the number of routes in the network.

Proof: Since we only consider the proportionally-fair algorithm in this proof, the superscript is ignored without causing any confusion. Define the Lyapunov function

$$W[t] = \sum_{r=1}^R \frac{1}{\lambda_r} n_r^2[t].$$

First, we note that

$$\begin{aligned} & E [W[t+1] - W[t] \mid \mathbf{n}[t]] \\ &= \left(\sum_{r=1}^R ((n_r[t] + 1)^2 - n_r[t]^2) \right. \\ & \quad \left. + \sum_{r=1}^R \frac{\mu_r}{\lambda_r} x_r[t] ((n_r[t] - 1)^2 - n_r[t]^2) \mathcal{I}_{n_r[t] > 0} \right), \end{aligned}$$

which implies that

$$\begin{aligned} & E [W[t+1] - W[t] \mid \mathbf{n}[0]] \\ &= E [E [W[t+1] - W[t] \mid \mathbf{n}[t]] \mid \mathbf{n}[0]] \\ &= R + \sum_{r=1}^R \frac{1}{\lambda_r} (E [2(\lambda_r - \mu_r x_r[t]) n_r[t] \mid \mathbf{n}[0]] + \\ & \quad E [\mu_r x_r[t] \mathcal{I}_{n_r[t] > 0} \mid \mathbf{n}[0]]). \end{aligned} \quad (3)$$

It has been proved in [10] that $\mathbf{n}[t]$ is positive-recurrent under the proportionally-fair algorithm given $\rho < 1$. Thus, we have

$$\begin{aligned} \lim_{t \rightarrow \infty} E [W[t+1] - W[t] \mid \mathbf{n}[0]] &= 0 \\ \lim_{t \rightarrow \infty} E [\mu_r x_r[t] \mathcal{I}_{n_r[t] > 0} \mid \mathbf{n}[0]] &= \lambda_r, \end{aligned}$$

where the second equality is simply a statement that the total departure rate is equal to the total arrival rate due to stability. Letting t in equality (3) goes to infinity, we can conclude that

$$\lim_{t \rightarrow \infty} E \left[\sum_{r=1}^R \frac{1}{\lambda_r} (\lambda_r - \mu x_r[t]) n_r[t] \mid \mathbf{n}[0] \right] = -R. \quad (4)$$

Recall that the resource allocation under the proportional-fairness is a solution to the following optimization problem:

$$\max_{\mathbf{x} \in \mathcal{F}} f(\mathbf{x}) = \sum_{r=1}^R n_r[t] \log x_r.$$

Since $f(\mathbf{x})$ is concave and \mathcal{F} is a convex set, we have

$$\nabla f \cdot (\mathbf{x} - \mathbf{x}[t]) \leq 0$$

where $\mathbf{x}[t]$ is the optimal solution given $\mathbf{n}[t]$, which implies that

$$\sum_{r=1}^R \frac{n_r[t]}{x_r} (x_r - x_r[t]) \leq 0. \quad (5)$$

Since $\{(1 + \epsilon)\lambda_r/\mu_r\} \in \mathcal{F}$, from inequality (5), we can get that

$$\sum_{r=1}^R n_r[t] \left(1 - \frac{\mu_r}{\lambda_r(1 + \epsilon)} x_r[t] \right) \leq 0,$$

which implies that

$$\sum_{r=1}^R \frac{n_r[t]}{\lambda_r} ((1 + \epsilon)\lambda_r - \mu_r x_r[t]) \leq 0,$$

which leads to

$$\sum_{r=1}^R \frac{1}{\lambda_r} (\lambda_r - \mu_r x_r[t]) n_r[t] \leq -\epsilon \sum_{r=1}^R n_r[t]. \quad (6)$$

Combining equation (4) and inequality (6), we finally conclude that

$$\lim_{t \rightarrow \infty} E \left[\sum_{r=1}^R n_r[t] \mid \mathbf{n}[0] \right] \leq \frac{R}{\epsilon}.$$

From Theorem 2, we can see that under the proportionally-fair resource allocation, the expected number of files linearly increases with the number of routes in the network. This result complements the work of Kang et al [4] who have conjectured the distribution on the number of files in the network in a heavy-traffic regime using Brownian motion models. Assuming there is a single-hop flow on each link (local traffic assumption), it has been shown that

$$E[n_r[\infty]] = \rho_r \sum_{l:l \in r} \frac{1}{c_l - \sum_{j:l \in j} \rho_j}, \quad (7)$$

where the subscript l indicates a link, r and j indicate routes, and, for simplicity, we use $n_r[\infty]$ to denote the steady-state value. Next, we examine the impact of the local traffic assumption on the accuracy of (7).

Now given a network with an arbitrary topology and arbitrary flows, we can use the result in Theorem 2 or equality (7)

to estimate the steady-state mean number of files. In networks with a few long routes, our bound could be tighter than (7). For example, consider a line network with L links and a single flow as in Figure 2, where we assume that $c_l = 1$ for all l and $\lambda/\mu(1 + \epsilon) = 1$. Our upper bound indicates

$$E[n[\infty]] \leq \frac{1}{\epsilon},$$

which is tight. However, letting the local traffic load go to zero on each link in (7) yields

$$E[n[\infty]] = \frac{L}{\epsilon}.$$




Fig. 2. A Line Network with L Links

Now we consider a star-network with L links as in Figure 3, where every pair of links forms a route. We assume that $c_l = 1$ for all l and $\rho_r(1 + \epsilon) = \lambda_r/\mu_r(1 + \epsilon) = 1/(L - 1)$. Since there are $L(L - 1)/2$ routes in the network, our upper bound shows that

$$E \left[\sum_{r=1}^{L(L-1)/2} n_r[\infty] \right] \leq \frac{L(L-1)}{2\epsilon}.$$

The result in [4] indicates that

$$E \left[\sum_{r=1}^{L(L-1)/2} n_r[\infty] \right] = \frac{L}{\epsilon},$$

which is tighter than our upper bound.




Fig. 3. A Star Network with L Links

Despite the above limitations, the expressions obtained in Theorem 2 and in [4] can be used to study the delay performance of proportionally-fair resource allocation. In [13], optimal policies for a line network with two links were proposed. The delay-performance of the optimal policies and the proportionally-fair policy were compared using simulations, and it was shown that the gap is less than 20%. We also

considered a symmetric star-network with three links and three flows in [14], where we identified the delay-optimal policy. Using the upper bound in Theorem 2, we showed that the expected number of files under the proportionally-fair policy is at most a factor of 1.5 more than the one under the optimal policy.

III. CONCLUSION

In this paper, we studied the long-term delay-performance of the proportionally-fair resource allocation. We obtained an upper bound on the expected number of files in the network, which linearly increases with the number of routes in the network. This upper bound is not tight in general, but complements the result in [4].

ACKNOWLEDGMENT

The research in the paper was supported by NSF Grants CNS 07-21286, CCF 06-34891 and DTRA grant HDTRA1-08-1-0016.

REFERENCES

- [1] T. Bonald and L. Massoulié. Impact of fairness on Internet performance. In *Proceedings of ACM Sigmetrics*, 2001.
- [2] T. Bonald and A. Proutiere. Insensitive bandwidth sharing in data networks. *Queueing Systems*, 44:69–100, 2003.
- [3] G. de Veciana, T.-J. Lee, and T. Konstantopoulos. Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Transactions on Networking*, 9(1):2–14, February 2001.
- [4] W. N. Kang, F. P. Kelly, N. Lee, and R. J. Williams. Product form stationary distributions for diffusion approximations to a flow level model operating under a proportional fair sharing policy. *Performance Evaluation Review*, 2007. To appear.
- [5] F. P. Kelly. Notes on effective bandwidths. in *Stochastic Networks: Theory and Applications*, (Editors F. P. Kelly, S. Zachary and I. B. Ziedins), pages 141–168, 1996.
- [6] F. P. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [7] P. R. Kumar and P. Varaiya. *Stochastic Systems, Estimation, Identification and Adaptive Control*. Prentice Hall, 1986.
- [8] X. Lin, N. Shroff, and R. Srikant. On the connection-level stability of congestion-controlled communication networks. *IEEE Transactions on Information Theory*, 2007. To appear.
- [9] J. Roberts and L. Massoulié. Bandwidth sharing and admission control for elastic traffic. In *Proc. ITC specialists seminar*, Yokohama, Japan, 1998.
- [10] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.
- [11] R. Srikant. Models and methods for analyzing Internet congestion control algorithms. In *Advances in Communication Control Networks in the series Lecture Notes in Control and Information Sciences (LCNCIS)*, C.T. Abdallah, J. Chiasson and S. Tarbouriech (eds.). Springer, 2004.
- [12] I. M. Verloop. Efficient flow scheduling in resource-sharing networks, 2005. Master Thesis, Department of Mathematics, Utrecht University.
- [13] I. M. Verloop, S. C. Borst, and R. Nunez-Queija. Delay optimization in bandwidth-sharing networks. In *Proceedings of the Conference on Information Sciences and Systems (CISS)*, Princeton University, 2006.
- [14] L. Ying, B. Tan, and R. Srikant. Don the delay optimality of proportional fairness. In *Proc. 2008 Information Theory and Applications Workshop (ITA)*, 2008.