

Stability of Data Networks Under An Optimization-Based Bandwidth Allocation ¹

Heng-Qing Ye ²

Abstract

It is known that a data network may not be stable at the connection level under some unfair bandwidth allocation policies, even when the normal offered load condition is satisfied, i.e., the average traffic load at each link is less than its capacity. In this paper, we show that, under the normal offered load condition, a data network is stable when the bandwidth of the network is allocated so as to maximize a class of general utility functions. Using the microscopic model proposed by Kelly [9] for a TCP congestion control algorithm, we argue that the bandwidth allocation in the network dominated by this algorithm can be modelled as our bandwidth allocation model, and hence that the network is stable under the normal offered load condition. This result may shed light on the stability issue of the Internet since the majority of its data traffic is dominated by the TCP.

Keywords: Data network, bandwidth allocation, TCP, stability, Lyapunov function.

Previous versions: July 2001, April 2002, October 2002

Current version: March 2003

1 Introduction

There is no doubt that the Internet has been one of the most exciting and revolutionary technological developments in the past decade. The information flows along the Internet are still increasing dramatically, and the traffic control of the information flows has been an important issue in both the academics and the telecommunication industry. Currently, the majority of the Internet traffic is dominated by various versions of TCP (the Transmission Control Protocol; see for example Jacobson [7]), and a lot of efforts have been placed on the study of the TCP congestion control. However, due to its complex nature, the behavior of the Internet traffic is not fully understood yet and remains an active research topic. A better understanding of the behavior of the Internet traffic will have great impact on the design and implementation of the Internet of next generation. In this study, we aim to provide some insights into the connection level behavior of networks supporting *elastic* traffic, which include the Internet.

It is known that a data network may not be stable under the normal offered load condition (the condition (2)), i.e., the long-term average bandwidth requirement (offered traffic load) at each link in the network is less than the bandwidth capacity of the link. Here the bandwidth of a link is the amount of data that can be transmitted through the link in unit time, and roughly

¹Supported in part by a grant from the Academic Research Fund and the Center for E-Business of the National University Singapore. Part of this work was done while the author was visiting the Statistical Laboratory, Cambridge University.

²The author is with the School of Business, National University of Singapore, 1 Business Link, Singapore 117591, email: bizyehq@nus.edu.sg.

the network is said to be stable (at the connection level) if the expected number of ongoing connections is finite. A simple counter-example for illustrating such unstability phenomenon is described in detail in Section 2. In the counter-example, the number of ongoing connections $x_{r_1}(t)$ on the long route r_1 approaches infinity as time $t \rightarrow \infty$, while the long-term average bandwidth requirements ρ_{l_1} and ρ_{l_2} for links l_1 and l_2 are less than the bandwidth capacities C_{l_1} and C_{l_2} of the two links respectively. A priority bandwidth allocation policy is employed in this example. This presents a situation that the bandwidth capacity of a network is sufficient for the data transmission requirement (i.e., the normal offered load condition $\rho_{l_1} < C_{l_1}$ and $\rho_{l_2} < C_{l_2}$) but is not fully utilized due to the unfair bandwidth allocation policy. Raised from the counter-example, a challenge is to answer the question related to the Internet, “is a network dominated by the TCP congestion control algorithm stable under the normal offered load condition?”

For this question, related results can be found in de Veciana, *et al.* [4] and Bonald and Massoulié [3]. In de Veciana, *et al.* [4], it is shown that networks employing a bandwidth allocation policy satisfying the (weighted) max-min fairness or the proportionally fairness criterion are stable under the normal offered load condition. Readers are referred to Bertsekas and Gallager [1] for detailed descriptions on the max-min fairness, and Kelly [8] and Kelly, *et al.* [10] on the proportionally fairness. In Bonald and Massoulié [3], a stability result is also established for networks employing a class of (p, α) -proportionally fair bandwidth allocation policies. This class of bandwidth allocation policies, first proposed by Mo and Walrand [14], include the bandwidth allocation policies satisfying the proportional fairness criterion and the minimal potential delay criterion as special cases. (We should brief these bandwidth allocation policies later.) These stability results strongly suggest that the network (such as the Internet) supporting elastic traffic would be stable under the normal offered load condition. In these papers, the data networks are modelled at the *connection (flow)* level, where a longer time scale is considered. Specifically, the queueing of packets at links and the transmission rate adaptation at sources (end users) are abstracted by assuming that the (equilibrium) bandwidth allocation adapts to the transition of ongoing connections immediately. Such an abstraction, often referred to as “separation of time scales”, is appealing when the time scale of the packet level rate control (e.g., the queueing time of packets at links and the propagation delay) is small compared with the time scale of the connection level dynamics (e.g., the transmission duration for a connection). Interested readers are referred to for example Hui [6] and de Veciana, *et al.* [4] for more details on the connection level models for data networks.

In this paper, we generalize the stability result in de Veciana, *et al.* [4] and Bonald and Massoulié [3] and show that the network under a bandwidth allocation that maximize a class of more general utility functions is stable under the normal offered load condition; see Section 2. By using a result in Kelly [9], which models the microscopic behavior of a TCP congestion control algorithm, we show that our bandwidth allocation model does capture some important characteristics of the macroscopic behavior of the network dominated by this TCP. Hence, it can be used to argue that such a network is stable under the normal offered load condition. This result may shed light on the stability issue of the Internet since the majority of its data traffic is dominated by the TCP; see Section 3. We conclude in Section 4.

Finally, we introduce some notation and convention that are used throughout this paper. The J -dimensional Euclidean space and its nonnegative orthant are denoted by \mathcal{R}^J and $\mathcal{R}_+^J := [0, \infty)^J$ respectively. Let $\mathcal{R} = \mathcal{R}^1$ and $\mathcal{R}_+ = \mathcal{R}_+^1$. Let $\mathcal{Z}_+ = \{0, 1, 2, \dots\}$ be the set of all nonnegative integers. Vectors are understood to be column vectors. The transpose of a vector

or a matrix is obtained by adding a T as a superscript to it. The notation e_r represents a vector with the r th element equal to 1 and all the other elements equal to 0. The dimension e_r can be easily deduced from the context. The operator $\partial_k f$ is a shorthand for the partial derivative of the first order differentiable function $f(\cdot)$ with respect to its k th variable, i.e.,

$$\partial_k f(x_1, \dots, x_k, \dots, x_n) = \frac{\partial f}{\partial x_k}(x_1, \dots, x_k, \dots, x_n).$$

2 A Markovian Connection Level Network Model Under An Optimization-Based Bandwidth Allocation

Consider a network with a set of links L . Each link $l \in L$ has a bandwidth capacity $C_l > 0$, i.e., the maximum amount of data that can be transmitted through the link l in unit time. Let R be the set of all possible routes, with each route r being a non-empty subset of L . On each route, connections are established and disconnected for data transmission dynamically. We assume that the interarrival times of connections (flows) to the route $r \in R$ are an i.i.d. sequence of exponentially distributed random variables with the mean arrival rate ν_r . We also assume that the amount of data (or the size of document) to be transmitted by a connection in the route r is exponentially distributed with the mean $1/\mu_r$. For and route $r \in R$, let $X_r(t) \in \mathcal{Z}_+$ be the number of ongoing connections in the route r at time $t \geq 0$. Let $\lambda_r(x) \in \mathcal{R}_+$ be the bandwidth allocated to each connection in the route r , and $\Lambda_r(x) := x_r \lambda_r(x)$ be the total bandwidth allocated to the route r , where x_r is the number of ongoing connections in the route r and $x = \{x_r, r \in R\}$. We call $X(t) = \{X_r(t), r \in R\}$ and $\Lambda(x) = \{\Lambda_r(x), r \in R\}$ the *ongoing connection process* and the (state-dependent) *bandwidth allocation policy* respectively. We also call $\Lambda(X(t))$ the *bandwidth allocation process*. A bandwidth allocation $\Lambda(\cdot)$ is feasible if conditions (4)-(6) are satisfied. For a network supporting elastic traffic, the bandwidth allocation policy $\Lambda(\cdot)$ is often (implicitly) imbedded in the network traffic control mechanism and can be modelled using an algorithm or an optimization problem (e.g., the proportionally fair, the minimal potential delay and the max-min fair allocation policies mentioned in the introduction section). Without loss of generality, we assume that each connection in the route r ($r \in R$) is allocated to a same bandwidth $\lambda_r(x)$ given the ongoing connection state x .

Given the (state-dependent, feasible) bandwidth allocation $\Lambda(\cdot)$, the ongoing connection process $X(t)$ is a continuous time Markov chain and its transition rates are

$$q(x, y) = \begin{cases} \nu_r & y = x + e_r, \\ \mu_r \Lambda_r(x) [= \mu_r x_r \lambda_r(x)] & y = x - e_r \text{ and } x_r \geq 1, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

for any ongoing connection states $x, y \in \mathcal{Z}_+^{|R|}$. The network model with the bandwidth allocation $\Lambda(\cdot)$ is said to be *stable* (at the connection level) if the ongoing connection process $X(t)$ is positive recurrent. An intuitive implication of the stability is that the number of ongoing connections will not blow up and has a finite expectation as time $t \rightarrow \infty$. Let $\rho_r = \nu_r / \mu_r$ be the offered data traffic load for the route $r \in R$. An obvious necessary condition for the positive recurrence of the Markov chain $X(t)$ is the *normal offered load condition*, i.e.,

$$\sum_{r:l \in r} \rho_r < C_l \quad \text{for } l \in L. \quad (2)$$

This condition implies that the *long-term average* bandwidth requirement at each link is within the link capacity.

However, the normal offered load condition is not sufficient for the continuous time Markov chain $X(t)$ to be positive recurrent if the bandwidth allocation $\Lambda(\cdot)$ is not chosen properly. This can be illustrated by the following counter-example, which is similar to an example in Bonald and Massoulié [3].

Example. Consider a network with two links $L = \{l_1, l_2\}$ and two routes $R = \{r_1, r_2\}$, where the route r_1 traverses both two links and the route r_2 traverses the link l_2 only; see Figure 1. Suppose that connections in the route r_2 have higher priority than those in the route r_1 when competing for the bandwidth of the shared link l_2 . That is, whenever there are connections in the route r_2 , all the bandwidth C_{l_2} of the link l_2 is allocated to the r_2 connections. Suppose that $C_{l_1} < C_{l_2}$. Then, the transition rates for this network read

$$q(x, y) = \begin{cases} \nu_r & y = x + e_r \text{ and } r \in R, \\ \mu_{r_2} C_{l_2} & y = x - e_{r_2} \text{ and } x_{r_2} \geq 1, \\ \mu_{r_1} C_{l_1} & y = x - e_{r_1} \text{ and } x_{r_2} = 0, x_{r_1} \geq 1, \\ 0 & \text{otherwise.} \end{cases}$$

By a careful thought, it is not difficult to see that the necessary and sufficient condition for the positive recurrence of the continuous time Markov chain $X(t) = (X_{r_1}(t), X_{r_2}(t))$ is

$$\rho_{r_1} < (1 - \frac{\rho_{r_2}}{C_{l_2}})C_{l_1} \text{ and } \rho_{r_1} + \rho_{r_2} < C_{l_2}.$$

This condition is strictly stronger than the normal offered load condition for this network, which is

$$\rho_{r_1} < C_{l_1} \text{ and } \rho_{r_1} + \rho_{r_2} < C_{l_2}.$$

□




Figure 1: A network with two links and two routes

Remarks. In this counter-example, the priority bandwidth allocation policy is employed at the link l_2 . Such priority phenomenon exists in the current Internet due to various reasons such as the followings. For example, some routes in the Internet differentiate the service among different classes of data flows by giving priority to specific classes of data flows; see, e.g., Blake, *et al.* [2]. □

In general, one would expect that the network be stable under the normal offered load condition. From the above counter-example, we have seen that the bandwidth capacity realization

problem of a network is non-trivial. It remains a challenge to investigate whether networks supporting elastic traffic such as the Internet are stable under the normal offered load condition. For this problem, we describe and investigate a bandwidth allocation model, which generalizes those in de Veciana, *et al.* [4] and Bonald and Massoulié [3]. In our model, the bandwidth allocation $\Lambda(x)$, for any ongoing connection state $x \in \mathcal{Z}_+^{|R|}$, is given as the solution to the optimization problem below.

$$\text{maximize} \quad \sum_{r \in R} U_r(x_r, \Lambda_r), \quad (3)$$

$$\text{subject to} \quad \sum_{r: l \in r} \Lambda_r \leq C_l \quad \text{for } l \in L, \quad (4)$$

$$\Lambda_r \geq 0 \quad \text{for } r \in R, \quad (5)$$

$$\Lambda_r = 0 \text{ if } x_r = 0, \quad \text{for } r \in R. \quad (6)$$

Here the utility function $U_r : \mathcal{R}_+ \times \mathcal{R}_+ \rightarrow \mathcal{R}_+$ ($r \in R$) is continuous and second order differentiable in $(0, \infty)^2$ satisfying the following assumptions:

$$U_r(x_r, 0) \equiv 0, \quad (7)$$

$$\partial_2 U_r(0, \Lambda_r) \equiv 0, \quad (8)$$

$$\partial_2 U_r(x_r, \Lambda_r) > 0 \text{ for any } x_r > 0, \quad (9)$$

$$\partial_1 \partial_2 U_r(x_r, \Lambda_r) \geq 0 \text{ for any } x_r, \Lambda_r \geq 0, \quad (10)$$

$$\partial_2 U_r(\cdot, \Lambda_r) \text{ is bounded for any } \Lambda_r > 0, \quad (11)$$

$$U_r(x_r, \cdot) \text{ is concave for any fix } x_r > 0. \quad (12)$$

The assumptions (7)-(9) are intuitively appealing. The assumption (10) may seem odd at the first sight. Consider the pairs (x_r^1, Λ_r) and (x_r^2, Λ_r) with $x_r^1 < x_r^2$, where more connections share the bandwidth Λ_r in the latter case. Then, $\partial_2 U_r(x_r^1, \Lambda_r) \leq \partial_2 U_r(x_r^2, \Lambda_r)$ (by the assumption (10)) merely says that increasing the bandwidth is more rewarded in the latter case. The assumption (11) is a technical condition. The concavity assumption (12) simply says that adding an additional bandwidth is more welcomed when the allocated bandwidth is smaller.

The following stability result shows that the network dominated by elastic sources characterized by the above utility is stable under the normal offered load condition.

Theorem 2.1 *Suppose that the normal offered load condition (2) is satisfied and the bandwidth allocation process $\Lambda(x)$ is determined by the allocation model (3)-(6). Then, the ongoing connection process $X(t)$ is positive recurrent.*

Proof. Define the *mean velocity vector* $\Delta x \in \mathcal{R}^R$ so that the equality

$$E[X(t) - X(0) | X(0) = x] = (\Delta x)t + o(t)$$

holds for any ongoing connection state x . Then, given the transition rate of the Markov chain $X(t)$ in (1), we have

$$(\Delta x)_r = \nu_r - \mu_r \Lambda_r(x), \quad (13)$$

where $\Lambda(x) = \{\Lambda_r(x), r \in R\}$ is the solution to the optimization problem (3)-(6). According to the drift criteria for the positive recurrence of a Markov chain in Foster [5] and Kingman [11], it is sufficient to show that there exists a differentiable function $V : \mathcal{R}_+^R \rightarrow \mathcal{R}_+$, satisfying

$$V(x) > 0 \text{ if and only if } x \neq 0, \quad (14)$$

$$V(x) \geq V(y) \text{ if } x \geq y \geq 0, \text{ and} \quad (15)$$

$$V'(\cdot) \text{ is bounded,} \quad (16)$$

such that, for some $\epsilon > 0$,

$$(\Delta x)^T V'(x) \leq -\epsilon \text{ for any } x \neq 0. \quad (17)$$

Here $V'(x)$ is the gradient with its i th component being $\partial V / \partial x_i$.

Fix x for the moment, and define $W(\eta) = \sum_{r \in R} U_r(x_r, \eta_r)$ for $\eta \in \mathcal{R}_+^{|R|}$. Since $\Lambda(x)$ is the solution to the optimization problem (3)-(6) and $U_r(x_r, \cdot)$ is concave for any fixed $x_r \geq 0$ (see the condition (12)), we have

$$(\eta - \Lambda(x))^T W'(\eta) \leq 0 \quad (18)$$

for any η satisfying constraints (4)-(6) with Λ replaced by η . Let $\eta^\delta = (\eta_r^\delta)$ with $\eta_r^\delta = \rho_r(1 + \delta)1_{\{x_r > 0\}}$. Due to the normal offered load condition (2), for a sufficiently small positive number δ , η^δ satisfies the constraints (4)-(6) with Λ replaced by η^δ . Hence, substituting η with η^δ in the inequality (18) and by the definition of W , we have

$$\sum_{r \in R} \left(\rho_r(1 + \delta)1_{\{x_r > 0\}} - \Lambda_r(x) \right) \cdot \partial_2 U_r(x_r, \rho_r(1 + \delta)1_{\{x_r > 0\}}) \leq 0.$$

Due to the assumption (8), this is equivalent to

$$\sum_{r \in R} (\rho_r(1 + \delta) - \Lambda_r(x)) \cdot \partial_2 U_r(x_r, \rho_r(1 + \delta)) \leq 0.$$

Hence, when $x \neq 0$, we have

$$\begin{aligned} \sum_{r \in R} (\rho_r - \Lambda_r) \cdot \partial_2 U_r(x_r, \rho_r(1 + \delta)) &\leq - \sum_{r \in R} \rho_r \cdot \delta \cdot \partial_2 U_r(x_r, \rho_r(1 + \delta)) \\ &\leq - \sum_{r \in R} \rho_r \cdot \delta \cdot \partial_2 U_r(1_{\{x_r > 0\}}, \rho_r(1 + \delta)) \leq - \min_{r \in R} \rho_r \cdot \delta \cdot \partial_2 U_r(1, \rho_r(1 + \delta)) \leq -\epsilon \end{aligned} \quad (19)$$

for some positive number ϵ . The last inequality above is due to the condition (10), and the third inequality is due to (9).

Now, define (the candidate Lyapunov function) $V(x)$ as

$$V(x) = \sum_{r \in R} \int_0^{x_r} \mu_r^{-1} \partial_2 U_r(y, \rho_r(1 + \delta)) dy.$$

It is direct to verify that $V(x)$ satisfies conditions (14) and (15). The condition (16) is satisfied due to the assumption (11). To verify (17), note that

$$\begin{aligned} (\Delta x)^T (\partial V(x)) &= \sum_{r \in R} (\nu_r - \mu_r \Lambda_r(x)) \cdot \mu_r^{-1} \partial_2 U_r(x_r, \rho_r(1 + \delta)) \\ &= \sum_{r \in R} (\rho_r - \Lambda_r(x)) \cdot \partial_2 U_r(x_r, \rho_r(1 + \delta)) \leq -\epsilon, \end{aligned}$$

where the first equality is due to the equality (13) and the last inequality is due to the inequality (19). \square

Remarks. In Bonald and Massoulié [3], the utility function,

$$U_r(x_r, \Lambda_r) = w_r x_r^\alpha \Lambda_r^{1-\alpha} / (1 - \alpha) \quad \text{for } r \in R,$$

is considered for the bandwidth allocation (3)-(6). Here the real number $\alpha > 0$ is a parameter, and w_r ($r \in R$) are some pre-determined weights related to the route r . When the parameter α varies at $\alpha = 0$, $\alpha \rightarrow 1$, $\alpha = 2$ or $\alpha \rightarrow \infty$, the utility function corresponds to a maximum throughput, the proportional fairness, the minimal potential delay or the max-min fair bandwidth allocation criterion, respectively; see Bonald and Massoulié [3] for more details. A stability result for the network with this class of utility functions is also presented in their paper. Their stability result, as well the earlier one in de Veciana *et al.* [4], is a significant step toward a better understanding of the stability issue of the Internet, noting that any utility function is only an approximation to the precise behavior of TCP. Moreover, it can be seen from our result that their result is robust to certain changes in the utility function originally used in their paper. \square

3 Is the Internet stable?

The Internet is a system of interconnected computer networks, which has spanned all over the world. The majority of data traffic on the Internet is controlled by various versions of TCP, wherein end user systems adapt their data transmission rate in response to the delay and loss of packets. In this section, we argue heuristically that the network dominated by a TCP congestion control algorithm is stable under the normal offered load condition. This stability result may shed new light on the connection level stability issue of the Internet, though we should realize that it is impossible to provide a complete and satisfactory answer to such an issue in this study.

In Kelly [9], the bandwidth allocation for TCP is modelled at the microscopic level as follows. [As pointed out in a contemporary work by Low [12], this model approximates the bandwidth allocation of a network dominated by (various versions of) TCP Reno.] For any given number x_r of ongoing connections at each route, the bandwidth allocation Λ_r is chosen so as to maximize the utility function

$$U(x, \Lambda) = \sum_{r \in R} \frac{\sqrt{2}x_r}{T_r} \arctan\left(\frac{\Lambda_r T_r}{\sqrt{2}x_r}\right) - \sum_{l \in L} \int_0^{\sum_{r:l \in r} \Lambda_r} p_l(y) dy, \quad (20)$$

where T_r is the round trip time of route $r \in R$, and $p_l(y)$ is the probability that a packet incurs a congestion signal at the link l when the bandwidth y has been allocated to the routes that traverse the link. In this formulation, the FIFO scheduling policy is implicitly assumed for the packet transmission at routers, and therefore all connections (on either the same route or some different routes) that go through a link would experience the same congestion probability at the link. The function $U(x, \cdot)$ is strictly concave for any $x > 0$, and hence, there exists a unique bandwidth allocation Λ satisfying

$$\frac{\partial U(x, \Lambda)}{\partial \Lambda_r} = \frac{1}{1 + \frac{\Lambda_r^2 T_r^2}{2x_r^2}} - \sum_{j \in r} p_j \left(\sum_{s:j \in s} \Lambda_s \right) = 0. \quad (21)$$

For a well designed network, the congestion probability $p_l(\sum_{s:l \in s} \Lambda_s)$ would be close to 0 when the total bandwidth allocated to routes that traverse the link l , $\sum_{s:l \in s} \Lambda_s$, is less than the capacity C_l . However, when $\sum_{s:l \in s} \Lambda_s$ approaches the capacity C_l , the congestion probability $p_l(\sum_{s:l \in s} \Lambda_s)$ would approach 1 quickly. Noting that the first term in the equation (21), $(1 + \frac{\Lambda_r T_r^2}{2x_r^2})^{-1}$, is always less than 1, we can regard the function (with respect to Λ)

$$\sum_{j \in r} p_j \left(\sum_{s: j \in s} \Lambda_s \right)$$

in the equation (21), or the function

$$\sum_{l \in L} \int_0^{\sum_{r: l \in r} \Lambda_r} p_l(y) dy$$

in the utility function (20), as a barrier function; see Luenberger [13]. Thus, the solution to the equation (21) (or the maximization problem (20)) can be approximated by the solution to the following constrained optimization problem:

$$\begin{aligned} & \text{maximize} && \sum_{r \in R} \frac{\sqrt{2}x_r}{T_r} \arctan \left(\frac{\Lambda_r T_r}{\sqrt{2}x_r} \right), \\ & \text{subject to} && \text{constraints (4) - (6),} \end{aligned}$$

where the vector x is a given ongoing connection state. This is exactly an instance of the bandwidth allocation model discussed in Section 2. According to Theorem 2.1, we claim that the network dominated by the TCP (with the bandwidth allocation characterized as the above optimization problem) is stable under the normal offered load condition.

4 Concluding Remarks

It was believed that the Internet would be stable (at the connection level) under the normal offered load condition. However, it is difficult, if not impossible, to obtain sufficiently supportive evidences from the current Internet. Hence, theoretical results to support or deny this belief are preferable. In this paper, we first propose a Markovian network model with the bandwidth allocation specified by an optimization problem. This is an idealized connection level, or macroscopic level model. Next, we prove its stability under the normal offered load condition. The stability result is consistent with the common expectation that a long-term average throughput with the bandwidth requirement being within the link capacity should be admissible. Finally, this stability result, combined with a recent result in the microscopic modelling of a TCP congestion control algorithm in Kelly [9], is used to heuristically show the stability of the network dominated by the TCP congestion control algorithm under the normal offered load condition.

An assumption of the bandwidth allocation in this paper is that the *instant* traffic load at each link can not exceed the link capacity; see the constraint (4). Concerning the practical operation of data networks supporting elastic traffic (including in particular the Internet), one may question whether such assumption would be realistic. For the Internet, when some links

are carrying multiple TCP connections, it is possible for there to be a steady-state packet drop rate on those links, and for the (instant) arrival rate at those links to equal or slightly exceed the link capacity. Future work will include investigating whether the results in this paper provide a good approximation to this environment.

It is known that the exponential assumption on the document size of a connection is often violated in the real network. Can we relax this assumption in our model so that the stability result would be more robust? We believe that the relaxation would be a significant but challenging step toward a better understanding of the network dynamics. Under the exponential assumption, the network can be modelled as a continuous time Markov chain for which analytical tools are available. To extend the model to allow more general document size assumption, it is necessary to keep track of remaining untransmitted document sizes on all connections in order to capture the network dynamics. For this purpose, more sophisticated stochastic model is required and studying the stability for the network model with exponential document size assumption would be helpful.

Acknowledgement. I would like to thank Frank P. Kelly who interested me to this research and shared his insights.

References

- [1] Bertsekas, D. and Gallager, R. (1992). *Data Networks*, Prentice Hall.
- [2] Blake S., D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss (1998). An Architecture for Differentiated Services. RFC 2475, Internet Engineering Task Force, December 1998. <ftp://ftp.ietf.org/rfc/rfc2475.txt>.
- [3] Bonald T. and Massoulié L. (2001). Impact of Fairness on Internet Performance. *Proc. ACM SIGMETRICS 2001*, Boston, MA, June 2001.
- [4] de Veciana G., Lee T.J. and Konstantopoulos T. (2001). Stability and Performance Analysis of Networks Supporting Elastic Services. *IEEE/ACM Transactions on Networking*, **9**, 2-14. Earlier version appeared in *the Proceedings of 18th IEEE INFOCOM'99*, New York, March 1999.
- [5] Foster F.G. (1953). On the stochastic matrices associated with certain queueing processes. *Ann. Math. Statist.*, **24**, 355-360.
- [6] Hui J. Y. (1988). Resource Allocation for Broadband Networks. *IEEE Journal on Selected Areas in Communications*, Vol. 6, No. 9, 1598-1608.
- [7] Jacobson V. (1988). Congestion Avoidance and Control. *Proceedings of the ACM SIGCOMM '88 Conference*, 314-329.
- [8] Kelly, F.P. (1997). Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, **29**, 1009-1016.

- [9] Kelly, F.P. (2001). Mathematical modeling of the Internet, in B. Engquist and W. Schmid (ed.), *Mathematics Unlimited - 2001 and Beyond*, 685-702, Springer-Verlag, Berlin.
- [10] Kelly, F.P., Maulloo, A. and Tan, D. (1998). Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, **49**, 237-252.
- [11] Kingman, J.F.C. (1961). The ergodic behaviour of random walks. *Biometrika*, **48**, 391-396.
- [12] Low, S. (2002). A Duality Model of TCP and Queue Management Algorithms. Working paper.
- [13] Luenberger, D.G. (1984). *Linear and Nonlinear Programming*, 369-371. Addison-Wesley Publishing Company, Reading, Massachusetts.
- [14] Mo J. and Walrand J. (1998). Fair End-to-End Window-based Congestion Control. *Proceedings of SPIE '98 International Symposium on Voice, Video and Data Communications*. available at <http://www.path.berkeley.edu/~jhmo>.