December 2003

Stability of Adaptive Congestion-Controllers for High Bandwidth Connections

Egemen Kavak
University of Pennsylvania

Srisankar S. Kunniyur
University of Pennsylvania, kunniyur@seas.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/ease_papers

---

Recommended Citation

Egemen Kavak and Srisankar S. Kunniyur, "Stability of Adaptive Congestion-Controllers for High Bandwidth Connections", . December 2003.

Copyright 2003 IEEE. Reprinted from *Proceedings of the 42nd IEEE Conference on Decision and Control 2003*, Volume 4, pages 3659-3664.

Publisher URL: <http://ieeexplore.ieee.org/xpl/tocresult.jsp?isNumber=28456&page=5>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ease_papers/60
For more information, please contact repository@pobox.upenn.edu.

Stability of Adaptive Congestion-Controllers for High Bandwidth Connections

Abstract

In this paper we use fluid-models to look at the adaptation of congestion-controllers to achieve higher throughputs and utilizations in high bandwidth connections. We first parameterize the congestion-controllers using a parameter called the multiplicative decrease parameter and study the adaptation of the congestion-controllers in terms of adapting this parameter. We then linearize the system to study the local stability properties and provide design rules for choosing the parameters of the congestion-controllers. We assume a general network topology and arbitrary round-trip delays in the analysis. Simulations that show the throughput increase that can be achieved using such adaptive congestion-controllers are also presented.

Comments

Copyright 2003 IEEE. Reprinted from *Proceedings of the 42nd IEEE Conference on Decision and Control 2003*, Volume 4, pages 3659-3664.

Publisher URL: <http://ieeexplore.ieee.org/xpl/tocresult.jsp?isNumber=28456&page=5>

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Stability of Adaptive Congestion-Controllers for High Bandwidth Connections

Egemen Kavak
kavak@seas.upenn.edu

Srisankar S. Kunniyur¹
kunniyur@seas.upenn.edu

Abstract

In this paper we use fluid-models to look at the adaptation of congestion-controllers to achieve higher throughputs and utilizations in high bandwidth connections. We first parameterize the congestion-controllers using a parameter called the multiplicative decrease parameter and study the adaptation of the congestion-controllers in terms of adapting this parameter. We then linearize the system to study the local stability properties and provide design rules for choosing the parameters of the congestion-controllers. We assume a general network topology and arbitrary round-trip delays in the analysis. Simulations that show the throughput increase that can be achieved using such adaptive congestion-controllers are also presented.

1 Introduction

Designing robust congestion-controllers have been an active research area in the Internet community [1, 2, 3, 4]. The users are associated with an utility function and the congestion-control problem is cast as a convex optimization problem. The steepest ascent algorithm to this convex optimization problem then acts as the congestion-controller for the user [1, 2, 3]. As a result, the congestion-controllers are modeled as ordinary differential equations with time-delayed feedback. The most common congestion-controller in the current Internet is the TCP congestion-controller. A TCP user (or congestion-controller) uses a window-flow control algorithm to implement its congestion-controller. The sender sends packets into the network which gets marked (or dropped) when a router detects incipient congestion. The marked or the dropped packet acts as the feedback from the router. The mark is then echoed back by the receiver to the sender. The sender then reduces the sending rate by reducing the window size.

In this paper we look at the adaptation of the congestion-controllers in the presence of high bandwidth links. We illustrate this using the current TCP congestion-control algorithm. Consider a huge file transfer using a TCP connection traversing a 10 Gbps router. A typical (approx-

Figure 1: Typical evolution of the TCP flow rate

imate) evolution of the TCP flow rate is shown in Figure 1. The TCP congestion controllers increases its rate aggressively initially using a *slow-start* algorithm in which the flow-rate increases exponentially. After this phase, the congestion-controller uses a linear increase to probe for additional bandwidth in the network (called the *congestion-avoidance* phase). When the flow rate causes a congestion on the link (say at time T_1), the link provides feedback to the user by either marking the packet or by dropping the packet. This feedback causes the TCP to reduce its sending rate by half (called the *multiplicative-decrease* phase). The TCP then tries to increase its congestion-window using the linear congestion-avoidance phase. The motivation behind the multiplicative decrease and congestion avoidance phase was to prevent TCP connections from causing a congestion collapse in the network. But this motivation assumes that at least one link in a flow's path is either utilized by a large number of flows and/or has a small capacity. However, in the case of large bandwidth links, the huge decrease coupled with the slow growth of the congestion window in the congestion avoidance phase limits the throughput of the TCP connection. Thus, one would like a multiplicative decrease phase that is small (i.e., reduce the flow-rate in small steps) when the available bandwidth at the links along its path is huge but conservative in a normal setting. To implement such a scheme, an user can either estimate the available bandwidth on its path or rely on feedback from the routers which specify the exact reduction in its rate. Such a feedback mechanism requires the routers to maintain the state of each flow. In the current Internet, this is infeasible due to the presence of a large number of flows. In this paper, we wish to adapt the multiplicative-decrease parameter of the

¹E. Kavak and S. S. Kunniyur are with the Department of Electrical and Systems Engineering at the University of Pennsylvania, Philadelphia, PA 19104

congestion-controller when the available bandwidth along the path of a flow is large, in order to yield high utilization while maintaining the normal operation of the congestion-controller in a normal setting. We use the current flow-rate of the user as an indication of the available bandwidth on its path. We assume a preset level (called x_{ref}) as the threshold to determine if a flow traverses a path of high bandwidth links. When the current flow-rate is above the threshold (x_{ref}), we assume uncongested high bandwidth links along the route and adapt the multiplicative decrease parameter to achieve a better utilization of the links along the path. When the current flow-rate is below the threshold (x_{ref}), we assume that the flow passes through congested links and hence use the prescribed value (for backward compatibility and to be fair to TCP users) of the multiplicative decrease parameter (eg., for TCP the prescribed value is $\frac{1}{2}$).

Adapting the multiplicative decrease parameter as a function of the current flow-rate can also be thought of as adapting the congestion-controller to the conditions in the network. One can think of this adaptation as the user adapting its utility function when the link capacities are high. The steady-state throughput of the congestion-controller is a function of the marking probability at the link. When the number of users at the link is sufficiently large (with a high throughput per user), the effect of an user on the marking probability can be assumed to be small. Consider a TCP congestion-controller with a multiplicative decrease parameter of β . The steady state throughput (using a fluid-model) of the TCP connection (x^*) can be written as: $x^* \propto \frac{1}{\sqrt{\beta p}}$, where p is the fraction of packets marked at the link. To achieve a high throughput, a very small marking probability at the link is required. As a result, one can think of adapting the multiplicative decrease parameter as modifying this function. In [4], the authors propose a protocol which changes this dependence from $\sqrt{p}$ to $p^{0.8}$ for TCP. In this paper, we look at such adaptation from a fluid-model perspective and study the local stability of such systems. We also assume a generalized congestion-controller and a generalized adaptation function for the decrease parameter. We then study the stability of this system assuming an arbitrary topology of the network and arbitrary round-trip delays for the users.

The rest of the paper is organized as follows: in Section 2 we describe the system model and derive the local stability properties of the system. We then present some preliminary simulations in Section 3 that agree with the analytical results in Section 2. We finally conclude in Section 4.

2 System Model

We adopt the model described in [1]. Consider a network with a set $\mathcal{L}$ of links and a set $\mathcal{R}$ of users. Associate a route r with each user where r is a non-empty subset of $\mathcal{L}$. The terms user, flow and route will be used interchangeably throughout the paper. Assume User r generates traffic at a

rate x_r . The rate x_r is assumed to have an utility $U_r(x_r)$ to flow r . We will assume that the utility functions are strictly concave functions and that $U_r(x) \rightarrow \infty$ as $x \rightarrow 0$ for all $r \in \mathcal{R}$. In this paper we will assume only utility functions of the form $\frac{1}{x_r^{v_r}}$ for some $v_r \geq 1$.

For each user r , let $d_1(r, j)$ be the delay from the source to link j , $d_2(r, j)$ be the feedback delay from link j back to the source and T_r be the total round-trip delay for route r . Note that

$$T_r = d_1(r, j) + d_2(r, j), \quad \forall j \in r.$$

Now, let each user $r \in \mathcal{R}$ employ the following congestion control algorithm

$$\dot{x}_r(t) = \kappa_r \left(w_r - \frac{\beta_r x_r(t - T_r)}{\tilde{U}'_r(x_r(t))} \sum_{l \in r} p_l \left(\sum_{\substack{j: \\ l \in j}} x_j(t - d_{lr}) \right) \right), \quad (1)$$

where $\tilde{U}'_r(x_r(t)) \triangleq x_r(t) U'_r(x_r(t))$, and $d_{lr} \triangleq d_1(j, l) + d_2(r, l)$, the parameter κ_r determines the speed of the congestion controller, $w_r (> 0)$ is the weight of user r , $p_l(\cdot)$ is the fraction of packets marked at link l and β_r denotes the multiplicative decrease parameter of user r . When $U_r(x_r) = \frac{1}{x_r}$, (1) results in a TCP-like congestion controller [3]. For a TCP congestion-controller, β_r is approximately equal to 0.5. In the absence of feedback delays in the system, it is shown in [1] that the above congestion control scheme converges to a unique equilibrium point which is a function of β . The local stability of the above scheme in the presence of feedback delays when proportionally-fair congestion-controllers are employed is shown in [5]. For the more general case of utility functions, the following theorem states a sufficient condition for the local stability of the congestion-controller.

Theorem 2.1 *The system described by the set of differential equations in (1) is locally asymptotically stable if for all $r \in \mathcal{R}$:*

$$\frac{\kappa_r \beta_r}{\tilde{U}'_r(\hat{x}_r)} \left(\sum_{l \in r} \hat{p}_l + \sum_{l \in r} \frac{\partial p_l}{\partial x_r} \left(\sum_{\substack{m: \\ l \in m}} \hat{x}_m \right) \sum_{\substack{j: \\ l \in j}} \hat{x}_j \right) < \frac{\pi}{2T_r}, \quad (2)$$

where $\hat{x}_r$ and $\hat{p}$ denote the equilibrium values of the flow-rates and marking probability.

Proof: The proof is an extension of the proof for log utility functions shown in [5]. Due to space constraints we omit the proof in this manuscript. ■

In this paper, we adapt the multiplicative decrease parameter β_r to achieve high utilizations when traversing high bandwidth links. Without loss of generality, we assume that each user starts with a multiplicative decrease parameter of $\beta_{\max}$. For example, $\beta_{\max} = 0.5$ for a TCP connection and each user starts its connection assuming this value

of the multiplicative decrease factor. We also assume that $\beta_r \in (0, \beta_{\max}]$.

Now, each user r updates β_r using the following algorithm

$$\dot{\beta}_r(t) = \begin{cases} \gamma_r (\beta_{\max} - \beta_r(t)), & x_r(t) \leq x_{ref} \\ \gamma_r (f_r(x_r(t)) - \beta_r(t)), & x_r(t) > x_{ref} \end{cases}, \quad (3)$$

where $f_r(\cdot)$ is a continuously differentiable, strictly decreasing function, γ denotes the speed of adaptation of the β_r parameter and $\beta_{\max}$ is the default value of the multiplicative decrease parameter. We also assume that $f_r(x_{ref}) = \beta_{\max}$.

Note that the particular form of the β adaptation is motivated by the desire to maintain the status-quo (or backward compatibility) of the congestion-controllers in the absence of high-bandwidth links. That is, in the absence of high bandwidth links (which is determined by the current rate of the user and x_{ref}), the equilibrium value of β_r is compatible with current standards. Also note that the choice of f_r decides the utility function that the user would like to employ at high bandwidths. As a result, we can consider the β adaptation as an adaptation of the congestion-controller to realize a new utility function for the user.

We now study the local stability of the system described by the set of differential equations in (1) and (3). Local stability results provide design rules for choosing parameter values in the congestion-control algorithm. The starting point of our analysis is the linearization of the system given in (1)-(3) about the equilibrium point. Let $x_r(t) = \hat{x}_r + y_r(t)$ and $\beta_r(t) = \hat{\beta}_r + z_r(t)$ where $\hat{x}_r$ and $\hat{\beta}_r$ denote equilibrium values of $x_r(t)$ and $\beta_r(t)$, respectively. As seen from the set of equations, the static relationship between these equilibrium values are given as;

$$w_r = \frac{\hat{\beta}_r \hat{x}_r}{\tilde{U}'_r(\hat{x}_r)} \sum_{l \in r} p_l \left(\sum_{j: l \in j} \hat{x}_j \right),$$

$$\hat{\beta}_r = \begin{cases} \beta_{\max}, & \hat{x}_r \leq x_{ref} \\ f_r(\hat{x}_r), & \hat{x}_r > x_{ref}, \end{cases}$$

We note that the case when even one $\hat{x}_j = x_{ref}$ is not amenable to local stability analysis since the adaptation function is not differentiable at the point $\hat{x}_r = x_{ref}$. We will simulation to consider this scenario.

In general, in steady state, some of the users will have an equilibrium rate higher than the threshold value and the rest of the users will have an equilibrium rate lower than the threshold value. Let $\mathcal{R}_1$ be the set of users who have an equilibrium rate smaller than the threshold value (i.e. $\hat{x}_r < x_{ref}$ for all $r \in \mathcal{R}_1$) and let $\mathcal{R}_2$ denote the set of users who have an equilibrium rate greater than the threshold (i.e., $\hat{x}_r > x_{ref}$ for all $r \in \mathcal{R}_2$). We will assume that $\mathcal{R}_1 \cup \mathcal{R}_2 = \mathcal{R}$.

Denote

$$p_l \left(\sum_{j: l \in j} \hat{x}_j \right) \text{ by } \hat{p}_l \text{ and } \frac{\partial p_l}{\partial x_k} \left(\sum_{j: l \in j} x_j \right) \Big|_{x_j = \hat{x}_j} \text{ by } \hat{p}_l^x.$$

Linearizing the system about the equilibrium point, we get

$$\dot{y}_r(t) = -\kappa_r \left(\frac{\hat{\beta}_r \hat{x}_r}{\tilde{U}'_r(\hat{x}_r)} \sum_{l \in r} \hat{p}_l^x \sum_{j: l \in j} y_j(t - d_1(j, l) - d_2(r, l)) \right. \\ \left. + \frac{\hat{\beta}_r}{\tilde{U}'_r(\hat{x}_r)} \sum_{l \in r} \hat{p}_l y_r(t - T_r) + \frac{\hat{x}_r}{\tilde{U}'_r(\hat{x}_r)} \sum_{l \in r} \hat{p}_l z_r(t) \right. \\ \left. - \frac{\tilde{U}''_r(\hat{x}_r) \hat{\beta}_r \hat{x}_r}{(\tilde{U}'_r(\hat{x}_r))^2} \sum_{l \in r} \hat{p}_l y_r(t) \right), \quad (4)$$

and

$$\dot{z}_r(t) = \gamma_r (f'_r(\hat{x}_r) y_r(t) - z_r(t)). \quad (5)$$

Note that $\frac{\tilde{U}''_r(\hat{x}_r)}{(\tilde{U}'_r(\hat{x}_r))^2} < 0$ when $U_r(x_r) = \frac{-1}{x_r^{\alpha_r}}$. Let $Y_r(s)$ denote the Laplace transform of $y_r(t)$ and $Z_r(s)$ denote the Laplace transform of $z_r(t)$. Taking the Laplace transform of the set of differential equations in (4) - (5), we obtain

$$sY_r(s) = -\kappa_r \left(\frac{\hat{\beta}_r \hat{x}_r}{\tilde{U}'_r(\hat{x}_r)} \sum_{l \in r} \hat{p}_l^x \sum_{j: l \in j} e^{-s(T_r - d_1(r, l) - d_1(j, l))} Y_j(s) \right. \\ \left. + \frac{\hat{\beta}_r}{\tilde{U}'_r(\hat{x}_r)} \sum_{l \in r} \hat{p}_l e^{-sT_r} Y_r(s) + \frac{\hat{x}_r}{\tilde{U}'_r(\hat{x}_r)} \sum_{l \in r} \hat{p}_l Z_r(s) \right. \\ \left. - \frac{\tilde{U}''_r(\hat{x}_r)}{(\tilde{U}'_r(\hat{x}_r))^2} \hat{\beta}_r \hat{x}_r \sum_{l \in r} \hat{p}_l Y_r(s) \right),$$

and

$$Z_r(s) = \frac{\gamma_r}{s + \gamma_r} f'_r(\hat{x}_r) Y_r(s).$$

Define $Y(s) := [Y_1(s), Y_2(s), \dots, Y_{|\mathcal{R}|}(s)]^T$. We can now write the Laplace transform of the linearized version of (1) and (3) in matrix form as

$$Y(s) = -H(s)Y(s), \quad (6)$$

where,

$$H(s) \triangleq$$

$$\text{diag} \left\{ \frac{e^{-sT_r}}{(s + \phi_r - \frac{\gamma_r}{s + \gamma_r} \zeta_r)} \frac{\kappa_r \hat{\beta}_r \hat{x}_r}{\tilde{U}'_r(\hat{x}_r)} \right\} [\bar{M}(s) + W X^{-2}],$$

$$\phi_r \triangleq -\kappa_r \frac{\tilde{U}''_r(\hat{x}_r)}{(\tilde{U}'_r(\hat{x}_r))^2} \hat{\beta}_r \hat{x}_r \sum_{l \in r} \hat{p}_l, \quad (7)$$

$$\zeta_r \triangleq -\kappa_r \frac{f'_r(\hat{x}_r)}{\tilde{U}'_r(\hat{x}_r)} \hat{x}_r \sum_{l \in r} \hat{p}_l 1_{\{r \in \mathcal{R}_2\}}, \quad (8)$$

$$\bar{M}_{jr}(s) \triangleq \sum_{l \in j \cap r} \hat{p}_l^x e^{-s(d_1(r, l) - d_1(j, l))},$$

$$W \triangleq \text{diag} \left\{ \hat{x}_r \sum_{l \in r} \hat{p}_l \right\} \text{ and } X \triangleq \text{diag} \{ \hat{x}_r \}.$$

Note that, ϕ_r and ζ_r are non-negative quantities given the definitions of the functions $\tilde{U}(x)$ and $f(x)$. We now state the main result of the paper.

Theorem 2.2 *The system described by the set of equations (1) and (3) is locally asymptotically stable if for all $r \in \mathcal{R}$*

$$f_r(\hat{x}_r) \geq \tilde{U}'_r(\hat{x}_r) \quad \text{and} \quad (9)$$

$$\frac{\kappa_r \hat{\beta}_r}{\tilde{U}'_r(\hat{x}_r)} \left(\sum_{l \in r} \hat{p}_l + \sum_{l \in r} \hat{p}_l^x \sum_{j: l \in j} \hat{x}_j \right) < \frac{\pi}{2T_{max}}. \quad (10)$$

Outline of the Proof: We will show that if conditions (9) and (10) are satisfied, then the eigenvalues of $H(jw)$ cannot encircle -1 . Hence the result follows from the Generalized Nyquist criterion. Let

$$L \triangleq \text{diag} \left\{ \frac{e^{-jwT_r}}{jw + \phi_r - \frac{\gamma_r}{jw + \gamma_r} \zeta_r} \right\}$$

$$Q \triangleq \text{diag} \left\{ \sqrt{\frac{\kappa_r \hat{\beta}_r \hat{x}_r}{\tilde{U}'_r(\hat{x}_r)}} [\bar{M}(jw) + W X^{-2}] \text{diag} \left\{ \sqrt{\frac{\kappa_r \hat{\beta}_r \hat{x}_r}{\tilde{U}'_r(\hat{x}_r)}} \right\} \right\}$$

Then,

$$\begin{aligned} \sigma(H(jw)) &= \sigma(QL) \\ &\subset \rho(Q) \text{Co}\{0 \cup \{l_i\}\} \end{aligned} \quad (11)$$

since L is a diagonal matrix and $Q = Q^* > 0$ [5]. Also the imaginary part of l_r can be written as

$$\begin{aligned} \text{Im}\{l_r\} &= -\frac{1}{\Lambda} \left(\sin(wT_r) \left(\phi_r - \frac{\gamma_r^2}{\gamma_r^2 + w^2} \zeta_r \right) + \right. \\ &\quad \left. w \cos(wT_r) \left(1 + \frac{\gamma_r}{\gamma_r^2 + w^2} \zeta_r \right) \right) \end{aligned} \quad (12)$$

where,

$$\Lambda = \left(\phi_r - \frac{\gamma_r^2}{\gamma_r^2 + w^2} \zeta_r \right)^2 + w^2 \left(1 + \frac{\gamma_r}{\gamma_r^2 + w^2} \zeta_r \right)^2.$$

From (9) and Lemma 2.1, we have $\phi_r \geq \zeta_r$. Therefore, for $w < \frac{\pi}{2T_r}$ the imaginary part of l_r given in (12) is always negative. Consequently, the imaginary part of the convex hull cannot be zero for $w < \frac{\pi}{2T_{max}}$, which means that the eigenvalues of $H(jw)$ cannot cross real axis for $w < \frac{\pi}{2T_{max}}$. On the other hand, for $w \geq \frac{\pi}{2T_{max}}$, we can write the following inequality

$$|\text{Re}\{l_r\}| \leq |l_r| = \frac{1}{\sqrt{\Lambda}} \leq \frac{1}{w} \leq \frac{2T_{max}}{\pi}. \quad (13)$$

This implies that for $w \geq \frac{\pi}{2T_{max}}$, the real part of the convex hull cannot be smaller than $-\frac{2T_{max}}{\pi}$. Let $\rho(Q)$ denote the

spectral radius of the matrix Q . Using the fact that the spectral radius of a matrix is bounded by its row sum and (10) we can write

$$\begin{aligned} \rho(Q) &= \rho(\text{diag} \left\{ \frac{\kappa_r \hat{\beta}_r}{\tilde{U}'_r(\hat{x}_r)} \right\} [M(jw) + W X^{-2}] X) \\ &\leq \frac{\kappa_r \hat{\beta}_r}{\tilde{U}'_r(\hat{x}_r)} \left(\sum_{l \in r} \hat{p}_l + \sum_{l \in r} \hat{p}_l^x \sum_{j: l \in j} \hat{x}_j \right) \\ &< \frac{\pi}{2T_{max}}. \end{aligned} \quad (14)$$

From (11), (13), and (14), the real part of any eigenvalue of $H(jw)$ cannot be smaller than -1 for $w \geq \frac{\pi}{2T_{max}}$. Combining with the result we obtained for $w < \frac{\pi}{2T_{max}}$, we can say that the eigenloci of $H(jw)$ cannot enclose -1 , which concludes the proof. ■

Lemma 2.1 *Given the definitions of ϕ_r and ζ_r , in (7),(8), $\phi_r \geq \zeta_r$ if*

$$f_r(\hat{x}_r) \geq \tilde{U}'_r(\hat{x}_r).$$

Proof: Due to space constraints, we omit the proof in this paper. ■

Now we state a corollary for TCP-like congestion controllers with the utility function $\frac{-1}{x_r}$.

Corollary 2.3 *The system described by the equations*

$$\begin{aligned} \dot{x}_r(t) &= \kappa_r \left(w_r - \beta_r(t) x_r(t) x_r(t - T_r) \sum_{l \in r} p_l \left(\sum_{j \in l} x_j(t - d_{lr}) \right) \right), \\ \dot{\beta}_r(t) &= \begin{cases} \gamma_r (\beta_{max} - \beta_r(t)), & x_r(t) \leq x_{ref} \\ \gamma_r \left(\frac{\beta_{max} x_{ref}}{x_r(t)} - \beta_r(t) \right), & x_r(t) > x_{ref} \end{cases} \end{aligned}$$

is locally asymptotically stable if

$$\kappa_i \hat{\beta}_i \hat{x}_i \left(\sum_{l \in i} \hat{p}_l + \sum_{l \in i} \hat{p}_l^x \sum_{r \in l} \hat{x}_r \right) < \frac{\pi}{2T_{max}}, \quad \forall r \in \mathcal{R}.$$

3 Simulations

In this section, we provide simulation results for the fluid model and the packet level implementation of the congestion controller studied in Section 2. We use the software package MATLAB for fluid-model simulations and the software package ns-2 for the packet model simulations. We also show simulations that indicate that the system is stable independent of the initial conditions. In all simulations we consider a single link accessed by TCP users. For the fluid model, we assume an utility function of $\frac{-1}{x}$, where x is the flow rate, to approximate the TCP congestion controller [3].

Fluid-Model Simulation: We first consider a link of capacity 1000 units with 5 users traversing the link. The round trip delay of each user is chosen uniformly between $[0, 4]$ time units. The values of the parameters are chosen according to Corollary 2.3. We choose $\kappa_r = 0.02$, $w_r = 1$, and $\beta_{\max} = 0.5$ for all users. The users start their transmission randomly between 0 and 1000 time units. The marking function at the link is chosen to be $\frac{(\lambda - C)^+}{\lambda}$, where λ is the total arrival rate at the link and C is the capacity of the link. The evolution of the flow rates of the users with β adaptation is shown in Figure 2. We can see that the flow-rates of all users converge to a small neighborhood around their equilibrium values when β is adapted. The evolution of the multiplicative decrease parameter β_r is shown in Figure 3. Once again we can see that the β_r values converge (around a small neighborhood) for all users.

Figure 2: Evolution of the flow rates of all users with time.

Figure 3: Evolution of the multiplicative decrease factor of all users with time.

In Figure 4, we compare the total utilization at the link when β_r is adapted to the total utilization at the link when β is not adapted (i.e., fixed at 0.50). Note that the total utilization at the link for the adaptive case is higher than the total utilization when no adaptation is done by the users. We can see the effect of a conservative decrease factor in Figure 4. This difference increases as the capacity of the link is increased to a more realistic value. We now present a simulation in

which the equilibrium value of the flow rates of each user equals the threshold x_{ref} . We set the capacity of the link to be 100 units while we maintain x_{ref} at 20 units. The evolution of the flow-rates is shown in Figure 5. We can see that the flow-rates of all users converges to the equilibrium point. Figure 6 shows the utilization at the link when adaptation is employed and compares it to the case when no adaptation is employed by the users. We can see that in this case, the evolution of the utilization is identical as required by the model (note the scale of the y -axis). Therefore, when the equilibrium rates are below x_{ref} the user behavior is identical to the case with no adaptation.

Figure 4: Magnified total utilization at the link.

Figure 5: Evolution of the flow-rates of all users with time when the equilibrium flow-rate equals the threshold value.

Packet level simulations: We consider a single link of capacity 100 Mbps. To provide feedback, we employ an Active Queue Management (AQM) scheme called the Adaptive Virtual Queue (AVQ) algorithm [3, 6]. One can also employ other AQM schemes like RED [7], REM [8] or PI [9]. In the first scenario, we consider a single user with a round trip delay of 100ms. We set x_{ref} to be equal to 6 Mbps and each user employs the TCP-Reno congestion control protocol. Figure 7 compares the evolution of the window size when the multiplicative decrease factor is adaptively set against the conventional TCP algorithm. We can

Figure 6: Evolution of the total utilization with time at the link when the equilibrium flow-rate equals the threshold value.

see that due to the adaptation, the TCP algorithm is able to utilize the link more efficiently than the conventional TCP congestion-control protocol. This difference is magnified as the capacity of the link increases.

Figure 7: Window Size(in packets) vs time(in seconds).

In the second scenario, we consider 5 users traversing the link. The round trip delay of each user is set at 100 msec. The total utilization of the link averaged over 0.5 sec is shown in Figure 8. We once again observe that in the adaptive case, the link utilization is higher when compared to the conventional non-adaptive scenario.

4 Conclusions and Future work

In this paper we use fluid-models to study the adaptation of the multiplicative decrease parameter in congestion-controllers to achieve higher utilizations in high bandwidth links. We consider arbitrary congestion-controllers and arbitrary adaptation functions. We then study the local stability of such adaptive congestion-controllers. We show through simulations that such adaptation indeed increases the throughput of the congestion-controller and leads to

Figure 8: Link Utilization vs Time for 5 users.

high utilizations at the link.

The adaptation of the multiplicative decrease parameter comes into effect only when the equilibrium rates of the user exceed a specific threshold. A high value of this threshold will reduce the utilization of high bandwidth links while a low value might give such controllers unfair advantage over non-adaptive controllers (say, current versions of TCP) on low bandwidth links. Choice of this threshold is a subject of future research. Global stability of such schemes is another direction that has to be explored.

References

- [1] F. P. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [2] S. H. Low and D. E. Lapsley, "Optimization flow control I: Basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, pp. 861–875, December 1999.
- [3] S. Kunniyur and R. Srikant, "End-to-end congestion control: utility functions, random losses and ECN marks," in *Proceedings of INFOCOM 2000*, Tel Aviv, Israel, March 2000. Also to appear in *IEEE/ACM Transactions on Networking*, 2003.
- [4] S. Floyd, "Highspeed TCP for large congestion windows," Internet draft draft-floyd-tcp-highspeed-00.txt, Preprint, June 2002.
- [5] G. Vinnicombe, "On the stability of end-to-end congestion control for the Internet," *Technical Report, CUED/F-INFENG/TR.398*, University of Cambridge, UK, December 2000.
- [6] S. Kunniyur and R. Srikant, "Analysis and design of an adaptive virtual queue (AVQ) algorithm for active queue management," in *Proceedings of SIGCOMM 2001*, San Diego, CA, August 2001. Also to appear in *IEEE/ACM Transactions on Networking*.
- [7] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, August 1993.
- [8] S. Athuraliya, V. H. Li, S. H. Low, and Q. Yin, "REM: Active queue management," *IEEE Network*, pp. 48–53, June 2001.
- [9] C. Hollot, V. Misra, D. Towse, and W. Gong, "On designing improved controllers for AQM routers supporting TCP flows," in *Proceedings of INFOCOM 2001*, Anchorage, Alaska, April 2001.