# Stability and Performance Analysis of Networks Supporting Elastic Services

Gustavo de Veciana, *Member, IEEE*, Tae-Jin Lee, *Member, IEEE*, and Takis Konstantopoulos, *Member, IEEE*

***Abstract*—We consider the stability and performance of a model for networks supporting services that adapt their transmission to the available bandwidth. Not unlike real networks, in our model, connection arrivals are stochastic, each has a random amount of data to send, and the number of ongoing connections in the system changes over time. Consequently, the bandwidth allocated to, or throughput achieved by, a given connection may change during its lifetime as feedback control mechanisms react to network loads. Ideally, if there were a fixed number of ongoing connections, such feedback mechanisms would reach an equilibrium bandwidth allocation typically characterized in terms of its "fairness" to users, e.g., max–min or proportionally fair. In this paper we prove the stability of such networks when the offered load on each link does not exceed its capacity. We use simulation to investigate performance, in terms of average connection delays, for various fairness criteria. Finally, we pose an architectural problem in TCP/IPs decoupling of the transport and network layer from the point of view of guaranteeing connection-level stability, which we claim may explain congestion phenomena on the Internet.**

***Index Terms*—ABR service, bandwidth allocation, Lyapunov functions, performance analysis, proportional fairness, rate control, stability, TCP/IP, weighted max–min fairness.**

## I. Introduction

FUTURE communication networks are likely to increasingly support *elastic* or rate adaptive applications that permit varying the data transmission rate to match the available network bandwidth while achieving a graceful degradation in the perceived quality of service [28]. Transport services compatible with such applications are already supported on the Internet. Indeed, TCP is based on end-systems adjusting their transmissions in response to delayed or lost packets, which can be an implicit indicator of available bandwidth [14]. Similarly, available bit rate (ABR) service, defined for ATM networks, draws on both the end-systems and network elements to implement such functionality through adaptive rate control mechanisms that strive to allocate the available bandwidth among ongoing connections [5]. Typically such mechanisms represent an efficient way to carry traffic corresponding to elastic applications, e.g., today's file transfers and future highly adaptive voice/video applications.

Since mechanisms to adapt transmission rate typically draw on delayed (implicit or explicit) feedback from the network, much work has been devoted to establishing their stability. These results have usually been developed for networks supporting a *fixed* number of connections. Stability, in this context, is usually interpreted as avoiding queue/delay buildups, and/or somewhat loosely as ensuring that transmission rates converge to an equilibrium corresponding to a bandwidth allocation among ongoing connections [2], [3], [6], [29], [19], [1], [16]. An equilibrium bandwidth allocation is usually characterized in terms of its "fairness" to users, e.g., max–min or proportional fairness [4], [15]. Thus given a fixed number of users and fixed network capacities, one can typically arrange (through an appropriate control mechanism) to achieve an equilibrium which represents, according to some criterion, an equitable allocation of resources among users [24], [17], [20].

By contrast very little is known about network stability and performance when the number of ongoing connections is in constant flux. Previous work along these lines has focused on studying transients, i.e., how quickly will the transmission rates reach a new equilibrium. In this paper we consider a novel model that includes stochastic arrivals and departures of elastic flows/connections. We abstract the queueing and rate adaptation that would be taking place in the network by assuming that an equilibrium, and thus appropriate bandwidth allocation is immediately achieved upon a change in the number of ongoing connections. Thus, in essence, we assume a *separation of time scales* between the time scale of connection arrivals and departures and that on which rate control processes converge to equilibria.

Paralleling models used in the circuit switched literature, we assume connection arrivals processes are Poisson and that each connection has a random, exponentially distributed, amount of data to send.[1] In contrast to circuit switched models, the bandwidth allocated to each user is a function of the global state of the network. Indeed, recall that the bandwidth allocated to a user depends on the equilibrium achieved by the rate control mechanisms and the number of ongoing connections. Our goal in this paper is to determine when this network model is stable and compare connection-level performance for networks using different types of rate control and thus operated under different fairness policies.

In general, one expects work conserving systems to be stable when the offered load to each link (queue) in the network

Manuscript received June 29, 1999; revised May 15, 2000; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Tripathi. This work was supported by National Science Foundation Career Award NCR 96-24230, by National Science Foundation under Grant ANI 99-03495, by the Texas Higher Education Coordinating Board under Grant ARP-224, by Southwestern Bell Company Technology Resources Inc., and by an Intel Technology for Education 2000 equipment grant. An earlier version of this work was published in the *Proceedings of IEEE INFOCOM*, New York, NY, 1999.

G. de Veciana and T. Konstantopoulos are with the Department of Electrical and Computer Engineering, University of Texas at Austin, Austin, TX 78712 USA (e-mail: gustavo@ece.utexas.edu; takis@ece.utexas.edu).

T.-J. Lee was with the Department of Electrical and Computer Engineering, University of Texas at Austin, Austin, TX 78712 USA. He is now with the Network Solution Lab, Corporate R&D Center, Samsung Electronics Company, Ltd. Suwon City, South Korea. (e-mail: tjlee@samsung.com).

Publisher Item Identifier S 1063-6692(01)01311-5.

[1]This arrival model is reasonable for connections generated by a large population of independent users. The exponential assumption simplifies our analysis but is likely not to be critical for the stability results in this paper.

does not exceed its capacity. However, given the complex network-wide interactions underlying the bandwidth allocation mechanism, a demonstration of this fact was deemed important. Indeed, this model can be said to be "nonwork conserving" in the sense that a link supporting active connections may not be operating at a full utilization because its connections are "bottlenecked" elsewhere—a typical sign of a potential for instability. In this paper we come to terms with this problem by showing the stability of our model when natural conditions are satisfied.

Since ours is a higher layer model, it is logical to consider network-level performance, say in terms of average connection delays. This is important because the goals of fairness and low connection delays may not be compatible, and should be examined prior to committing to a particular architecture for large-scale broadband networks. Moreover, network designers might want to dimension capacities to achieve a reasonable responsiveness, say for web browsing, when the network is subject to typical loads. Our preliminary simulations suggest that indeed it may be of interest to examine more carefully the impact of a given fairness criterion and topology on the overall network performance.

Consideration of the stability of this system model points to an insidious architectural problem in networks supporting adaptive services of this type. To achieve connection layer stability we must ensure that connection-level loads do not exceed link capacities. This in turn requires that the routing layer be aware of the connection-level offered loads. However, today's routing algorithms, if at all, draw on link averages of utilization and/or packet delays. Such metrics reflect the connection-level offered loads quite poorly, since connections are adapting their transmission rates depending on link congestion. Loosely speaking, the router is indifferent to the fact that a 90% link utilization may be due to a single traffic source or a thousand sources transmitting at a thousandth of the latter's rate. Herein lies yet another possible explanation for the types of congestion currently experienced on the Internet, i.e., connection-level instability.

To our knowledge only the work in [27] has attempted to tackle this type of system model. Their work provides an explicit analysis of the performance of linear networks under the proportionally fair bandwidth allocation, as well as limited performance comparisons via simulation. The authors argue for the need to perform light-weight call admission if performance targets are to be met. The main contributions of our work is a rigorous analysis of stability for a general network model under several bandwidth allocation criteria, and additional simulations investigating additional performance characteristics for linear networks. This paper is an extended version of [8] including some additional results for proportional fairness, discussions, and simulations.

The paper is organized as follows. In Section II, we present our model and define the max–min, weighted max–min and (weighted) proportionally fair bandwidth allocations. Next, in Section III we show the stability of the model by constructing appropriate Lyapunov functions. Performance issues are discussed in Section IV. In Section V we return to our question concerning possible connection-level instabilities in current networks and discuss future work.

TABLE I
SUMMARY OF KEY PARAMETERS FOR SYSTEM MODEL

| | |
|---|---|
| $c_\ell$ | capacity of link $\ell \in \mathcal{L}$ (bits/sec) |
| $\lambda_r$ | arrival rate of connections on route $r \in \mathcal{R}$ (connections/sec) |
| $v_r$ | mean volume for connections on route $r \in \mathcal{R}$ (bits) |
| $\rho_r$ | offered load $\rho_r = \lambda_r v_r$ to route $r \in \mathcal{R}$ (bits/sec) |
| $M$ | route link incidence 0-1 matrix |
| $n_r$ | state of route $r$ (connections) |
| $n$ | network state vector $n = (n_r, r \in \mathcal{R})$ (connections) |
| $\mu_r(n)$ | bandwidth allocated to route $r$ (connections/sec) |
| $\mu(n)$ | bandwidth allocation vector $\mu(n) = (\mu_r(n), r \in \mathcal{R})$ (connections/sec) |

## II. NETWORK MODEL AND BANDWIDTH ALLOCATION SCHEMES

Table I summarizes the key parameters for our system model. Our network model consists of a set of links $\mathcal{L}$ with fixed capacities $c = (c_\ell, \ell \in \mathcal{L})$ in bits/s shared by a collection of routes $\mathcal{R}$. Routes are undirected and may traverse several links in the network. A 0-1 matrix $M = (M_{\ell r}, \ell \in \mathcal{L}, r \in \mathcal{R})$ indicates which links a route traverses. In other words, $M_{\ell r} = 1$ if route $r$ uses link $\ell$ and zero otherwise.

The dynamics of the model are as follows. New connections are initiated on route $r \in \mathcal{R}$ at random times forming a Poisson process $A_r$ with arrival rate $\lambda_r$ connections/s. The collection of processes $A = \{A_r, r \in \mathcal{R}\}$, with rates $\lambda = (\lambda_r, r \in \mathcal{R})$ are assumed to be independent. Each connection on route $r \in \mathcal{R}$ has a volume of data (in bits) to transmit, which is assumed to be an exponentially distributed random variable with mean $v_r$ bits. We let $v = (v_r, r \in \mathcal{R})$. The random variables representing connection volumes are thus i.i.d. and assumed to be independent of $A$. We let

$$\rho_r = \lambda_r v_r$$

denote the offered load along route $r$, expressed in bits/s, and let $\rho = (\rho_r, r \in \mathcal{R})$.

The "state" of the network is denoted by $n = (n_r, r \in \mathcal{R})$ where $n_r$ denotes the number of ongoing connections on route $r$. We assume that the bandwidth allocated to each ongoing connection depends only on the current state $n$ of the system. Let $\mu_r(n)$ denote the total bandwidth allocated to connections on route $r$ when the system state is $n$, expressed as a service rate in connections/s. As explained in the sequel $\mu_r(n)$ is given by the bandwidth in bits/s allocated to route $r$ divided by the mean volume $v_r$ for connections on this route. The choice of the functions $\mu = (\mu_r\colon \mathbb{Z}_+^{\mathcal{R}} \to \mathbb{R}_+, r \in \mathcal{R})$ will be described in the sequel. If the state of the system changes during the sojourn of a connection (e.g., due to the establishment of a new connection or the termination of an existing one), then there may be a corresponding change (speed-up or slow-down) in its service rate. Indeed, since no arriving connections are blocked, new connections must be accommodated by changing the bandwidth allocation and when bandwidth becomes available due to departing connections it is reallocated to the remaining ones. We assume that ongoing connections are *greedy* in the sense that they will use whatever network bandwidth is made available to them. Note that in reality a given connection may have a limit on the rate at which it can transmit, e.g., may be limited by the access network or network interface card. Herein we shall assume that such bottlenecks have been explicitly modeled by incorporating limited capacity access links in the network.

The overall evolution of the system can be defined as follows. Recall that $A_r(t)$ denotes the number of connections arriving on route $r$ on the time interval $(0, t]$—a rate-$\lambda_r$ Poisson counting process. We let $D_r(t)$ be another independent *unit rate* Poisson process associated with departures from route $r$. Letting $N(t) = (N_r(t), r \in r \in \mathcal{R})$ denote the random process capturing the number of connections on the routes, for $r \in \mathcal{R}$ and $t \geq 0$, we have

$$N_r(t) = N_r(0) + A_r(t) - D_r\left(\int_0^t \mu_r(N(s))\, ds\right) \quad (1)$$

i.e., arrivals minus departures. Note that the state dependent service rates along each route in the network are captured by rescaling the time axis for the departure process, i.e., speeding it up or slowing it down depending on the bandwidth $\mu_r(N(s))$ allocated to that route. It should be clear that given an initial state $N(0)$, this evolution equation has a unique solution. Moreover, if the initial condition $N(0)$ is selected independently of the arrivals and service processes then the $\mathbb{Z}_+^{\mathcal{R}}$-valued process $N = (N(t), t \geq 0)$ is Markovian.

In the sequel, we describe various bandwidth allocation schemes, or, equivalently, various possible functions $\mu$. In particular we will use $\mu^m$, $\mu^w$, and $\mu^p$ to denote the max–min, weighted max–min, and weighted proportionally fair bandwidth allocation functions, respectively. As will be seen, these are functions, of the state $n$, the capacity vector $c$, the mean connection sizes $v$, the routing matrix $M$, and the type of rate control used on the network. In contrast to standard queueing models, which track packets and queues throughout the network, it is through the dependence of the allocated bandwidth on the number of ongoing connections that the evolution (1) captures the dynamics of the system. Also note that we have assumed that connections are not rerouted once they are initiated. One could in principle account for rerouting or splitting of flows across the network but this will not be considered here. Finally, as explained below, we reiterate that the allocated bandwidth $\mu_r$ will be measured in units of connections/s rather than bits/s.

### A. Max–Min Fair Bandwidth Allocation

We first consider max–min fair bandwidth allocation. An allocation is said to be max–min fair if the bandwidth allocated to a connection cannot be increased without also decreasing that of a connection having a smaller or equal allocation [4]. For a single link network this translates to giving each connection traversing the link the same amount of bandwidth. In general one first determines what would be the maximum minimum bandwidth one could assign to any connection in the network and allocates it to the most poorly treated connections. One then removes these connections and the allocated bandwidths from the network, and iteratively repeats the process of maximizing the minimum bandwidth allocation for the remaining connections. More formally the max–min fair allocation can be defined in terms of a hierarchy of optimization problems, described in detail in [12], which are easily solved via the above procedure. Below we briefly review how given the state $n$ of the network one determines the max–min fair bandwidth allocations per connection and in turn determines the bandwidth allocations $(\mu_r^m(n), r \in \mathcal{R})$ per route.

Let the vector $a^* = (a_r^*, r \in \mathcal{R})$ be the max–min fair allocation where $a_r^*$ denotes the bandwidth, in bits/s, allocated to a single connection on route $r$. Notice that we have suppressed the dependence of $a^*$ on $n$. Since all connections on the same route get the same allocation so $\mu_r^m(n) = v_r^{-1} n_r a_r^*$ now measured in connections/s. We determine $a^*$ as follows. First for all routes $r \in \mathcal{R}$ such that $n_r = 0$ we set $a_r^* = 0$ and thus $\mu_r^m(n) = 0$. Next we solve a hierarchy of optimization problems starting with

$$f^{(1)}(n) := \max_a \left\{ \min_{r\in\mathcal{R}} a_r \colon \sum_{r\in\mathcal{R}} M_{\ell r} n_r a_r \leq c_\ell,\ \ell \in \mathcal{L} \right\} \quad (2)$$

which corresponds to maximizing the minimum bandwidth per connection subject to the link capacity constraints. It can be shown, see [12], that the solution to this problem is given by

$$f^{(1)}(n) = \min_{\ell\in\mathcal{L}} f_\ell^{(1)}(n) \quad \text{with} \quad f_\ell^{(1)}(n) := \frac{c_\ell}{\sum_{r\in\mathcal{R}} M_{\ell r} n_r} \quad (3)$$

where $f_\ell^{(1)}(n)$ can be thought of as the *fair share* at link $\ell$, i.e., the bandwidth per connection at link $\ell$ if its capacity were equally divided among the connections traversing the link.

Let $\mathcal{L}^{(1)}$ be the set of links $\ell$ such that $f_\ell^{(1)}(n) = f^{(1)}(n)$. This is the set of first-level *bottleneck links*. The set of first-level *bottleneck routes* $\mathcal{R}^{(1)}$ is the set of routes traversing a link in $\mathcal{L}^{(1)}$. These two sets make up the first-level of the *bottleneck hierarchy*. Finally, for each route $r \in \mathcal{R}^{(1)}$, let $a_r^* = f^{(1)}(n)$. The remaining, if any, components of $a^*$ are determined by repeating this process on a reduced network as explained next.

In the second step, if it arises, the algorithm replaces the sets $\mathcal{L}$ and $\mathcal{R}$ by $\mathcal{L} \setminus \mathcal{L}^{(1)}$ and $\mathcal{R} \setminus \mathcal{R}^{(1)}$, respectively where, $A \setminus B$ is the difference of between sets $A$ and $B$. The new state of the system is simply the projection $(n_r, r \in \mathcal{R} \setminus \mathcal{R}^{(1)})$, and a new link capacity vector, $c^{(1)}$ is defined on $\mathcal{L} \setminus \mathcal{L}^{(1)}$, where $c_\ell$ is reduced to

$$c_\ell^{(1)} = c_\ell - \sum_{r\in\mathcal{R}^{(1)}} M_{\ell r} n_r f^{(1)}(n) = c_\ell - f^{(1)}(n) \sum_{r\in\mathcal{R}^{(1)}} M_{\ell r} n_r.$$

From (2) and the definition of $\mathcal{L}^{(1)}$ it is clear that the reduced capacities are nonnegative. A new problem paralleling (2) but on the reduced network (with reduced sets or routes and links, reduced state, and reduced capacities—as described above) is then defined and solved to obtain a new value $f^{(2)}(n)$, and second-level bottleneck sets $\mathcal{L}^{(2)}$ and $\mathcal{R}^{(2)}$. Finally, for $r \in \mathcal{R}^{(2)}$ we set $a_r^* = f^{(2)}(n)$. If necessary, this process is once again repeated, but, since the sets $\mathcal{R}^{(1)}, \mathcal{R}^{(2)}, \ldots$ are nonempty, it terminates in a finite number of steps, uniquely specifying the vector $a^*$ and thus $\mu^m(n)$.

Notice that, in the above procedure, $n$ need not be integer valued, hence $\mu^m(n)$ can be easily extended for real-valued arguments. We shall use the same notation to denote the extension of $\mu^m$ from $\mathbb{Z}_+^{\mathcal{R}}$ to $\mathbb{R}_+^{\mathcal{R}}$. Some straightforward properties of this function are summarized below.

*Proposition 2.1:* The function $\mu^m: \mathbb{R}_+^{\mathcal{R}} \to \mathbb{R}_+^{\mathcal{R}}$ is radially homogeneous, in the sense that

$$\mu^m(\alpha x) = \mu^m(x), \qquad x \in \mathbb{R}_+^{\mathcal{R}}, \qquad \alpha > 0.$$

In the interior of the positive orthant $\mathbb{R}_+^{\mathcal{R}}$, the function $\mu^m$ is continuous, and has strictly positive components. Finally, $\mu^m$ is bounded.

The proof of this proposition can be shown by induction on the bottleneck hierarchy, considering the dependence on $x$ of the max–min fair bandwidth allocation.

Notice that although the bandwidth allocation policy reflected in $\mu^m$ is max–min fair this may not lead to an "optimal" network performance, say in terms of connection delays. In general to address this issue one might want to give different priorities to connections based on their routes, whence, in the next section we briefly discuss weighted max–min fair bandwidth allocations.

### B. Weighted Max–Min Fair Bandwidth Allocation

Let $w = (w_r, r \in \mathcal{R})$ be positive "weights" associated with each route in the network, and $a^{w*} = (a_r^{w*}, r \in \mathcal{R})$ denote the weighted max–min fair bandwidth allocation vector. For a given state $n$ we determine $a^{w*}$ in a similar fashion to the max–min fair allocation. First for all routes $r \in \mathcal{R}$ such that $n_r = 0$ set $a_r^{w*} = 0$. Next, replace (2) with

$$f^{(1),w}(n) := \max_a \left\{ \min_{r \in \mathcal{R}} \{a_r/w_r\} : \sum_{r \in \mathcal{R}} M_{\ell r} n_r a_r \le c_\ell,\ \ell \in \mathcal{L} \right\}$$

which can again be solved by first defining the *weighted fair share* on link $\ell$ as

$$f_\ell^{(1),w}(n) := \frac{c_\ell}{\sum_{r \in \mathcal{R}} M_{\ell r} w_r n_r} \tag{4}$$

and then setting $f^{(1),w}(n) = \min_{\ell \in \mathcal{L}} f_\ell^{(1),w}(n)$. Paralleling the max–min fair case, the first-level bottleneck links and routes, denoted $\mathcal{L}^{(1),w}$ and $\mathcal{R}^{(1),w}$ respectively, can be defined, and one can proceed iteratively to determine the bandwidth allocation for connections on all routes. We will let $\mu^w(n)$ denote the vector of bandwidths allocated to each route where $\mu_r^w(n) = w_r v_r^{-1} n_r a_r^{w*}$ connections/s, and let $\mu^w = (\mu_r^w: \mathbb{Z}_+^{\mathcal{R}} \to \mathbb{R}_+, r \in \mathcal{R})$.

One can again extend $\mu^w$ for real-valued arguments, i.e., from $\mathbb{Z}_+^{\mathcal{R}}$ to $\mathbb{R}_+^{\mathcal{R}}$, and show that

$$\mu^w(x) = \mu^m(Dx) \tag{5}$$

where $\mu^m$ corresponds to the unweighted max–min fair allocation discussed in the previous section, and $D = \operatorname{diag}(w)$, i.e., a square matrix with components $(w_r, r \in \mathcal{R})$ along its diagonal. Thus one way to view the weighted max–min fair allocation is as a max–min fair allocation where the "effective number" of ongoing connections is $Dx$. Moreover, one can easily see that the results in Proposition 2.1 also apply to $\mu^w$.

A weighted max–min fair allocation can be used to differentiate among connections following different routes and thus give priority based on geographic, administrative, or service requirements by grouping like connections on a route. However, in order to do so specific criteria for the selection of weights need to be developed. In principle one can consider control policies which adjust the weights based on the state of the network—a simple example is briefly considered in [18].

### C. Proportionally Fair Bandwidth Allocation

As a final alternative we consider a framework where utility functions $U_r: \mathbb{R}_+ \to \mathbb{R}$, $r \in \mathcal{R}$ have been associated with users whose connections follow particular routes. Here $U_r(a_r)$ is the utility to a user/connection on route $r$ of a bandwidth allocation $a_r$. A bandwidth allocation policy which maximizes the total network utility when the state is $n$ can be obtained by solving the following optimization problem:

$$\max_{a \ge 0} \left\{ \sum_{r \in \mathcal{R}} n_r U_r(a_r) : \sum_{r \in \mathcal{R}} M_{\ell r} n_r a_r \le c_\ell,\ \ell \in \mathcal{L} \right\} \tag{6}$$

where we assume that the utility functions are strictly concave and so the optimizer is unique. This approach to allocating bandwidth is pleasing in the sense that it finds an appropriate compromise between the resources required by a user's connection and the *overall* network utility.

In general it is unclear what types of utility functions would appropriately model the users. However, [15] and others, have considered the case where $U_r(a_r) = w_r \log a_r$. In this case they have shown that the maximizer $a^{p*} = (a_r^{p*}, r \in \mathcal{R})$ corresponds to a (weighted) *proportionally fair* bandwidth allocation, in the sense that the optimal allocation satisfies the link capacity constraints, and for any other feasible rate $a' = (a'_r, r \in \mathcal{R})$, the aggregate weighted proportional change is nonnegative, i.e.

$$\sum_{r \in \mathcal{R}} w_r n_r \frac{a'_r - a_r^{p*}}{a_r^{p*}} \le 0. \tag{7}$$

Determining the maximizer of (6) for log utility functions can be done explicitly for simple networks. Alternatively, as with max–min fairness, one can design rate control mechanisms that converge to the associated bandwidth allocation [16]. We will let

$$\mu_r^p(n) = w_r v_r^{-1} n_r a_r^{p*}$$

denote the total bandwidth allocated to connections along route $r \in \mathcal{R}$, in connections/s, and $\mu^p(n) = (\mu_r^p(n), r \in \mathcal{R})$. Again $\mu^p$ can be easily extended for real-valued arguments. We shall use the same notation to denote the extension of $\mu^p$ from $\mathbb{Z}_+^{\mathcal{R}}$ to $\mathbb{R}_+^{\mathcal{R}}$.

*Proposition 2.2:* The function $\mu^p: \mathbb{R}_+^{\mathcal{R}} \to \mathbb{R}_+^{\mathcal{R}}$ is radially homogeneous, in the sense that

$$\mu^p(\alpha x) = \mu^p(x),\ x \in \mathbb{R}_+^{\mathcal{R}}, \qquad \alpha > 0. \tag{8}$$

In the interior of the positive orthant $\mathbb{R}_+^{\mathcal{R}}$, the function $\mu^p$ is continuous, and has strictly positive components. Finally, $\mu^p$ is bounded.

*Proof:* The continuity of $\mu^p$ follows by considering the functional dependence on $x$ of the proportionally fair bandwidth allocation. The radial homogeneity can be shown as follows. Consider a change of variables $b_r = x_r a_r$ in (6) for the case of weighted logarithmic utility functions. Here $b_r$ denotes the bandwidth, in bits/s, allocated to route $r$. With this change of variables one finds that

$$\begin{aligned} b^*(x) &:= \arg\max_b \left\{ \sum_{r\in\mathcal{R}} x_r w_r \log(b_r/x_r) \colon Mb \le c;\ b \ge 0 \right\} \\ &= \arg\max_b \left\{ \sum_{r\in\mathcal{R}} x_r w_r \log(b_r) \colon Mb \le c;\ b \ge 0 \right\} \end{aligned} \tag{9}$$

from which it follows that $b^*(\alpha x) = b^*(x)$ and since $\mu_r^p(x) = w_r v_r^{-1} b_r^*(x)$ it follows that $\mu^p(\alpha x) = \mu^p(x)$. ■

## III. Stability of the Stochastic Network

In this section we will consider the stability of the stochastic network model defined in Section II for various types of bandwidth allocation. Assuming $\{A_r, D_r, r \in \mathcal{R}\}$ are independent Poisson processes on $[0, \infty)$, where $A_r$ has rate $\lambda_r$ and $D_r$ has rate 1, the evolution equation (1) defines a Markov chain in $\mathbb{Z}_+^{\mathcal{R}}$ with transition rates

$$q(n, m) = \begin{cases} \lambda_r, & m = n + e^r,\ r \in \mathcal{R} \\ \mu_r(n), & m = n - e^r,\ r \in \mathcal{R} \\ 0, & \text{otherwise} \end{cases} \tag{10}$$

for $m \ne n$ with $e^r = (e_s^r, s \in \mathcal{R})$ and $e_s^r = 1(r = s)$, and where $1(\,)$ denotes the indicator function. Thus, when the state is $n$, route $r$ sees arrivals with rate $\lambda_r$ and departures with rate $\mu_r(n)$. Note that when $n_r = 0$ we have $\mu_r(n) = 0$, thus $q(n, n - e^r) = 0$, and so the rates are supported on the positive orthant.

We use the notation $Q$ for the infinitesimal generator (*viz.*, rate matrix) of this continuous-time Markov chain. For a function $\varphi \colon \mathbb{R}_+^{\mathcal{R}} \to \mathbb{R}$, we write

$$\begin{aligned} Q\varphi(n) &:= \sum_{m\in\mathbb{Z}_+^{\mathcal{R}}} q(n, m)\varphi(m) \\ &= \sum_{m\in\mathbb{Z}_+^{\mathcal{R}}} q(n, m)[\varphi(m) - \varphi(n)] \end{aligned} \tag{11}$$

where the latter equality follows from the fact that $Q$ is conservative:

$$q(n, n) = -\sum_{m\ne n} q(n, m).$$

Note that $Q\varphi(n)$ can be interpreted as the expected drift, i.e., the change in $\varphi(N(t))$ when $N(t) = n$.

Clearly the Markov chain $\{N(t), t \ge 0\}$ is irreducible, and we say that it is stable, if and only if it is positive recurrent. We will show positive recurrence by constructing a Lyapunov function [23], [10]. For our system, a Lyapunov function is any function $V \colon \mathbb{Z}_+^{\mathcal{R}} \to \mathbb{R}_+$ such that there exists a finite set $K \subseteq \mathbb{Z}_+^{\mathcal{R}}$, where

$$\sup_{n\notin K} QV(n) < 0 \tag{12}$$

with $QV$ as defined in (11). Using our formula (10) for the transition rates we can rewrite $QV$ as

$$\begin{aligned} QV(n) = \sum_{r\in\mathcal{R}} \{&\lambda_r[V(n + e^r) - V(n)] \\ &+ \mu_r(n)[V(n - e^r) - V(n)]\}. \end{aligned} \tag{13}$$

Intuitively, (12) means that when the process $N(t)$ lies outside $K$, it is such that on average $V(N(t))$ is decreasing, i.e., has negative drift. Searching for appropriate Lyapunov functions can be a tedious procedure. Particularly in our case since the transition rates of our Markov chain are defined via optimization problems associated with the various bandwidth allocation criteria to be considered.

### A. Stability Under Max–Min Fair Bandwidth Allocation

We first consider the stability of the network when bandwidth is allocated according to the max–min fair criterion and thus the dynamics of the system are captured by (1) with $\mu$ replaced by $\mu^m$ as defined in Section II-A.

We begin by proposing the following *candidate* Lyapunov function:

$$V(x) = \max_{\ell\in\mathcal{L}} \left\{ c_\ell^{-1} \sum_{r\in\mathcal{R}} M_{\ell r} v_r x_r \right\}, \qquad x \in \mathbb{R}_+^{\mathcal{R}}$$

which can be interpreted as the maximum time to finish off the work load on any link in the network. For convenience we introduce the vectors

$$\xi^\ell = (\xi_r^\ell, r \in \mathcal{R}), \quad \xi_r^\ell := c_\ell^{-1} M_{\ell r} v_r, \quad \ell \in \mathcal{L} \tag{14}$$

and let $\varphi^\ell(x) = \langle \xi^\ell, x\rangle, \in \mathcal{L}$ where $\langle\cdot, \cdot\rangle$ denotes the standard inner product in $\mathbb{R}^{\mathcal{R}}$. With this notation we have that

$$V(x) = \max_{\ell\in\mathcal{L}} \varphi^\ell(x) = \max_{\ell\in\mathcal{L}} \langle \xi^\ell, x\rangle. \tag{15}$$

Thus $V$ is a piecewise linear function. Since the vectors $\xi^\ell$ have nonnegative components, the sets $\{x \in \mathbb{R}_+^{\mathcal{R}} \colon V(x) \le \alpha\}$ are compact convex polytopes, for all $\alpha \ge 0$. For a fixed $x$, one or more of the indices $\ell$ achieve the maximum in (15)—these are the first-level bottleneck links defined earlier. We will use $\mathcal{L}^{(1)}(x)$ to denote the dependence of the first-level bottleneck links on $x$. Similarly $\mathcal{R}^{(1)}(x)$ will denote the first-level bottleneck routes and $c^{(1)}(x)$ the remaining available capacity when the network state is $x$, see Section II-A.

Since for first-level bottleneck links the link capacity is fully utilized among ongoing connections, we would expect that, on average, the number of connections $\sum_{r\in\mathcal{R}} M_{\ell r} n_r$ on such a link will decrease as long as the average arrival rate does not exceed the link capacity. The following lemma makes this clear.

*Lemma 3.1:* Assume that the stability condition $M\rho < c$, i.e., $\sum_{r\in\mathcal{R}} M_{\ell r}\lambda_r v_r < c_\ell$, for all $\ell \in \mathcal{L}$ holds. Then, there is a constant $d > 0$, such that for all $x \in \mathbb{R}_+^{\mathcal{R}}$, and all $\ell^* \in \arg\max_{\ell \in \mathcal{L}} \varphi^{\ell}(x)$, i.e., first-level bottleneck links $\ell^* \in \mathcal{L}^{(1)}(x)$, we have

$$Q\varphi^{\ell^*}(x) = \left\langle \xi^{\ell^*}, \lambda - \mu^m(x) \right\rangle \leq -d. \tag{16}$$

Note that $\ell^*$ depends on $x$, but we suppress this dependence in our notation $Q\varphi^{\ell^*}(x)$.

*Proof:* First, using (13) and the definition (14) of $\xi^{\ell}$ we have that

$$\begin{aligned} Q\varphi^{\ell^*}(x) &= \sum_{r \in \mathcal{R}} \xi_r^{\ell^*}(\lambda_r - \mu_r^m(x)) = \left\langle \xi^{\ell^*}, \lambda - \mu^m(x) \right\rangle \\ &= \sum_{r \in \mathcal{R}} c_{\ell^*}^{-1} M_{\ell^* r} v_r (\lambda_r - \mu_r^m(x)). \end{aligned}$$

Next, since $\ell^*$ is a first-level bottleneck link, it follows that for routes $r$ traversing link $\ell^*$ we have $\mu_r^m(x) = v_r^{-1} x_r a_r^*$ where $a_r^*$ is given by (3). Thus

$$\begin{aligned} Q\varphi^{\ell^*}(x) &= c_{\ell^*}^{-1} \left( \sum_{r \in \mathcal{R}} M_{\ell^* r} \lambda_r v_r - \sum_{r \in \mathcal{R}} \frac{M_{\ell^* r} x_r c_{\ell^*}}{\sum_{s \in \mathcal{R}} M_{\ell^* s} x_s} \right) \\ &= c_{\ell^*}^{-1} \left( \sum_{r \in \mathcal{R}} M_{\ell^* r} \rho_r - c_{\ell^*} \right) \leq -d \end{aligned}$$

where $d := \max_{\ell \in \mathcal{L}} \{c_\ell^{-1}(c_\ell - \sum_{r \in \mathcal{R}} M_{\ell r} \rho_r)\}$ is positive by the stability condition. ■

Despite the promise of Lemma 3.1 it is unclear whether $V$ is an appropriate Lyapunov function. Indeed, the lemma only suggests that as long as the state makes transitions on regions having the *same* first-level bottleneck links, $V(N(t))$ will experience a negative drift. To make this more precise we will explicitly identify these regions and for clarity present an example in Section III-C. Let $\mathcal{M}$ be a nonempty subset of $\mathcal{L}$ and let

$$C_{\mathcal{M}} = \left\{ x \in \mathbb{R}_+^{\mathcal{R}} \colon \mathcal{L}^{(1)}(x) = \mathcal{M} \right\}. \tag{17}$$

It is clear that if $\alpha > 0$, $x \in C_{\mathcal{M}} \Rightarrow \alpha x \in C_{\mathcal{M}}$, i.e., these sets are cones, and that

$$\bigcup_{\mathcal{M} \subseteq \mathcal{L},\, \mathcal{M} \neq \emptyset} C_{\mathcal{M}} = \mathbb{R}_+^{\mathcal{R}}.$$

Suppose that $n \in C_{\mathcal{M}}$, for some nonempty $\mathcal{M}$, then the drift $QV(n)$ can easily be computed [see (13)], provided $n + e^r$, $n - e^r \in C_{\mathcal{M}}$, for all $r \in \mathcal{R}$. In this case, with $\ell$ any element of $\mathcal{M}$, we have $QV(n) = \langle \xi^{\ell}, \lambda - \mu^m(n) \rangle \leq -d$, by Lemma 3.1. However, when $n$ and ($n + e^r$ or $n - e^r$) belong to different cones an explicit verification of the negative drift requirement becomes difficult. Indeed, when this is the case a transition causes a change in the links that are bottlenecked, i.e., we are "crossing the boundary" of one of the cones.

Intuitively one might argue that this effect is negligible, since it occurs at a relatively small fraction of points in the state space. To make this intuition into a rigorous statement observe that Lemma 3.1 also implies that there is a $c > 0$, such that

$$\langle \nabla V(x), \lambda - \mu^m(x) \rangle \leq -d \tag{18}$$

for all $x$ at which the gradient $\nabla V(x) := (\partial V(x)/\partial x_r,\, r \in \mathcal{R})$ exists. It is easy to see that this gradient exists almost everywhere, and, when it exists, it equals $\xi^{\ell}$, for some $\ell$. In order to obtain an appropriate Lyapunov function in Lemma 3.2 below we follow the result of [9] which is based on showing the existence of a smoothened version $W$ of the function $V$ that satisfies a drift condition in the sense of (18) for all $x \in \mathbb{R}_+^{\mathcal{R}}$. The proof of this lemma follows [9] and can be found in [18].

*Lemma 3.2:* If $M\rho < c$, then there is a nonnegative function $W$, defined on $\mathbb{R}_+^{\mathcal{R}} \setminus \{0\}$, that is at least twice-continuously differentiable, has a Hessian $\nabla^2 W(x)$, such that $\nabla^2 W(x) \to 0$, as $|x| \to \infty$, and which satisfies the following drift condition: there is a $d > 0$, such that

$$\langle \nabla W(x), \lambda - \mu^m(x) \rangle \leq -d$$

for all $x \neq 0$. The Hessian $\nabla^2 W(x)$ is the $|\mathcal{R}| \times |\mathcal{R}|$ matrix with entries $\{(\partial^2 W/\partial x_r \partial x_s)(x),\, r,\, s \in \mathcal{R}\}$.

Given this results we can show that the network is indeed positive recurrent as follows.

*Theorem 3.1:* If $M\rho < c$ then the Markov chain $\{N(t),\, t \geq 0\}$ associated with the max–min fair bandwidth allocation is positive recurrent.

*Proof:* Since $W$ is twice differentiable it follows by the Mean Value Theorem that for $n,\, m \in \mathbb{Z}_+^{\mathcal{R}}$ there exists a $\theta$, $0 \leq \theta \leq 1$ such that

$$\begin{aligned} &W(n+m) - W(n) \\ &\quad = \langle \nabla W(n), m \rangle + \tfrac{1}{2} m^T \nabla^2 W(n + \theta m) m \\ &\quad := \langle \nabla W(n), m \rangle + \beta(n, m). \end{aligned}$$

Recall that $\nabla^2 W(n) \to 0$ and thus $\beta(n, z) \to 0$ as $|n| \to \infty$. Now, using this approximation to compute $QW$, as in (13), we have

$$QW(n) = \langle \nabla W(n), \lambda - \mu^m(n) \rangle + \sum_m q(n, m) \beta(n, m - n).$$

It follows by Lemma 3.2 that the first term is at most $-d$. The second term, is a sum of a finite number of terms, and can be made smaller than $d/2$ for all $|n|$ sufficiently large. Thus noting that $\sup_{|n| > \gamma} QW(n) < 0$, for sufficiently large $\gamma$, and letting $K = \{n \colon |n| \leq \gamma\}$ we satisfy the drift condition (12) which as discussed earlier implies positive recurrence. ■

### *B. Stability Under Weighted Max–Min Fair Bandwidth Allocation*

While the previous result is intuitive, in that the number of connections on bottleneck links must be decreasing, it is not easily extended to show the stability of networks under weighted max–min fair bandwidth allocation. Thus, we develop an alternative approach which, instead of focusing on links, focuses on the relative states of each route. Suppose that a set of weights $w$ is selected and the network is operated subject to the bandwidth allocation function $\mu^w$ defined in Section II-B. We will let

$$\varphi^r(x) = \rho_r^{-1} w_r x_r,\ r \in \mathcal{R}$$

and consider a candidate Lyapunov function

$$V(x) = \max_{r \in \mathcal{R}} \varphi^r(x) = \max_{r \in \mathcal{R}} \{\rho_r^{-1} w_r x_r\}. \tag{19}$$

The following lemma shows why this particular function is useful. The proof has been relegated to the Appendix.

*Lemma 3.3:* Assume that $M\rho < c$ then there is a constant $d > 0$, such that for all $x \in \mathbb{R}_+^{\mathcal{R}}$ and for all $r^* \in \arg\max_{r \in \mathcal{R}} \varphi^r(x)$ we have

$$Q\varphi^{r^*}(x) = \rho_{r^*}^{-1} w_{r^*}(\lambda_{r^*} - \mu_{r^*}^w(x)) \leq -d.$$

*Theorem 3.2:* If $M\rho < c$ then Markov chain $\{N(t), t \geq 0\}$ associated with weighted max–min fair bandwidth allocation is positive recurrent.

*Proof:* Based on Lemma 3.3, and the technique used in Lemma 3.2, it should be clear that an appropriately smooth Lyapunov function $W$ can be constructed from $V$ in (19). Positive recurrence then follows as in Theorem 3.1. ■

Note that since max–min fairness is a special case of weighted max–min fairness, Theorem 3.2 also establishes the stability of both. These results establish that $M\rho < c$ is a *sufficient* condition for stability. In fact, ignoring equality ($M\lambda = c$), this is a *necessary* condition. A brief argument can be made as follows. Say there exists a link $\ell$ such that $\sum_{r \in \mathcal{R}} M_{\ell r}\rho > c_\ell$. Clearly such a link in isolation is unstable, i.e., on average will tend to drift off to infinity. When the link is incorporated within a network, the situation can in fact only get worse, since other links can only slow down the departures for connections on $\ell$.

### C. Example Network

In this example we consider max–min fair bandwidth allocation for the network shown in Fig. 1. It consists of two routes $\mathcal{R} = \{1, 2\}$ and three links, $\mathcal{L} = \{1, 2, 3\}$, as shown in the figure. Fig. 2 shows the vector field $\lambda - \mu^m(x)$ corresponding to the case with $\lambda = (1.5, 1.5)$, $v = (1, 1)$ and $c = (5, 6, 4)$. We have shown the boundaries $x_1 = 5x_2$ and $2x_1 = x_2$ between three cones $C_{\{1\}}$, $C_{\{2\}}$ and $C_{\{3\}}$ corresponding to links 1, 2, or 3 being bottlenecks, respectively. Also shown on the figure is a level set of the function $V$. From the figure it is clear that on each cone the network dynamics push inwards, i.e., have negative drift with respect to $V$. By smoothing $V$ as in Lemma 3.2 we obtain a Lyapunov function $W$ from which the stability of the system follows.

### D. Stability Under Proportionally Fair Bandwidth Allocation

Unlike max–min fair bandwidth allocations that aim at maximizing the worst-case individual utility/performance criteria such as (weighted) proportional fairness attempt to maximize the overall network utility at any point in time. This is reflected in our candidate Lyapunov function:

$$W(x) = \sum_{r \in \mathcal{R}} \int_0^{x_r} v_r U_r'\left(\frac{\rho_r}{y_r}\right) dy_r = \sum_{r \in \mathcal{R}} \frac{w_r x_r^2}{2\lambda_r}$$

where utility functions are logarithmic i.e., $U_r(a_r) = w_r \log(a_r)$. Note that $W$ is continuous and twice differentiable, thus there is no need for the smoothing process required used previously. In Lemma 3.4 and Theorem 3.3 below we show

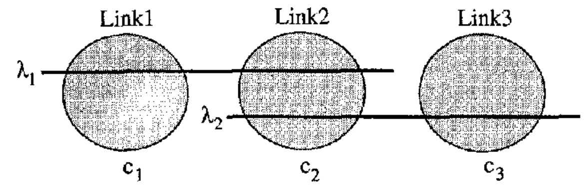


Fig. 1. Example network with three links and two routes.

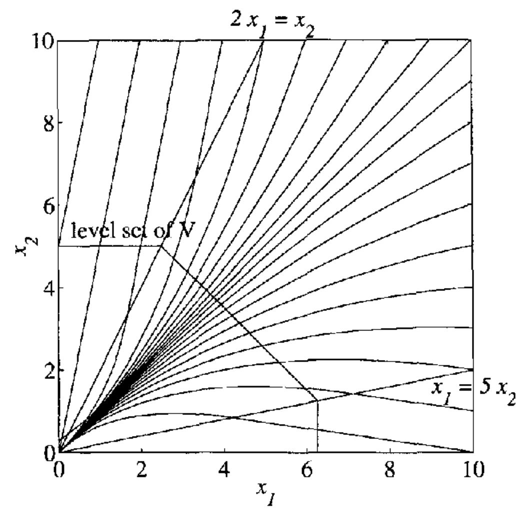


Fig. 2. A vector field of the example network.

that this function is indeed a Lyapunov function and show the stability of the Markov Chains associated with (weighted) proportionally fair bandwidth allocations.

*Lemma 3.4:* Assume that $M\rho < c$. Then given an $\varepsilon > 0$ there exists a constant $d > 0$ such that for all $x \in \mathbb{R}_+^{\mathcal{R}} \setminus \{0\}$ the following holds

$$\langle \nabla W(x), \lambda - \mu^p(x) \rangle = \sum_{r \in \mathcal{R}} v_r U_r'\left(\frac{\rho_r}{x_r}\right)(\lambda_r - \mu_r^p(x)) \leq -d.$$

*Proof:* Let $U(x, b) = \sum_{r \in \mathcal{R}} x_r U_r(b_r/x_r)$ denote the overall network utility of a network supporting $x_r$ users sharing a bandwidth allocation $b_r$ on route $r$. Note that $U(x, b)$ is strictly concave in $b$ and has a unique maximizer $b^*(x)$ over all feasible bandwidth allocations, see (9). Also note that $b^*(x) \neq \rho$ for any $x \neq 0$. Indeed, although $\rho$ is feasible, i.e., $M\rho < c$, it can not correspond to a utility maximizing vector since each route must have at least one bottlenecked link. Thus $U(x, b^*(x)) > U(x, \rho)$ and in fact the function $U(x, \alpha\rho + (1-\alpha)b^*(x))$ for $\alpha \in [0, 1]$ is strictly concave and decreasing in $\alpha$, thus

$$\left.\frac{\partial U(x, \alpha\rho + (1-\alpha)b^*(x))}{\partial \alpha}\right|_{\alpha=1}$$
$$= \sum_{r \in \mathcal{R}} U_r'\left(\frac{\rho_r}{x_r}\right)(\rho_r - b_r^*(x))$$
$$= \sum_{r \in \mathcal{R}} v_r U_r'\left(\frac{\rho_r}{x_r}\right)(\lambda_r - \mu_r^p(x)) < 0.$$

So for a given $x \neq 0$ by continuity of $\langle \nabla W(x), \lambda - \mu^p(x) \rangle$ in $x$ there exists $\delta$ and $d(x)$ such that for all $y$ in a ball $\|y -$

$x\| < \delta(x)$, we have $\langle \nabla W(y), \lambda - \mu^p(y)\rangle \le -d(x)$. The same property can be shown to hold uniformly on a compact set given by a sphere intersected with the positive orthant, i.e., $\langle \nabla W(x), \lambda - \mu^p(x)\rangle \le -d$ for all $x \in S(0, \varepsilon) \cap \mathbb{R}_+^{\mathcal{R}}$. This uniformity can in turn be extended to the set $\mathbb{R}_+^{\mathcal{R}} \setminus \{0\}$ using the following property. For $\alpha > 0$

$$\langle \nabla W(\alpha x), \lambda - \mu^p(\alpha x)\rangle = \alpha\langle \nabla W(x), \lambda - \mu^p(x)\rangle$$

since $\partial W(x)/\partial x_r = w_r x_r/\lambda_r$ whence $\nabla W(\alpha x) = \alpha \nabla W(x)$ and by radial homogeneity of $\mu^p$. It follows that

$$\langle \nabla W(x), \lambda - \mu^p(x)\rangle \le -d$$

for all $x$ in $\mathbb{R}_+^{\mathcal{R}} \setminus \{0\}$ and a constant $d > 0$. ■

*Theorem 3.3:* If $M\rho < c$ the Markov chain $\{N(t), t \ge 0\}$ associated with proportionally fair bandwidth allocation is positive recurrent.

*Proof:* The method of proof for this theorem is analogous to that of our previous results. Since $W$ is twice differentiable it follows by the Mean Value Theorem that for $n, m \in \mathbb{Z}_+^{\mathcal{R}}$ there exists a $\theta$, $0 \le \theta \le 1$ such that

$$W(n+m) - W(n) = \langle \nabla W(n), m\rangle + \tfrac{1}{2} m^T \nabla^2 W(n+\theta m)m.$$

Noting that the Hessian is diagonal i.e.,

$$\nabla^2 W(x) = \mathrm{diag}(w_r/\lambda_r, r \in \mathcal{R}),$$

and $m = e^r$ in (11) we have

$$QW(n) = \langle \nabla W(n), \lambda - \mu^p(n)\rangle + \langle h, \lambda - \mu^p(n)\rangle,$$

where $h = (w_r/2\lambda_r, r \in \mathcal{R})$. By Lemma 3.4 for $n \in \mathbb{Z}_+^{\mathcal{R}} \setminus \{0\}$ we have that

$$\begin{aligned} QW(\alpha n) &= \langle \nabla W(\alpha n), \lambda - \mu^p(\alpha n)\rangle + \langle h, \lambda - \mu^p(\alpha n)\rangle \\ &= \alpha\langle \nabla W(n), \lambda - \mu^p(n)\rangle + \langle h, \lambda - \mu^p(n)\rangle \\ &\le -\alpha d + \langle h, \lambda\rangle \\ &\le -\alpha d + |\mathcal{R}|/2, \end{aligned}$$

where $\alpha > 0$ and $|\mathcal{R}|/2$ is a finite constant. Thus for sufficiently large $\alpha$ or equivalently $|n|$, the drift can be made negative. Letting $K = \{n: |n| \le \gamma\}$ with large enough $\gamma$, we have $\sup_{n \notin K} QW(n) < 0$, which satisfies the drift condition (12) and implies positive recurrence.

The vector field corresponding to proportionally fair bandwidth allocation on the network in Fig. 1 along with a level set for the function $W$ are exhibited in Fig. 3 when $\lambda = (1.5, 1.5)$, $v = (1, 1)$, $c = (5, 6, 4)$ and $w = (1, 1)$.

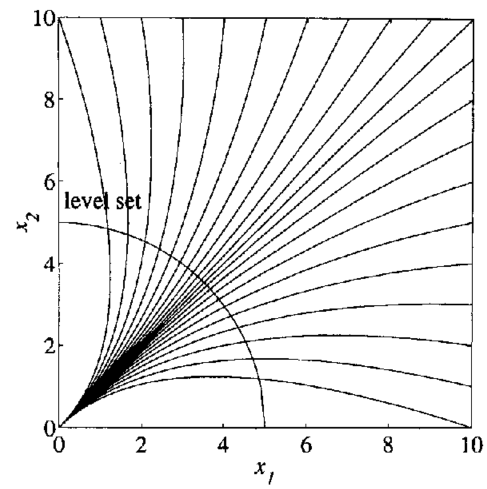


Fig. 3. A vector field corresponding to proportional bandwidth allocation of the example network.

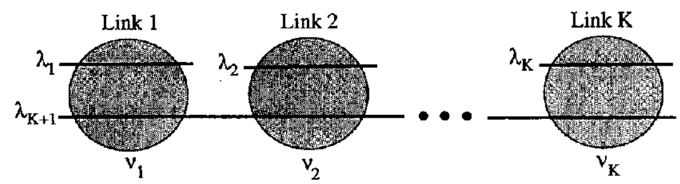


Fig. 4. A network for simulations.

TABLE II
SIMULATION ENVIRONMENT (SYMMETRIC LOADS ON ALL ROUTES)

| Load | $\lambda_r$ (conn./sec) | $v_r$ (kbits) | $\rho_r$ (kbps) | $c_\ell$ (kbps) |
|---|---|---|---|---|
| Light | 0.2 | 32 | 6.4 | 76.8 |
| Moderate | 2 | 32 | 64 | 192 |
| Heavy | 20 | 32 | 640 | 1344 |

## IV. PERFORMANCE

Quantifying the performance of dynamic networks supporting services with adaptive bandwidth allocations is a challenging task. Nevertheless this is an exceedingly interesting problem that network designers will eventually need to face. In this section we resort to simulation in an effort to further investigate the performance implications of various bandwidth allocation criteria. We focus on average connection delay as our performance metric. This metric is of interest in dimensioning networks to provide a reasonable call-level quality of service.

We shall consider a network consisting of $K$ links in series; see Fig. 4. A long route traverses each link in the network, while short single-link routes, model "cross traffic." This basic network serves to study the impact of short (local) on long (transit) traffic and vice versa. To investigate the degradation in performance as connections traverse an increasing number of links we have simulated several configurations with $K = 2, 3, 4$, and 5. To assess the impact of the bandwidth allocation mechanism we simulated networks operated under max–min, weighted max–min, and proportionally fair bandwidth sharing. In the case of weighted max–min fairness, short and long connections were given weights 1 and 2, respectively, i.e., $w_r = 1, r = 1, \ldots, K$ and $w_{K+1} = 2$. Thus priority was given to connections traversing several links as they are likely to experience the poorest performance. We considered various symmetric load conditions wherein long and short routes have the same traffic loads, i.e., $\lambda_r = \lambda_s, \forall r, s \in \mathcal{R}$. The load conditions are summarized in Table II.

Average overall connection delays as well as those on short and long routes, under max–min, weighted max–min, and proportionally fair allocation were measured as the number of links

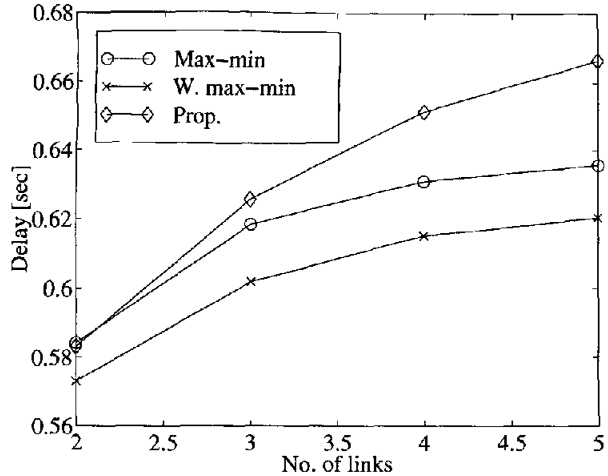


Fig. 5. Average overall delay (moderate load).

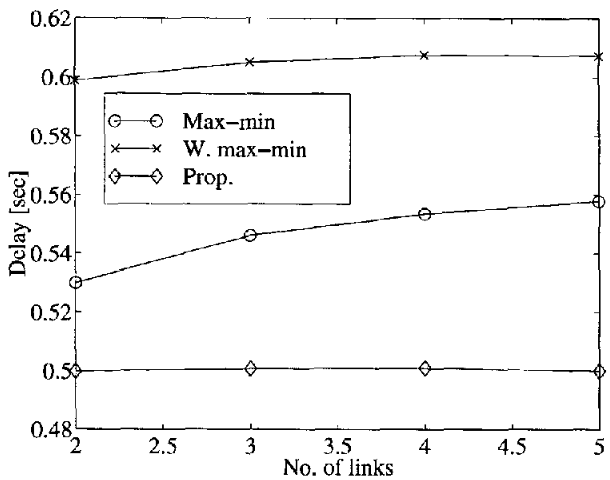


Fig. 6. Average delay on short routes (moderate load).

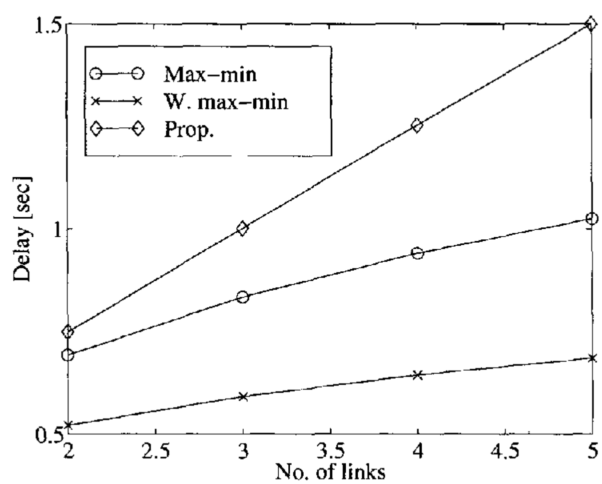


Fig. 7. Average delay on long routes (moderate load).

$K$ increases; see Figs. 5–7 for networks with moderate loads. All the simulations were conducted using an event-driven simulation program written in C. In our simulations, the 95% confidence interval of average connection delay was within range of 0.001–0.005 seconds of the sample mean. Not surprisingly, the results suggest that as traffic load becomes heavier, and long

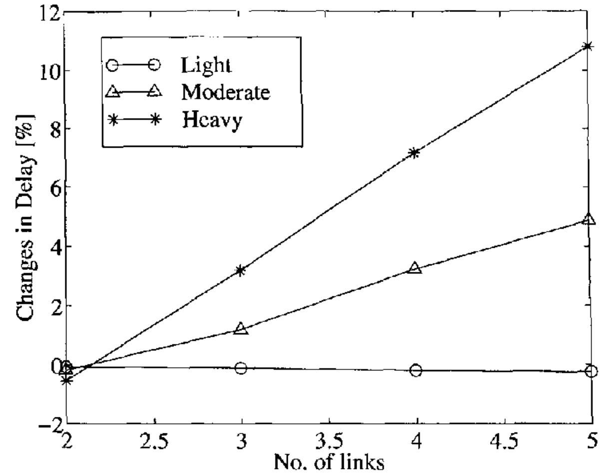


Fig. 8. Change in delays, prop. over max–min (overall).

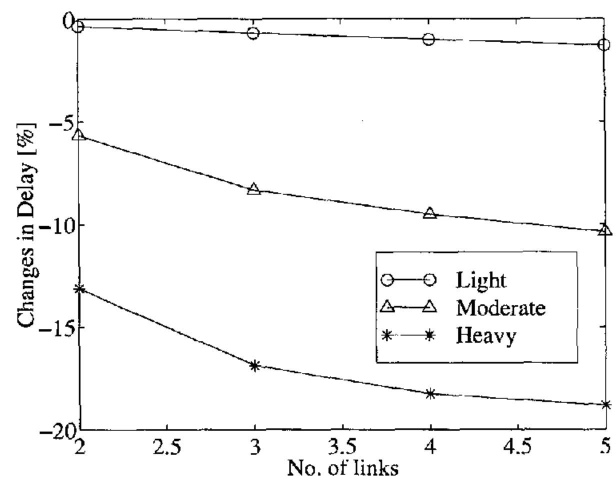


Fig. 9. Change in delays, prop. over max–min (short routes).

routes traverse a larger number of links, average overall connection delay becomes large, regardless of the bandwidth allocation policy.

We first contrast the performance of max–min fair bandwidth allocation, which strives to maximize the worst-case individual performance versus proportional fairness, which strives to maximize the overall network utility. The latter tends to give more bandwidth to connections crossing a small number of links, as they are more efficient in terms of their resource requirements. As a result long routes may linger in the network possibly degrading the overall performance. For example, for $K = 5$ and moderate load, the relative change in delays for proportional versus max–min fair bandwidth allocation is $-10\%$ on short routes, $+46\%$ on long routes, and $+5\%$ overall; see Figs. 8–10.

This result demonstrates that the max–min outperforms the proportionally fair allocation in terms of delays on long routes and overall delays. The same observation is made in [27]. Moreover, the change in delays for proportional versus max–min becomes larger as the size of network grows and the load of traffic becomes heavier. This suggests that maximizing the overall network utility is not necessarily compatible with minimizing connection delays. Note that as the number of links increases, proportional fairness leads to a surprisingly flat average delay on

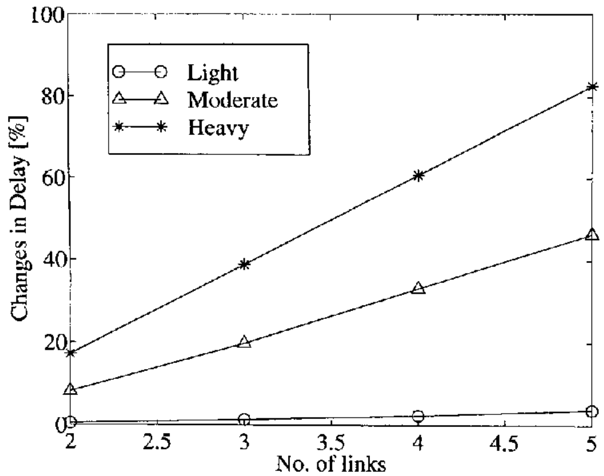


Fig. 10. Change in delays, prop. over max–min (long routes).

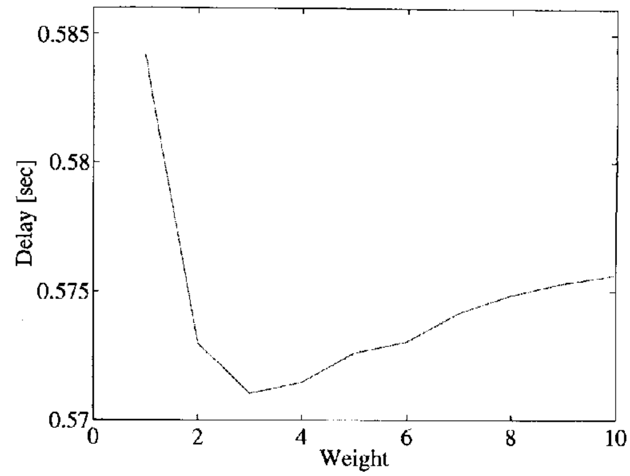


Fig. 11. Overall delay as the weight on a long route increases (moderate load, $K = 2$).

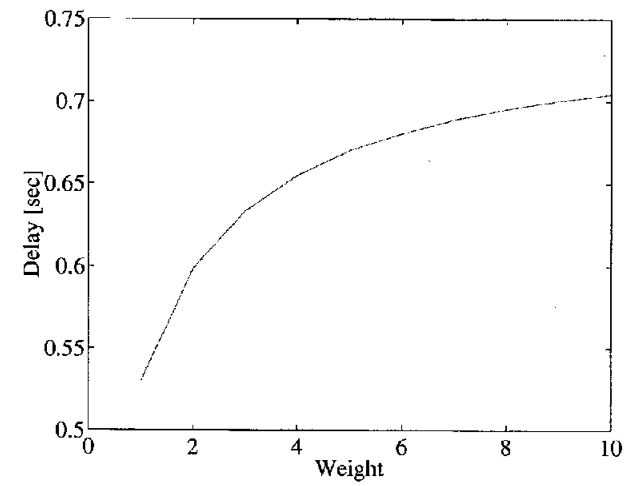


Fig. 12. Average delay on short routes as the weight on a long route increases (moderate load, $K = 2$).

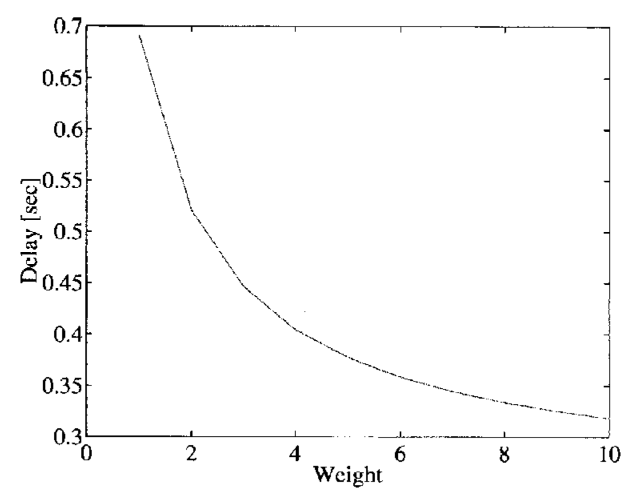


Fig. 13. Average delay on long routes as the weight on a long route increases (moderate load, $K = 2$).

short routes, while long routes see a linear growth in average delay; see Figs. 6 and 7. Thus the overall delay is not linear since it is an average of delays on short and long routes. Since the relative total load on short versus long routes is increasing with $K$, the overall delay behavior does not grow linearly. This suggests that proportional fairness may provide a clean performance differentiation among routes that have different lengths.

To make this point clearer we considered the impact that using a weighted max–min fair bandwidth allocation has on delays, if weights are selected so as to expedite connections on long routes. Clearly, weighted max–min fair allocation can provide additional flexibility in allocating bandwidth over max–min fair allocation. Continuing with our example, when $K = 5$ and the network load is moderate, the relative change in delays for the weighted versus the max–min fair bandwidth allocation is $+9\%$ on short routes, $-33\%$ on long routes, and $-2\%$ overall. Thus, one can not only dramatically improve the delays experienced on long routes, but also marginally improve the overall performance.

In order to see the impact of weight selection on network performance, we have measured the performance of a network with fixed $K = 2$ and load condition as weights for long routes vary. It turns out that overall performance is not continuously improved in proportion with the increase of weights given to long routes, although average delay on long routes decreases. For the moderate load condition and $K = 2$, performance is illustrated in Figs. 11–13. The overall delay is minimized when the weight $w_3 = 3$, and then degrades as the weight $w_3$ increases. This result shows that there is a tradeoff between improving delay performance for long connections and maximizing the overall delay performance, which can be achieved by optimizing weights (priorities).

These results exhibit the potential impact that a fairness criterion selected by designers may have on network performance. However, a better characterization of network performance and tools to "optimally" select weights, or route connections, will need to be developed if a call-level quality of service such as that considered here is deemed important in future networks. Also note that one could in theory introduce weights on a proportionally fair allocation in order to also enhance the performance seen on long connections. Hence our results do not suggest that a particular mechanism is best, we merely suggest that a consideration of these issues is warranted.

## V. Could the Internet be Unstable?

Internet traffic has been growing dramatically for the last few years. As of July 2000, the number of hosts advertised in the Domain Name Server (DNS) reached more than 93 million [30]. In many places the increase in demand is outpacing resources, leading to congestion and degradation in performance. Since performance of Internet traffic is closely linked to the behavior of TCP congestion avoidance algorithm [13], [22], it is crucial to understand the impact of TCP on the macroscopic network-level performance.

However, due to complicated interactions of Internet traffic and TCP transport algorithms [26], most research on the performance of TCP has relied on simulations for various TCP mechanisms. In an attempt to quantify throughputs of TCP connections more precisely and predictably, some researchers have started to consider analytical models and throughputs of TCP connections under various operating conditions, see e.g., [13], [11], [22], [25]. Recently, this approach has drawn much attention and relevant work is ongoing.

Mathis *et al.* [22] formulate a simple TCP model under the assumptions that 1) TCP is running over lossy path with constant Round Trip Time (RTT), and 2) packet loss is random with constant probability of $p$. The TCP throughput, $BW(p)$, is derived as

$$BW(p) = \frac{1}{\mathrm{RTT}}\sqrt{\frac{3}{2p}} \qquad \text{(packets/s)}.$$

The model is shown to match with real traffic when assumptions 1) and 2) hold. This model does not apply in some situations, e.g., when "timeout" behavior is dominant or for the case of short connections which require only a few cycles of congestion avoidance. In fact, real-life Internet traffic exhibits many timeouts compared with congestion avoidance behavior, i.e., retransmission.

A recently developed model by Padhye *et al.* [25] improves upon the previous one. The model captures not only congestion avoidance but also timeout behaviors that many real-life TCP traces exhibit. Moreover, their model is shown to fit a wider range of operating conditions, i.e., loss regimes. They assume that packet losses are correlated based on the fact that most current Internet employs drop-tail queueing policy and thus packets are likely to be lost again once previous packets experienced losses due to a full buffer. Their approximate model for TCP[2] throughput $BW(p)$, in packets/s, as a function of loss probability is given by

$$\min\left(\frac{W_{\max}}{\mathrm{RTT}}, \frac{1}{\mathrm{RTT}\sqrt{\frac{2bp}{3}}+T_0\min\left(1,3\sqrt{\frac{3bp}{8}}\right)p(1+32p^2)}\right)$$

where $W_{\max}$ is the maximum congestion window size, $b$ is the number of packets that are acknowledged by a received ACK, and $T_0$ is the time interval a sender waits before it starts retransmitting unacknowledged packets when a timeout occurs. Although the model may not fit all TCP traces under different implementations such as TCP-Tahoe or the Linux TCP implementation, it has been shown to match a broad range of real TCP traces and to predict the TCP throughput.

[2]They model TCP-Reno which is the most popular implementation of TCP in the Internet.

These models have focused on finding throughput of TCP connections given loss rates and RTTs. It is of interest to consider performance of a model for networks with stochastic arrivals and departures since connections are likely to be dynamically setup or terminated. In this paper, we have considered the stability and performance of an idealized model for a network supporting services that adjust their transmissions to network loads. The model is only a *rough caricature* of the Internet today, in that it assumes TCP operates efficiently by immediately achieving an *average* throughput related to a (weighted) proportionally fair bandwidth allocation. For a single congested link, weighted max–min or weighted proportionally fair allocation can model TCP appropriately [22], [7]. Indeed, it has been shown in [21] that a flow control mechanism with linear increase and multiplicative decrease (e.g., TCP) results in (weighted) proportionally fair bandwidth allocation. So a connection's throughput is dictated by a weighted allocation of resources at congested or bottleneck links. The average RTT experienced by connections and loss rate can be captured by weights given to connections which in turn impact the equilibrium throughput achieved by TCP connections. This model also parallels the one proposed and validated via simulation in [22]. We also assume that packets associated with a given TCP connection typically follow the same route, and connections send data in a greedy manner and depart. Subject to these, perhaps fanciful assumptions, one can show that network stability cannot be guaranteed unless the connection-level offered loads do not exceed the network's link capacities.

While this result is not entirely surprising, it presents an interesting architectural dilemma for future networks. Since routing algorithms on the Internet base their decisions on short term measures, i.e., are not explicitly tracking the long-term averages or number of ongoing connections required to assess the connection-level offered loads, there is no reason to believe that the Internet would satisfy a connection-level stability requirement. Instability would be perceived by users as an unacceptably low throughput, or inordinate delays, and typically cause them to abandon, thus in some sense solving the problem. To avoid such extremes one might overprovision the network. Unfortunately, this may result in a network which is still unstable, resulting in sporadic long-lasting congestion events that are challenging to explain.

The simulations we have presented also suggest that there is a need to consider other types of performance objectives such as minimizing overall connection delays in addition to providing fairness among connections. However, which performance objective we should choose in a dynamic network environment may depend on various aspects, e.g., network topology, application types, and priorities to connections. It would of course be interesting to look at congestion patterns on the Internet today and attempt to explain them in terms of a connection-level instability. However, given the typically nonstationary demands on today's networks and the detailed data that would be required to provide a conclusive answer to this question, this appears to be a challenging task.

## APPENDIX A
## PROOF OF LEMMA 3.3

Start proof by observing that, due to (13)

$$Q\varphi^r(x) = \frac{w_r}{\rho_r}(\lambda_r - \mu_r(x)). \tag{20}$$

Suppose the network state is $x$ and let

$$r^* \in \arg\max_{r\in\mathcal{R}} \{\rho_r^{-1} w_r x_r\}$$

then for all $r \in \mathcal{R}$, we have that

$$\frac{w_r x_r}{\rho_r} \le \frac{w_{r^*} x_{r^*}}{\rho_{r^*}} \tag{21}$$

or equivalently that $\rho_{r^*} w_r x_r \le \rho_r w_{r^*} x_{r^*}$. Now summing over all routes traversing a link $\ell \in \mathcal{L}$ we have that

$$\rho_{r^*} \sum_{r\in\mathcal{R}} M_{\ell r} w_r x_r \le w_{r^*} x_{r^*} \sum_{r\in\mathcal{R}} M_{\ell r}\rho_r$$

which one can rearrange to obtain

$$\begin{aligned}\rho_{r^*} &\le \frac{w_{r^*} x_{r^*}}{\sum_{r\in\mathcal{R}} M_{\ell r} w_r x_r} \sum_{r\in\mathcal{R}} M_{\ell r}\rho_r \\ &= \frac{w_{r^*} x_{r^*}}{\sum_{r\in\mathcal{R}} M_{\ell r} w_r x_r} c_\ell \\ &\quad - \frac{w_{r^*} x_{r^*}}{\sum_{r\in\mathcal{R}} M_{\ell r} w_r x_r}\left(c_\ell - \sum_{r\in\mathcal{R}} M_{\ell r}\rho_r\right).\end{aligned}$$

Given (21) and the stability condition one can easily show the existence of a positive lower bound, $\varepsilon > 0$, for the second term on the right-hand side

$$\begin{aligned}&\frac{w_{r^*} x_{r^*}}{\sum_{r\in\mathcal{R}} M_{\ell r} w_r x_r}\left(c_\ell - \sum_{r\in\mathcal{R}} M_{\ell r}\rho_r\right) \\ &\ge \min_{\ell\in\mathcal{L}} \min_{s\in\mathcal{R}} \left\{ \frac{M_{\ell s}\rho_s}{\sum_{r\in\mathcal{R}} M_{\ell r}\rho_r}\left(c_\ell - \sum_{r\in\mathcal{R}} M_{\ell r}\rho_r\right)\right\} = \varepsilon.\end{aligned}$$

Thus we have that

$$\rho_{r^*} \le w_{r^*} x_{r^*} f_\ell^{(1),w}(x) - \varepsilon,$$

where we recognize a term corresponding to the fair share at link $\ell$, $f_\ell^{(1),w}(x)$; see (4). Moreover, since this is true for all $\ell$, and multiplying through by $v_{r^*}^{-1}$ we have that

$$\lambda_{r^*} \le w_{r^*} v_{r^*}^{-1} x_{r^*} f^{(1),w}(x) - v_{r^*}^{-1}\varepsilon \tag{22}$$

where $f^{(1),w}(x) = \min_{\ell\in\mathcal{L}} f_\ell^{(1),w}(x)$ is the fair share at first-level bottleneck links $\mathcal{L}^{(1),w}(x)$.

Now if $r^*$ is a first-level bottleneck route, i.e., $r^* \in \mathcal{R}^{(1),w}(x)$, then $\mu_{r^*}^w(x) = w_{r^*} v_{r^*}^{-1} x_{r^*} f^{(1),w}(x)$, and it follows by (22) that $\lambda_{r^*} - \mu_{r^*}^w(x) \le -v_{r^*}^{-1}\varepsilon$. If $r^*$ is not a first-level bottleneck route, we will show that its bandwidth allocation must exceed $w_{r^*} v_{r^*}^{-1} x_{r^*} f^{(1),w}(x)$ and so again by (22) we have that $\lambda_{r^*} - \mu_{r^*}^w(x) \le -v_{r^*}^{-1}\varepsilon$.

We begin by showing that $f^{(2),w}(x) \ge f^{(1),w}(x)$. Suppose $\ell \in \mathcal{L} \setminus \mathcal{L}^{(1),w}(x)$ and note that

$$\sum_{r\in\mathcal{R}} M_{\ell r} w_r x_r f^{(1),w}(x) \le \sum_{r\in\mathcal{R}} M_{\ell r} w_r x_r f_\ell^{(1),w}(x) = c_\ell$$

so it follows that

$$\begin{aligned}c_\ell - f^{(1),w}(x) &\sum_{r\in\mathcal{R}^{(1),w}(x)} M_{\ell r} w_r x_r \\ &\ge f^{(1),w}(x) \sum_{r\in\mathcal{R}\setminus\mathcal{R}^{(1),w}(x)} M_{\ell r} w_r x_r.\end{aligned}$$

Rearranging terms and recalling the definition of fair share for the links in the second level of the bottleneck hierarchy, we have that

$$\begin{aligned}f_\ell^{(2),w}(x) &= \frac{c_\ell^{(1),w}(x)}{\sum_{r\in\mathcal{R}\setminus\mathcal{R}^{(1),w}(x)} M_{\ell r} w_r x_r} \\ &= \frac{c_\ell - f^{(1),w}(x)\sum_{r\in\mathcal{R}^{(1),w}(x)} M_{\ell r} w_r x_r}{\sum_{r\in\mathcal{R}\setminus\mathcal{R}^{(1),w}(x)} M_{\ell r} w_r x_r} \ge f^{(1),w}(x).\end{aligned}$$

Thus $f^{(2),w}(x) = \min_{\ell\in\mathcal{L}} f_\ell^{(2),w}(x) \ge f^{(1),w}(x)$. Similarly it follows by induction that $f^{(i+1),w}(x) \ge f^{(i),w}(x)$, until the bottleneck hierarchy is exhausted.

Now since $\mu_{r^*}^w(x) = w_{r^*} v_{r^*}^{-1} x_{r^*} f^{(j),w}(x)$ for some level $j$ in the bottleneck hierarchy, it follows that

$$\mu_{r^*}^w(x) \ge w_{r^*} v_{r^*}^{-1} x_{r^*} f^{(1),w}(x)$$

and so $\lambda_{r^*} - \mu_{r^*}^w(x) \le -v_{r^*}^{-1}\varepsilon$. The lemma follows by selecting $d = \varepsilon \min_{r\in\mathcal{R}} \{\lambda_r^{-1} w_r\}$ and by (20). ∎

## ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers and editor S. Tripathi for their constructive comments.

## REFERENCES

[1] S. P. Abraham and A. Kumar, "A stochastic approximation approach for max–min fair adaptive rate control of ABR sessions with MCRs," in *Proc. IEEE INFOCOM*, vol. 3, 1998, pp. 1358–1365.
[2] E. Altman, F. Bacelli, and J.-C. Bolot, "Discrete-time analysis of adaptive rate control mechanism," in *Proc. 5th Int. Conf. Data and Communications*, 1993, pp. 121–140.
[3] L. Benmohamed and S. M. Meerkov, "Feedback control of congestion in packet switching networks the case of a single congested node," *IEEE/ACM Trans. Networking*, vol. 1, pp. 693–709, Dec. 1993.
[4] D. Bertsekas and R. Gallager, *Data Networks*. Englewood Cliffs, NJ: Prentice Hall, 1992.

[5] F. Bonomi and K. W. Fendick, "The rate-based flow control framework for the available bit rate ATM service," *IEEE Network Mag.*, pp. 25–39, Mar./Apr. 1995.
[6] F. Bonomi, D. Mitra, and J. B. Seery, "Adaptive algorithms for feedback-based flow control in high-speed wide-area ATM networks," *IEEE J. Select. Areas Commun.*, vol. 13, pp. 1267–1283, Sept. 1995.
[7] J. Crowcroft and P. Oechslin, "Differentiated end-to-end Internet services using a weighted proportional fair sharing TCP," *Comput. Commun. Rev.*, vol. 28, no. 3, 1998.
[8] 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 *Proc. IEEE INFOCOM*, vol. 2, 1999, pp. 802–810.
[9] D. Down and S. P. Meyn, "Piecewise linear test functions for stability and instability of queueing networks," *Queueing Syst.*, vol. 27, pp. 205–226, 1997.
[10] G. Fayolle, V. A. Mayshev, and M. V. Menshikov, *Topics in Constructive Theory of Countable Markov Chains*. Cambridge and New York: Cambridge University Press, 1995.
[11] S. Floyd, "Connections with multiple congested gateways in packet-switched networks—Part 1: One-way traffic," *Comput. Commun. Rev.*, vol. 21, no. 5, 1991.
[12] E. M. Gafni and D. Bertsekas, "Dynamic control of session input rates in communication networks," *IEEE Trans. Automat. Contr.*, vol. 29, pp. 1009–1016, 1984.
[13] V. Jacobson, "Congestion avoidance and control," in *Proc. ACM SIGCOMM*, Aug. 1988.
[14] V. Jacobson and M. Karels, "Congestion avoidance and control," in *Proc. ACM SIGCOMM*, Aug. 1988, pp. 314–329.
[15] F. Kelly, "Charging and rate control for elastic traffic," *Eur. Trans. Telecommun.*, vol. 8, pp. 33–37, 1997.
[16] F. P. Kelly, A. K. Mauloo, and D. K. H. Tan, "Rate control in communication networks shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 15, no. 49, pp. 237–255, 1998.
[17] R. J. La and V. Anantharam, "Charge-sensitive TCP and rate control in the Internet," in *Proc. IEEE INFOCOM*, vol. 3, 2000, pp. 1166–1175.
[18] T.-J. Lee, "Traffic management and design of multiservice networks: The Internet and ATM networks," Ph.D. dissertation, Department of Electrical and Computer Engineering, University of Texas, Austin, 1999.
[19] T.-J. Lee and G. de Veciana, "A decentralized framework to achieve max–min fair bandwidth allocation for ATM networks," in *Proc. IEEE GLOBECOM*, vol. 3, 1998, pp. 1515–1520.
[20] S. Low and D. Lapsley, "Optimization flow control—I: Basic algorithm and convergence," *IEEE/ACM Trans. Networking*, vol. 7, pp. 861–874, Dec. 1999.
[21] L. Massoulié and J. Roberts, "Bandwidth sharing: Objectives and algorithms," in *Proc. IEEE INFOCOM*, vol. 3, 1999, pp. 1395–1403.
[22] M. Mathis, J. Semke, J. Mahdavi, and T. Ott, "The macroscopic behavior of the TCP congestion avoidance algorithm," *Comput. Commun. Rev.*, vol. 27, no. 3, pp. 67–82, 1997.
[23] S. P. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability*. Berlin, Germany: Springer-Verlag, 1993.
[24] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," IEEE/ACM Trans. Networking, pp. 556–567, Oct. 2000, submitted for publication.
[25] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP throughput: A simple model and its empirical validation," presented at the ACM SIGCOMM, Vancouver, B.C., Canada, Sept. 1998.
[26] V. Paxson and S. Floyd, "Why we don't know how to simulate the Internet," in *Proc. 1997 Winter Simulation Conf.*, 1997.
[27] J. W. Roberts and L. Massoulié. Bandwidth sharing and admission control for elastic traffic. ITC Specialist Seminar 1998. [Online]. Available: http://www-sop.inria.fr/mistral/pub/massoulie.html
[28] S. Shenker, "Fundamental design issues for the future Internet," *IEEE J. Select. Areas Commun.*, vol. 13, pp. 1176–1188, Sept. 1995.
[29] C.-F. Su, G. de Veciana, and J. Walrand, "Explicit rate flow control for ABR services in ATM networks," *IEEE/ACM Trans. Networking*, vol. 8, no. 3, pp. 350–361, Jun. 2000.
[30] Internet domain survey, Network Wizards. (1999). [Online]. Available: http://www.nw.com

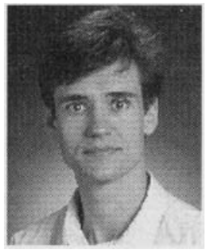

**Gustavo de Veciana** (S'88–M'94) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of California, Berkeley, in 1987, 1990, and 1993, respectively.

In 1993, he joined the Department of Electrical and Computer Engineering, University of Texas at Austin, where he is currently an Associate Professor. His research focuses on issues in the design and control of telecommunication networks.

Dr. de Veciana is an Editor for the IEEE/ACM TRANSACTIONS ON NETWORKING. He is the recipient of a General Motors Foundation Centennial Fellowship in Electrical Engineering and a 1996 National Science Foundation CAREER Award.

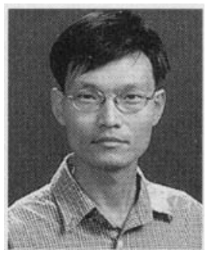

**Tae-Jin Lee** (S'89–M'91) received the B.S. and M.S. degrees in electronics engineering from Yonsei University, Seoul, Korea, in 1989 and 1991, respectively, and the M.S.E. degree in electrical engineering and computer science from the University of Michigan, Ann Arbor, in 1995. He received the Ph.D. degree in electrical and computer engineering from the University of Texas, Austin, in 1999.

In 1999, he joined the Corporate R & D Center, Samsung Electronics, Korea. His research interests include performance evaluation, traffic management, and design of communication networks, and wireless communication systems.

**Takis Konstantopoulos** (M'90) received the Diploma in electrical engineering from the National Technical University of Athens, Greece, and the M.Sc. and Ph.D. degrees from the University of California, Berkeley, in 1983, 1985, and 1989, respectively.

He held the position of Charge de Recherce in INRIA, Sophia Antipolis, France, between 1989 and 1990. He has held several visiting faculty positions including the University of California, Berkeley, Cornell University, Ithaca, NY, the University of Maryland, College Park, the University of Patras, Greece, and the Technical University of Braunschweig, Germany. He is currently an Associate Professor with the Department of Electrical and Computer Engineering and Computational and Applied Mathematics at the University of Texas at Austin. His main research interests are in performance analysis, modeling and control of communication networks, applied probability and stochastic processes, stochastic systems, and operations research.

Dr. Konstantopoulos is a recipient of several National Science Foundation awards. He is currently Associate Editor of *Queueing Systems: Theory and Applications*. He is a member of the American Mathematical Society, the Mathematical Association of America, and the Institute of Mathematical Statistics.