The 4th Asian Control Conference
September 25-27, 2002
Singapore

Analysis of a Flow Control Scheme for Rate Adjustment by Managing Inflows

Lachlan Andrew, Stephen Hanly and Rami Mukhtar
ARC Special Research Center for Ultra-Broadband Information Networks
Department of Electrical and Electronic Engineering
The University of Melbourne, Victoria 3010, Australia

Abstract

In this paper we study the stability of a recently proposed flow control algorithm for fair bandwidth sharing of a bottleneck access link. We present necessary conditions for stability for an arbitrary set of propagation delays. The analysis presented also facilitates proper configuration of the algorithms parameters.

1 Introduction

Due to economies of scale, network backbones have grown in speed and capacity much more rapidly than access networks. Access links will increasingly serve intranets. Each client would expect a fair share of the link capacity. Moreover, it is often known *a priori* that an access link is likely to be the bottleneck for most connections it carries.

Existing Internet congestion control mechanisms [1] do not attempt to optimize the use of individual bottleneck links, and are unfair to flows with disparate round trip times [2].

If there is a significant chance of brief but large increases in the capacity of a bottleneck link, such as when good radio propagation conditions exist, then the throughput can be significantly increased by maintaining sufficient buffer occupancy. This motivated the algorithm presented in [3], with the objectives of (a) maintaining a fixed buffer at the bottleneck, and (b) allocating bandwidth fairly to each flow independent of round trip time. In the present paper we provide preliminary results on the characterisation of the stability region of this algorithm.

The most common variant of the Transmission Control Protocol (TCP) used today, based on [1, 4], focuses on preventing congestion, with little attempt to optimize the performance of the bottleneck link.

TCP Vegas [5, 6, 7] keeps constant the amount of data each flow has in the network. This approach requires an estimate of the propagation time devoid of queueing delays. However, the queueing increases linearly with the number of flows. This makes the estimation increasingly difficult as the number of active flows increases, and causes a bias against new flows.

Much subsequent work has focused on optimizing the performance of the entire network [8, 9, 10]. Performance and stability issues have been investigated in [11, 9], but with limited attention to the effect of propagation delays.

The need to maintain a target queue at the bottleneck link whilst providing fair capacity allocation to each flow has independently been pointed out in [12]. That paper proposes a rate based controller, centralised at a bottleneck in an ATM network. Unlike the algorithm of [3], that of [12] requires per flow information, making it less suitable for the Internet. Although motivated by bottleneck links, the algorithm of [3] remains in the *end-to-end* Internet framework. Specifically, end nodes require only aggregate (as opposed to flow by flow) congestion information from the network, and do not rely on estimates of propagation delays.

TCP has a receiver window which may be used to improve the performance of a low bandwidth access link. In [13, 14, 15], feedback from the access router is used to set the TCP receiver window. This can increase utilization, improve fairness and provide the ability to prioritize between flows. However, [14] also suffers the weaknesses mentioned for [5, 6, 7].

The algorithm of [3] achieves the objectives of [12] using a idea similar to those of [14]: that of controlling the receiver window.

2 The Algorithm

Consider k flows of data packets share a single bottleneck link with output rate μ_c. Each flow, i, has a sending node, S_i, and a receiving node R_i (see Figure 1). Let $d_i = d_{if} + d_{ir}$ denote the total transmission delay, including propagation and all queueing delays, except that at the access router. Further assume each d_i is constant.

Each sending node has a sliding window of $w_i(t)$ packets in the network at time t. The algorithm of [3] calculates $w_i(t)$, for all i, in a decentralized way, such that each flow receives an equal rate, μ_c/k, and the equilibrium buffer occupancy of the access router, $q(t)$, can be controlled as discussed below.

The algorithm is best implemented in the receiving node, as that is where congestion signals from the access router

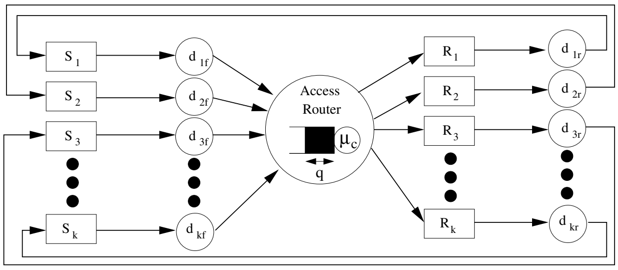

Figure 1: Model of k flows sharing a bottleneck access link

are received. However, [3] shows how it can be incorporated into TCP. In the present paper, we assume that it is implemented at the receiver, and that there is no other flow control.

2.1 Access Router Agent

A software agent in the access router samples the queue length, q, at regular intervals. A convex, monotonic increasing function of q, $p(q)$, is evaluated, and the value passed to each receiver. In this paper, we use the affine function

$$p(q) = \frac{bq - a}{\mu_c}, \tag{1}$$

where b determines how sensitive the bottleneck queue size is to the number of flows. The parameters a and b control the equilibrium mean queue size, q^*. As will be seen in Section 4, as b increases with a/b fixed, $q^* \to a/b$.

2.2 Window Update Algorithm

Each sender-receiver pair, i, operates sliding window flow control. They attempt to maintain a total of w_i packets and acknowledgements in flight at any time.

For simplicity, we assume all packets are of equal length. Each time receiver i receives a packet, compute a new value for w_i as follows.

Let $t_i(k)$ denote the time instant when the kth packet is received by user i. At time $t_i(k)$ the window is updated as

$$\begin{aligned} w_i(t_i(k)) \leftarrow{} & w_i(t_i(k-1)) + \\ & \left[\tau - p(q(t_i(k)))\,\tilde{\mu}_i(t_i(k))\right]\big(t_i(k) - t_i(k-1)\big), \end{aligned} \tag{2}$$

where τ is a constant and $\tilde{\mu}_i$ is an estimate of the received rate. The term in τ tries to increase the window at a constant rate, while the term in $\tilde{\mu}_i$ reduces it at a rate which increases with the occupancy of the queue and with the proportion of traffic due to flow i. In this paper, the current received rate, $\tilde{\mu}_i$, is estimated using a sliding window averaging function,

$$\tilde{\mu}_i(t_i(k)) = \frac{\alpha}{t_i(k) - t_i(k-\alpha)},$$

where the integer α is a smoothing factor.

The *prime start* algorithm of [3] is used to increase $w_i(t)$ rapidly from $w_i(0) = 0$ to a value which achieves full link utilisation.

3 Analytic model

In order to determine the region of convergence of the proposed algorithm, consider the following fluid flow model. Let B_i be the occupancy of the queue consisting of packets from source i, so that

$$q(t) = \sum_{i=1}^{k} B_i(t). \tag{3}$$

As long as it does not become zero, $B_i(t)$ changes with the flow into and the flow out of the buffer, as affected by the change in the window size. With the change in window size given by the fluid equivalent of (2), write

$$\begin{aligned} f_i(t) &= \left(\left.\frac{dw_i}{dt}\right|_{t-d_i} + \mu_i(t-d_i) \right) - \mu_i(t) \\ &= \tau - p(q(t-d_i))\mu_i(t-d_i) + \mu_i(t-d_i) - \mu_i(t). \end{aligned}$$

Then

$$\frac{dB_i}{dt} = \begin{cases} f_i(t) & \text{if } B_i > 0 \\ \max(0, f_i(t)) & \text{if } B_i = 0 \end{cases}. \tag{4}$$

Combining (3) and (4) gives the following expression for the total rate of change of queue occupancy:

$$\frac{dq}{dt} = |\kappa(t)|\,\tau + \left(\sum_{i\in\kappa(t)} (1 - p(q(t-d_i)))\mu_i(t-d_i) \right) - \mu_c, \tag{5}$$

where $\kappa(t)$ is the set of users for which $B_i(t) > 0$ or $dB_i/dt > 0$, and $|\,\boldsymbol{.}\,|$ is the number of elements in the set.

The flow rate out of the buffer can be approximated by

$$\mu_i(t) = \frac{B_i(t)}{q(t)}\mu_c. \tag{6}$$

Differentiating (6) and substituting for $B_i(t)$ gives

$$\frac{d\mu_i}{dt} = \frac{1}{q(t)} \left(\frac{dB_i}{dt}\mu_c - \mu_i(t)\frac{dq}{dt} \right). \tag{7}$$

Letting

$$\begin{aligned} g_i(t) = \frac{1}{q(t)} \Bigg[& (\tau + (1 - p(q(t-d_i)))\mu_i(t-d_i))\mu_c - \\ & \mu_i(t) \left(|\kappa(t)|\,\tau + \sum_{j\in\kappa(t)} (1 - p(q(t-d_j))\mu_j(t-d_j) \right) \Bigg] \end{aligned} \tag{8}$$

and substituting for (4) and (5) in (7), and canceling the terms in $\mu_i(t)\mu_c$ gives

$$\frac{d\mu_i}{dt} = \begin{cases} g_i(t) & \text{if } B_i > 0 \\ \max(0, g_i(t)) & \text{if } B_i = 0. \end{cases} \tag{9a}$$

Substituting the affine cost function, (1), into (8) and (5) gives

$$\begin{aligned} g_i(t) = \frac{1}{q(t)} \Bigg[& \left(\tau + \left(1 - \frac{bq(t-d_i) - a}{\mu_c} \right) \mu_i(t-d_i) \right) \mu_c - \\ \mu_i(t) & \left(|\kappa(t)|\,\tau + \sum_{j\in\kappa(t)} \left(1 - \frac{bq(t-d_j) - a}{\mu_c} \right) \mu_j(t-d_j) \right) \Bigg] \end{aligned} \tag{9b}$$

and

$$\frac{dq}{dt} = |\kappa(t)|\,\tau + \left(\sum_{i\in\kappa(t)} \left(1\text{-}\frac{bq(t\text{-}d_i)\text{-}a}{\mu_c}\right)\mu_i(t\text{-}d_i)\right) - \mu_c. \tag{9c}$$

4 Conditions for stability of linear systems

The coupled system of equations (9) has a fixed point at

$$q^* = p^{-1}\left(\frac{k\tau}{\mu_c}\right) = \frac{k\tau + a}{b}, \qquad \mu_i^* = \frac{\mu_c}{k}. \tag{10}$$

This is the unique equilibrium in which all users have access to the network, $\mu_i > 0$ for all i. We would like to characterize conditions under which this point is stable.

Equations (9) may be stable for a given b, but not for a smaller $b' < b$. Thus we focus on conditions on b_0 such that (9) is stable for all $b < b_0$. In the following, we fix k, τ and a, but let b vary. Denote the solution of (10) by $q^*(b)$. When it is clear from context, we will continue to denote the fixed point by q^*.

The nonlinearity of (9) makes it difficult to analyze. Instead, let us analyze some simpler, but related linear equations, which give considerable insight into the complete system. First, consider the case that the queue size, q, has converged.

4.1 Constant queue size

When $q(t) = q^*$ is constant, and all flows have converged to nonzero rates, the second term of (7) vanishes. Thus, the equations for the μ_is, (9b), decouple to form the linear constant coefficient delay differential equations

$$\begin{aligned}\frac{d\mu_i}{dt} &= \frac{\tau\mu_c}{q^*} + \frac{(1-p(q^*))\mu_c}{q^*}\mu_i(t-d_i)) - \frac{\mu_c}{q^*}\mu_i(t) \\ &= \frac{\tau\mu_c}{q^*} + \frac{\mu_c - k\tau}{q^*}\mu_i(t-d_i)) - \frac{\mu_c}{q^*}\mu_i(t). \end{aligned} \tag{11}$$

Proposition 1 *A necessary and sufficient condition for equation (11) to converge for all* $b < b_0$*, is that*

$$k\tau \le 2\mu_c, \tag{12}$$

or

$$d_i \le \frac{1}{b_0}\cos^{-1}\left(\frac{\mu_c}{\mu_c - k\tau}\right)\frac{k\tau + a}{\sqrt{k^2\tau^2 - 2\mu_c k\tau}}. \tag{13}$$

Proof: The stability of equation (11) for any value of b, is equivalent to the poles of the unilateral Laplace transform of the unforced equation,

$$\begin{aligned} M_i(s) &= \frac{\mu_i(0^-)}{s + (\mu_c/q^*) + (\mu_c - k\tau/q^*)e^{-sd_i}} \\ &= \frac{\mu_i(0^-)}{A(s) + C(s)e^{-sd_i}} \end{aligned} \tag{14}$$

all being in the left half plane. The stability analysis now follows that of [16]. For sufficiently small d_i this system is stable, as it is a retarded system. Let $d_i(b)$ be the smallest positive delay such that $s = j\omega$, $\omega \in \mathbb{R}$, is a pole. By the continuity of poles, the system will be stable for all $d_i < d_i(b)$. Since poles occur in complex conjugate pairs, this is further equivalent to $A(j\omega)A(-j\omega) = C(j\omega)C(-j\omega)$ [16], or equivalently

$$\omega^2 = \frac{k^2\tau^2 - 2\mu_c k\tau}{(q^*)^2}.$$

This has no non-zero real solution for $k\tau \le 2\mu_c$, and therefore (11) is always stable. This establishes the sufficiency of condition (12). It remains to show that, if $k\tau > 2\mu_c$ then (13) is necessary and sufficient.

When $k\tau > 2\mu_c$, $d_i(b)$ can be found by equating the real part of the denominator of (14) to zero:

$$\frac{\mu_c}{q^*} - \frac{\mu_c - k\tau}{q^*}\cos\left(\frac{d_i(b)\sqrt{k^2\tau^2 - 2\mu_c k\tau}}{q^*}\right) = 0$$

or equivalently

$$\cos\left(\frac{d_i(b)\sqrt{k^2\tau^2 - 2\mu_c k\tau}}{q^*}\right) = \frac{\mu_c}{\mu_c - k\tau}.$$

Taking the smallest positive argument argument of $\cos(\cdot)$,

$$d_i(b) = \cos^{-1}\left(\frac{\mu_c}{\mu_c - k\tau}\right)\frac{q^*(b)}{\sqrt{k^2\tau^2 - 2\mu_c k\tau}}.$$

Note that $d_i(b)$ increases to infinity as b decreases to zero. Define

$$\beta \equiv \cos^{-1}\left(\frac{\mu_c}{\mu_c - k\tau}\right)\frac{k\tau + a}{\sqrt{k^2\tau^2 - 2\mu_c k\tau}}.$$

Necessity: if $d_i > \beta/b_0$ then there exists a $b < b_0$ such that $d_i(b) = d_i$, and hence (11) is not stable. *Sufficiency:* if $d_i < \beta/b_0$ then for all $b < b_0$, $d_i < \beta/b$, and hence (11) is stable. ■

Note that the sufficient condition $k\tau < 2\mu_c$ implies that the system is always stable as long as the total rate at which packets are introduced due to stations expanding their windows does not exceed twice the bottleneck packet rate.

Note that when $k\tau = 2\mu_c$, the second constraint, (13), is merely $b_0 \le \infty$. To see how it behaves as $k\tau$ increases beyond $2\mu_c$, let

$$F(x) \equiv \frac{1}{d_i}\cos^{-1}\left(\frac{\mu_c}{\mu_c - x}\right)\frac{x + a}{\sqrt{x^2 - 2\mu_c x}}.$$

It is not difficult to verify that $F(\infty) = \pi/(2d_i)$ and $F'(x) < 0$ for all $x > 2\mu_c$. Thus, for $k\tau > 2\mu_c$: the larger $k\tau$, the tighter will be the second bound, and the tightest it can be is $b_0 < \pi/(2d_i)$.

A weaker sufficient condition for the stability of (11) is given by the following proposition:

Proposition 2 *A sufficient condition for (11) to be stable is that* $b < \pi/(2d_i)$.

Proof: If $k\tau \le \mu_c$ then by Proposition 1 the system (11) is stable. It remains to consider the case when $k\tau > \mu_c$. In this case, by Proposition 1 it is sufficient that

$$d_i b < \cos^{-1}\left(\frac{\mu_c}{\mu_c - k\tau}\right)\frac{k\tau + a}{\sqrt{k^2\tau^2 - 2\mu_c k\tau}}.$$

Since $k\tau > 2\mu_c$, it follows that

$$\frac{\pi}{2} < \cos^{-1}\left(\frac{\mu_c}{\mu_c - k\tau}\right) < \pi.$$

Also

$$\frac{k\tau + a}{\sqrt{k^2\tau^2 - 2\mu_c k\tau}} > 1.$$

Thus it is sufficient that $d_i < \pi/(2b)$, as required. ■

4.2 Constant flow rates

Under the assumption that $\mu_i(t)$ have converged to $\mu_i^* = \mu_c/k$ for all i, the queue dynamics, (9c), become

$$\frac{dq}{dt} = k\tau - \frac{b}{k}\sum_{i=1}^{k} q(t - d_i) + a\mu_c. \qquad (15)$$

Unlike the previous section, it is possible for the linear equation (15) to be unstable, yet the nonlinear, coupled system (9) be stable. The reason is that the evolution of the k equations (8) are not independent of the evolution of $q(t)$, and it is impossible for $q(t)$ to vary, whilst the $\mu_i(t)$ remain fixed. In fact, $\mu_i(t)$ can act as stabilizing controllers, preventing the coupled system from going unstable, even when (15) is itself unstable. Thus we cannot derive necessary conditions for the coupled system, under the above assumption. Nevertheless, it will prove insightful to obtain necessary and sufficient conditions on the stability of (15).

The convergence of $q(t)$ under (15) to any value implies that the poles of the unilateral Laplace transform,

$$Q(s) = \frac{q(0^-)}{s + \frac{b}{k}\sum_{i=1}^{k} e^{-sd_i}}, \qquad (16)$$

all have negative real part.

In the two flow case, a simple necessary and sufficient condition on b_0 will be derived for arbitrary delays, d_1, and d_2. After that, a procedure for obtaining a necessary and sufficient condition on b_0 in the general case of an arbitrary number of flows will be outlined, under the innocuous constraint that the delays be commensurate. We emphasize that these conditions on b_0 are necessary and sufficient for the stability of (15) for all values of b less than b_0, but not for the original system of nonlinear equations (9).

The proofs will again follow the approach described in [16]. In all cases, the first step of showing that the poles of (16) are in the left half plane for small delays is a direct result of the following lemma [17, p. 139], which applies to (16).

Lemma 1 Consider a linear constant coefficient retarded system, which without any delays would be stable. Then there exists an $\epsilon > 0$ such that if all delays, d_i, satisfy $0 < d_i < \epsilon$, then all the poles of the delayed system have negative real part.

Thus, in each case, stability can be established by finding conditions under which the poles are not purely imaginary.

4.2.1 Two users:

Consider the case of $k = 2$ users, with delays d_1, d_2. Denote $\text{sinc}(x) = \sin(\pi x)/(\pi x)$.

Proposition 3 *In the two flow scenario:*

1. *a necessary condition on* b_0 *for (15) to converge for all* $b < b_0$*, is that*
$$b_0 < \frac{1}{d_1\,\text{sinc}(d_1/(d_1+d_2))}. \qquad (17)$$
2. *If* $\tau < \mu_c$*, the condition*
$$b < \frac{1}{d_1\,\text{sinc}(d_1/(d_1+d_2))} \qquad (18)$$
is sufficient for (15) to be stable.

Proof: *Part 1*: if the first part of the proposition were false, the the system would converge for $b = 1/(d_1\,\text{sinc}(d_1/(d_1+d_2)))$. However, it is easily verified that in that case, $jb\sin(\pi d_1/(d_1+d_2))$ is an unstable pole of (16).

Part 2: For $0 < h < 1$, define $d_i(h) = d_i h$, for $i = 1, 2$. Suppose that there is a pole of (16) in the right half plane. By continuity of the poles with respect to h, there exists a minimum value of h, which we denote by h', which gives a pole on the imaginary axis, for the system with delays replaced by the $d_i(h')$ values. Note that $h' < 1$. Then there exists an $\omega_0 > 0$, such that

$$j\omega_0 + \frac{b}{2}e^{-j\omega_0 d_1 h'} + \frac{b}{2}e^{-j\omega_0 d_2 h'} = 0.$$

Equating real parts (using $\cos(x) = \cos(\pi - x)$) and imaginary parts gives

$$\omega_0(d_1 + d_2)h' = \pi,$$
$$\omega_0 = b\sin\left(\frac{\pi d_1}{d_1 + d_2}\right).$$

Combining these two equations, we obtain

$$h' = \frac{1}{bd_1\,\text{sinc}(d_1/(d_1+d_2))}.$$

But by assumption the right hand side is greater than unity, which is a contradiction. The proposition follows. ■

Note that as $d_1/d_2 \to 0$ the condition (17) becomes $b_0 < 1/d_1$, and the stability is dominated by the shortest delay, rather than the longest delay, which seems somewhat surprising.

As expected, the bound in (17) lies between the bounds obtained for the systems $d_1' = d_2' = \min(d_1, d_2)$ and $d_1'' = d_2'' = \max(d_1, d_2)$.

4.2.2 Commensurate delays: The special case of commensurate delays, in which $d_i = n_i h$, $n_i \in \mathbb{N}$, is covered in [16]. In this case,

$$Q(s) = \frac{q(0^-)}{s + \frac{b}{k}\sum_{i=1}^{k} e^{-sn_i h}}, \qquad (19)$$

Lemma 1 again shows that for sufficiently small h, all poles have negative real part. The value of h, h_θ, for which the poles of (19) first touch the imaginary axis, $s = j\omega_0$, can be found by finding roots of a polynomial, whose coefficients can be determined in k recursive steps, as described in [16]. This provides a necessary condition: $b_0 < b_0^*$ for the stability of (15), for all $b < b_0$, where b_0^* is computed by the k recursive steps. In simple cases, b_0^* can be found explicitly. This provides necessary and sufficient conditions for the stability of (15) over all such values of b.

For sufficiency only, consider the following proposition:

Proposition 4 If $b < \pi/(2d_i)$ for all i, then $q(t)$ is stable under (15).

This proposition is a simple corollary of the following lemma, and the fact that the poles of (16) are continuous functions of b.

Lemma 2 No poles of (16) can lie on the imaginary axis if $b < \pi/(2d_i)$ for all i.

Proof: Assume, with a view to obtaining a contradiction, that there is a pole at $j\omega$. This pole must satisfy

$$j\omega + \frac{b}{k}\sum_{i=1}^{k} e^{-j\omega d_i} = 0.$$

Writing $\omega' = \omega/b$, and equating real and imaginary parts gives

$$\frac{1}{k}\sum_{i=1}^{k} \cos(d_i b\omega') = 0 \qquad (20a)$$

$$\frac{1}{k}\sum_{i=1}^{k} \sin(d_i b\omega') = \omega'. \qquad (20b)$$

By (20b), ω' is the average of samples of the sine function, and hence $-1 \le \omega' \le 1$. However, by hypothesis $0 \le d_i b < \pi/2$. Thus

$$\frac{-\pi}{2} < d_i b\omega' < \frac{\pi}{2}$$

whence $\cos(d_i b\omega') > 0$ for all i. This contradicts (20a), and the lemma is proved. ∎

5 Characterization of stability

In this section, we give numerical results from the direct computation of the solutions of the coupled, nonlinear system of equations (4) in the case of $k = 2$. We attempt to characterize the region of stability, and relate the region to the theoretical results presented in the preceding section. Finally, we present numerical results from the simulation of the algorithm itself, and compare with those obtained from the theoretical models.

Figure 2 shows in grey the stability region of the coupled nonlinear system (4) for four cases: $\tau/\mu_c = 0.005$, 1, 10 and 100. These were obtained via numerical solution of (4), and a test for convergence. All graphs show the lines $bd_i = \pi/2$, $i = 1, 2$. Note that in (a) and (b) condition (12) holds, and (c), (d) show the line satisfying (13) with equality.

In all cases, the conditions $bd_i < \pi/2$, $i = 1, 2$ ensure stability. We have tested this for $k > 2$, up to $k = 8$, and taken hundreds of thousands of trials, and always obtain stability when $bd_i < \pi/2$.

Conjecture 1 *The condition $b < \pi/(2d_i)$ for all i, is a sufficient condition for stability of the entire coupled system (9).*

Note also that (13) appears to be necessary for stability when $k\tau > 2\mu_c$. Finally, note that the broken line, depicting (18), appears to characterise the stability region when $\tau = \mu_c$, noting $k = 2$. When $\tau > \mu_c$, they appear to give a reasonably good approximation, when coupled with the condition (13). Finally, Figure 3 presents preliminary evidence that the theoretical model (4) closely approximates the performance of the algorithm. This is the subject of continuing investigation.

6 Conclusion

We have shown that our algorithm performs well. It maintains a constant queue at the access router, insensitive to the number of flows, and does not bias against flows with large propagation delays. We have provided preliminary results which suggest strongly that, provided b is chosen judiciously as a function of the worst case propagation delay, then the algorithm is stable.

References

[1] V. Jacobson, “Congestion avoidance and control,” in *Proc. ACM SIGCOMM*, (Stanford, CA), 1988.

[2] S. Floyd, “Connections with multiple congested gateways in packet-switched networks, part I: One-way traffic,” *ACM Comp. Commun. Rev.*, vol. 21, pp. 30–47, Oct. 1991.

[3] S. V. Hanly, R. G. Mukhtar, and L. L. H. Andrew, “A novel algorithm for receiver-based management of low bandwidth access links,” in *(to appear)*, 2002.

[4] R. Jain, K. K. Ramakrishnan, and D.-M. Chiu, “Congestion avoidance in computer networks with a connectionless network layer,” tech. rep., Digital Equipment Corperation, Littleton, MA, 1988.

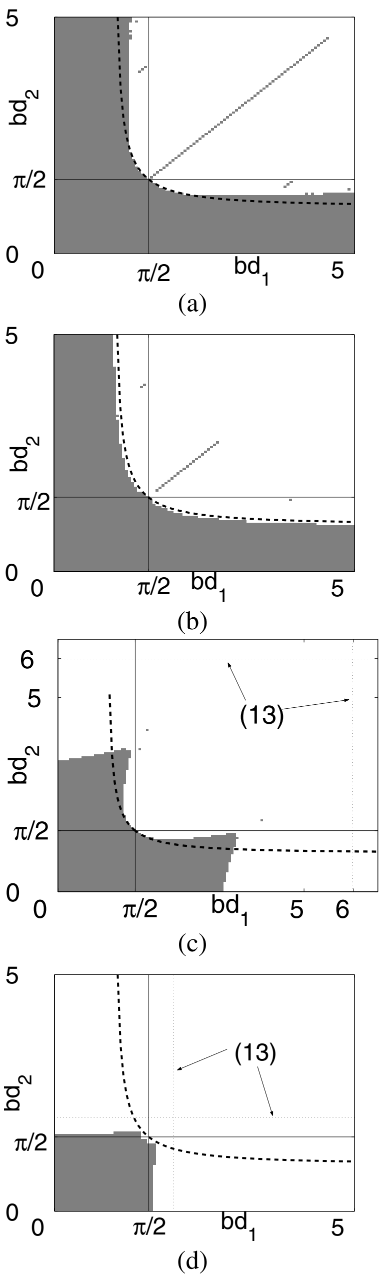

Figure 2: Stability region $\tau/\mu_c =$ (a) 0.005 (b) 1 (c) 10 (d) 100.

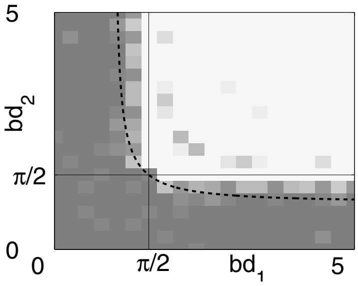

Figure 3: Algorithm's stability region, $\tau/\mu_c = 0.005$.

[5] R. Jain, "A delay-based approach for congestion avoidance in interconnected heterogeneous computer networks," *ACM Comp. Commun. Rev.*, vol. 19, pp. 56–71, Oct. 1989.

[6] L. S. Brakmo, S. W. O'Malley, and L. L. Peterson, "TCP Vegas: New techniques for congestion detection and avoidance," in *Proc. ACM SIGCOMM*, (London, UK), pp. 24–35, 1994.

[7] L. S. Brakmo and L. L. Peterson, "TCP Vegas: End to end congestion avoidance on a global internet," *IEEE J. Select. Areas Commun.*, vol. 13, pp. 1465–1480, Oct. 1995.

[8] S. H. Low and D. E. Lapsley, "Optimization flow control I: Basic algorithm and convergence," *IEEE/ACM Trans. Networking*, vol. 7, pp. 861–875, Dec. 1999.

[9] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: Shadow prices, proportional fairness and stability," *J. Op. Res. Soc.*, vol. 49, pp. 237–378, 1998.

[10] S. H. Low, "A duality model of TCP and queue management algorithms," in *Proc. ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, (Monterey, CA), 2002.

[11] F. Paganini, J. C. Doyle, and S. H. Low, "Scalable laws for stable network congestion control," in *Proc. IEEE Conf. Decision Contr. (CDC)*, (Orlando, FL), pp. 185–90, 2001.

[12] P. F. Quet, B. Ataşlar, A. Îftar, H. Özbay, S. Kalyanaraman, and T. Kang, "Rate-based flow controllers for communication networks in the presence of uncertain time-varying multiple time-delays," *Automatica*, vol. 38, pp. 917–928, 2002.

[13] J. Aweya, M. Ouellette, D. Y. Montuno, and Z. Yao, "Enhancing network performance with TCP rate control," in *Proc. IEEE Globecom*, (San Francisco, CA), pp. 1712–1718, 2000.

[14] N. T. Spring, M. Chesire, M. Berryman, V. Sahasranaman, T. Anderson, and B. N. Bershad, "Receiver based management of low bandwidth access links," in *Proc. IEEE INFOCOM*, pp. 245–254, 2000.

[15] R. Mukhtar, S. Hanly, M. V. Ivanovich, H. Vu, and P. G. Fitzpatrick, "Analysis of TCP performance over hybrid fast fixed-to-slow wireless links," in *Proc. IEEE Globecom 2001*, (San Antonio, TX), pp. 1776–1780, 2001.

[16] J. E. Marshall, H. Górecki, K. Walton, and A. Korytowski, *Time-Delay Systems: Stability and Performance Criteria with Applications*. New York, NY: Ellis Horwood, 1992.

[17] L. E. El'sgol'ts and S. B. Norkin, *Introduction to the Theory and Application of Differential Equations with Deviating Arguments*. New York, NY: Academic Press, 1973. Translated from Russian by J.L. Casti.