

End-to-End Congestion Control for the Internet: Delays and Stability

Ramesh Johari and David Tan
Statistical Laboratory, University of Cambridge

March 16, 2000

Abstract

Under the assumption that queueing delays will eventually become small relative to propagation delays, we derive stability results for a fluid flow model of end-to-end Internet congestion control. The theoretical results of the paper are intended to be decentralized and locally implemented: each end system needs knowledge only of its own roundtrip delay. Criteria for local stability and rate of convergence are completely characterized for a single resource, single user system. Stability criteria are also described for networks where all users share the same roundtrip delay. Numerical experiments investigate extensions to more general networks. Finally, through simulations, we are able to evaluate the relative importance of queueing delays and propagation delays on network stability.

Introduction

Changes in communication networks over the last five years have forced researchers to closely examine network congestion control, particularly for the Internet. However, the task has been significantly complicated by the globalization of networks: “trial-and-error” methods employed on small testbeds do not necessarily yield results scalable to a network as large as the Internet. Instead, an increasing emphasis is being placed on theoretical predictions of robust behavior, *prior* to implementation.

At the level of theoretical abstraction we will consider, a network is comprised of *users* and *resources*. Each user wishes to employ one or more resources of the network. Within this framework, our goal will be to formulate a congestion control algorithm which can remain stable under communication delays.

Any suggestion of new congestion control algorithms for the Internet must address the following basic question: will the control be implemented by the resources, or by the users? As an example of the differences between these choices, we compare the TCP/IP model of the Internet with ATM networks. In TCP/IP, the “intelligence” of the congestion control is implemented by the end users: to signal congestion, routers simply drop packets. It is left to the end systems to detect these packet drops, and act accordingly. Conversely, in ATM networks, “quality-of-service”(QoS) is a key concern, and ensuring that QoS requirements are met is a task handled primarily by the links in the network.

Which viewpoint should we adopt? Arguments have been made at an economic level regarding the fairness benefits of end system congestion control [5]: it is suggested that end users are best equipped to adjust their flows so as to maximize utility, and hence congestion control should be implemented by asking resources to convey information on incipient congestion to intelligent end systems.

However, it is not immediately clear that this type of congestion control is robust under communication delays in the network. Communication delay comprises two elements: propagation delay, and queueing delay. The propagation delay is the physical delay in transmission of data along a length of fiber (or through space, for a wireless network). The queueing delay is the delay experienced by data waiting to be served at resources within the network.

We will work from the assumption that, in the future, queueing delays will become small relative to propagation delay, as Kelly argues in [10]. Router hardware and network capacity continue to improve rapidly, reducing queueing delays; on the other hand, propagation delays are fixed by the distances between nodes on the network, and the speed of light. In [10], Kelly uses several scaling regimes at queues in a network to support the claim that queueing delays will eventually become relatively small.

Under this assumption, the choice of congestion control mechanism becomes much easier. If propagation delays are the primary source of communication delay in the network, then the choice of congestion control mechanism must be able to estimate propagation delays, and adjust the sending rate in response to these delays. The end users are naturally placed to perform this task; on the other hand, resources would be required to estimate the delays experienced by each flow passing through them—a very difficult problem in a network on the scale of the Internet.

The work of Chong et al. in [2, 3] is related to this point. In [2, 3], Chong et al. analyzed the equilibrium and stability of their proposed first-order rate-based flow control (FRFC) algorithm for a single bottleneck resource of multiple users with diverse roundtrip delays. In the FRFC algorithm, the rate allocated to each user is the difference between the observed queue length and the queue threshold, multiplied by a control gain. To maintain stability of such an algorithm, however, the resource needs knowledge of the propagation delays experienced by all users. For this reason, we will instead consider a model which implements end-to-end congestion control, in response to simple congestion indication signals from resources.

In [11], Kelly et al. proposed two complementary congestion control algorithms. The FRFC algorithm is similar to the “dual algorithm” of [11]: it relies on the resource to implement the congestion control. In this paper, we will study the “primal algorithm” of [11], where the end user implements a TCP-like rate control algorithm which responds to congestion indication signals from resources. The packet-level queueing behavior at the resources is not modeled; instead, a deterministic *fluid flow* approximation is considered.

We will represent a network by a set J of *resources*. Let a *route* r be a nonempty subset of J , and denote the set of all routes by R . We will ignore routing choices, and thus assume we can identify each user with one and only one route. Consider the primal algorithm of [11], with strictly positive gain parameters κ_r , $r \in R$:

$$\frac{d}{dt}x_r(t) = \kappa_r \left(w_r - x_r(t) \sum_{j \in r} \mu_j(t) \right), \quad r \in R \quad (1)$$

where

$$\mu_j(t) = p_j \left(\sum_{s:j \in s} x_s(t) \right), \quad j \in J. \quad (2)$$

The first equation describes the time evolution of x_r , the sending rate of user r . The second equation describes the generation of congestion indication signals at resource j , by means of a congestion indication function $p_j(y)$. We will assume $p_j(\cdot)$ is increasing, nonnegative, and not identically zero.

We may motivate the algorithm (1)-(2) as follows. Suppose that user r generates $x_r(t)$ packets at time t . If the total flow through resource j is y , we interpret $p_j(y)$ as the probability a packet at resource j receives a ‘‘mark’’—a congestion indication signal. Then $\mu_j(t)$ is the probability a packet at resource j receives a mark at time t . We assume a packet may only be marked at most once. In this case, $\sum_{j \in r} \mu_j(t)$ is the probability a packet from user r receives a mark at time t , and $x_r(t) \sum_{j \in r} \mu_j(t)$ is just the expected number of marks user r receives at time t . User r then adjusts $x_r(t)$ to reach a target number of marks per unit time w_r . (Although the preceding discussion is at the packet level, notice that (1)-(2) is a fluid flow approximation.)

In [11], the system (1)-(2) is shown to be asymptotically stable in the large, and hence convergent to a unique equilibrium point x given by:

$$x_r = \frac{w_r}{\sum_{j \in r} p_j \left(\sum_{s:j \in s} x_s \right)}. \quad (3)$$

The system (1)-(2) does not model propagation delays. We will add propagation delays as follows. Given a route r , for each resource $j \in r$ we define a *forward delay* $d_1(j, r)$, and a *return delay* $d_2(j, r)$. The forward delay is the delay incurred in communication from the user to the resource; the return delay is the delay incurred in communication from the resource back to the user. In the current Internet, each route is subject to a roundtrip delay. We model this delay by assuming each route has an associated delay D_r , such that $d_1(j, r) + d_2(j, r) = D_r$ for each $j \in r$. Consider now the following delayed difference equations analogous to the primal algorithm, where we assume that $d_1(j, r)$ and $d_2(j, r)$ are integer valued:

$$x_r[t+1] = x_r[t] + \kappa_r \left(w_r - x_r[t - D_r] \sum_{j \in r} \mu_j[t - d_2(j, r)] \right), \quad r \in R, \quad (4)$$

where

$$\mu_j[t] = p_j \left(\sum_{s:j \in s} x_s[t - d_1(j, s)] \right), \quad j \in J. \quad (5)$$

These are just discrete-time equations corresponding to the continuous-time equations (1)-(2). Suppose again that we consider $x_r[t]$ to model the number of packets generated by user r at time t . These packets experience a delay of $d_1(j, r)$ before arriving at resource j ; therefore, as before, $\mu_j[t]$ gives the probability a packet at resource j receives a mark at time t . Packets leaving resource j experience a delay of $d_2(j, r)$ before returning to user r . Thus, the total expected number of marks received by user r at time t is now given by $x_r[t - D_r] \sum_{j \in r} \mu_j[t - d_2(j, r)]$. Notice that the number

of marks received by user r at time t is proportional to the rate D_r units of time ago: this is a direct result of the fact that $d_1(j, r) + d_2(j, r) = D_r$ for all $j \in r$. As before, user r adjusts $x_r[t]$ to reach a target number of marks per unit time w_r .

The delay structure we have assumed here is very similar to the current operation of TCP [9]. In TCP, the sender transmits packets to a receiver; for every packet received, an acknowledgment is transmitted back to the sender. All packets make a round trip from sender to receiver, and back to sender. As a result, the sum of forward and return delays is fixed for all resources on a given route. This sum is the roundtrip delay D_r encountered in our previous development.

We will focus in this paper on the engineering requirement that the network remain stable under communication delays; in particular, since queueing delays are assumed to be small, the key source of instability will be the propagation delay, as captured by the roundtrip delay D_r . We primarily seek decentralized conditions for stability. Given any network with many users, the results of [11] give a global condition for stability of the lagged system. Here, we will pursue conditions which each individual user must satisfy to ensure that the system remains stable.

We will consider stability conditions of the form (approximately):

$$\kappa_r D_r < \text{route-dependent constant, } r \in R. \quad (6)$$

Such a condition says that if the product $\kappa_r D_r$ is less than some constant dependent on route topology and/or the network operating point, then the system (4)-(5) will be locally stable. In a typical network where the roundtrip delays may not be known, this condition means that *self-clocking* rate control algorithms will work well to maintain network stability. In a self-clocking algorithm, the sender uses an acknowledgement from the receiver to prompt a step forward. This gives the inverse dependence of the gain κ on the roundtrip delay D . (For more details on self-clocking algorithms, see [5] and [9].)

The type of constraint described by (6) is very desirable in a large network such as the Internet. Implementing end-to-end congestion control would be very difficult if the “route-dependent constant” involved information not available to the user. Instead, the condition (6) allows stable behavior to result without any central control. We will study the form of the route-dependent constant, and also consider what this constant implies for resource design.

The paper is organized as follows. In Section 1, we study a simplified model with one resource, and give a condition on κ which ensures the system is asymptotically stable. Next, we discuss rate of convergence for the one resource model in Section 2, completely characterizing convergence to equilibrium for the region of stability.

We next turn our attention to the network scenario. In Section 3, we develop conditions for stability of a network where all routes have the same roundtrip delay ($D_r = D$ for all $r \in R$), by linearization of the equations (4)-(5). The condition that $d_1(j, r) + d_2(j, r) = D_r$ proves to be essential in developing these stability criteria. We also suggest two extensions to this model: one where we add completely nonadaptive users, and another where we add instantly adaptive users. In Section 4, we discuss a conjecture extending our results to networks with diverse roundtrip delays. This conjecture is investigated with numerical experiments.

In Section 5, we give simulation results which add queueing and packet-level behavior to the theoretical model. In particular, we are able to further investigate how queueing delays and propagation delays affect network stability, and what it means for queueing delays to become negligible. Finally, we comment on the implications of our results for future networks in Section 6.

1 Stability criterion: one route, one resource

Consider the network depicted in Figure 1, consisting of one resource and one route. The difference equation describing this system is:

$$x[t + 1] = x[t] + \kappa(w - x[t - D]p(x[t - D])), \quad (7)$$

where $D = d_1 + d_2$. We begin by considering the linearization of (7). Let the stable point be x , where $w = xp(x)$, and suppose p is differentiable at this point with derivative p' . Then linearizing with $x[t] = x + y[t]$, we obtain:

$$y[t + 1] = y[t] - \kappa(p + xp')y[t - D], \quad (8)$$

neglecting higher order terms. Given a complex number λ , we try a solution of the form $y[t] = \lambda^t$. If such a solution exists, it is called a *normal mode*. Substituting in (8), we find that λ must satisfy the following equation—the *characteristic equation* of (8):

$$\lambda^{D+1} - \lambda^D + \kappa(p + xp') = 0. \quad (9)$$

Thus, a solution $y[t] = \lambda^t$ exists only for those λ which are roots of the characteristic equation; but such a solution converges to zero only if $|\lambda| < 1$. Because any solution of (8) is a linear combination of normal modes, a sufficient condition for the system (7) to be locally stable is that all roots of (9) have modulus less than unity. Conversely, if any one root of (9) has modulus greater than unity, then the system (7) will be unstable.

Our first theorem will state conditions which are essentially necessary and sufficient for the fixed point of the system (7) to be locally stable. The strategy of proof is as follows. We first show that the maximum modulus of the roots of equation (9) is continuous in κ . Then we show that equation (9) has all roots with modulus less than unity for κ near the origin. Finally, we look for the smallest κ such that equation (9) has a root of unit modulus.

Lemma 1 *Let $p(\lambda, \kappa)$ be a polynomial in λ , with coefficients which are continuous functions of κ (where κ may be vector valued). Then the maximum modulus of the roots λ of $p(\lambda, \kappa) = 0$ is continuous in κ .*

Proof. The roots of any polynomial are continuous functions of the coefficients [8]. If the coefficients are continuous functions of κ , this immediately implies the conclusion of the lemma. \square

Lemma 2 *For sufficiently small κ , equation (9) has all roots with modulus less than unity.*

Proof. We define the polynomial $p(\lambda, \kappa)$ by:

$$p(\lambda, \kappa) = \lambda^{D+1} - \lambda^D + \kappa a,$$

where $a = p + xp' > 0$. Then the characteristic equation (9) is $p(\lambda, \kappa) = 0$. When $\kappa = 0$, the roots are $\lambda = 0$ with multiplicity D and $\lambda = 1$ with multiplicity 1. By direct computation, we find that $\partial p / \partial \lambda$ is nonzero at $\lambda = 1$, $\kappa = 0$. Thus, we can apply the Implicit Function Theorem to find an

open interval $(-\varepsilon, \varepsilon)$ and a differentiable complex-valued function $g(\kappa)$ such that $1 = g(0)$, and $\lambda = g(\kappa)$ satisfies $p(\lambda, \kappa) = 0$ for $-\varepsilon < \kappa < \varepsilon$. For such κ , differentiating $p(\lambda, \kappa) = 0$ with respect to κ yields:

$$(D+1)g(\kappa)^D g'(\kappa) - Dg(\kappa)^{D-1} g'(\kappa) + a = 0.$$

Evaluating this through at $\kappa = 0$ yields $g'(0) = -a$. We take a Taylor series expansion of $g(\kappa)$ around $\kappa = 0$:

$$g(\kappa) = 1 - \kappa a + o(\kappa).$$

Thus, as κ increases away from zero, the root $\lambda = 1$ initially moves approximately as $1 - \kappa a$, and hence for sufficiently small κ this root will have modulus less than unity. For the root $\lambda = 0$ of multiplicity D , we note that by taking κ even smaller if necessary, we can appeal to Lemma 1 to ensure that these D roots remain of modulus less than unity as κ increases away from zero. Thus, for sufficiently small κ , equation (9) has all roots of modulus less than unity. \square

Theorem 3 *The system (7) is locally stable if:*

$$\kappa(p + xp') < 2 \sin\left(\frac{\pi}{2(2D+1)}\right),$$

and unstable if:

$$\kappa(p + xp') > 2 \sin\left(\frac{\pi}{2(2D+1)}\right).$$

Proof. We can easily check the result for $D = 0$ since the linearized system becomes:

$$y[t+1] = (1 - \kappa(p + xp'))y[t];$$

so assume $D \geq 1$ for the rest of the proof. By Lemma 1, the maximum modulus of the roots of equation (9) varies continuously with κ . Also we know that the system is locally stable for small κ from Lemma 2. Hence, for the first part of the theorem, it suffices to look for the smallest κ such that equation (9) has a root of unit modulus. Let the root be $\lambda = e^{2i\theta}$; then equation (9) may be rewritten as:

$$2 \sin \theta e^{i((2D+1)\theta - \pi/2)} = a,$$

where $a = \kappa(p + xp')$. Hence we conclude that:

$$2|\sin \theta| = a, \quad \theta = \frac{\pi}{2(2D+1)} + \frac{2\pi n}{2D+1}$$

where n is an integer. We choose $n = 0$ since we are looking for the smallest positive a such that λ is of unit modulus. Hence, substituting for θ in $2|\sin \theta| = a$, we see that there are no solutions for θ if $a < 2 \sin\left(\frac{\pi}{2(2D+1)}\right)$. This gives the first part of the theorem.

For the second part of the theorem, it suffices to prove that there exists a root $\lambda = re^{i\theta}$ of modulus $r > 1$ to equation (9) when $a > 2 \sin(\theta^*/2)$, where $\theta^* = \pi/(2D+1)$. From (9), note that $\lambda = 1 - a\lambda^{-D}$, and that as vectors in the complex plane, the angle between λ and 1 is θ . Thus, applying the cosine rule, we have:

$$\cos \theta = \frac{1 + r^2 - a^2 r^{-2D}}{2r}. \quad (10)$$

Further, by equating imaginary parts in (9), we have:

$$r = \frac{\sin(D\theta)}{\sin((D+1)\theta)}. \quad (11)$$

We know that for $a = 2 \sin(\theta^*/2)$, $\cos \theta^* = 1 - \frac{a^2}{2}$ and $r = 1$. Suppose $a > 2 \sin(\theta^*/2)$. Then the RHS of (10) is less than $\cos \theta^*$ for $r = 1$. Notice that the LHS of (10) decreases from $\cos \theta^*$ to $\cos(\pi/(D+1))$ for $\theta \in [\theta^*, \frac{\pi}{D+1}]$. To prove that a root of modulus greater than unity exists, it suffices to show that r increases from 1 to ∞ for $\theta \in [\theta^*, \frac{\pi}{D+1}]$ and that the RHS of (10) is increasing in r for $r \geq 1$.

From equation (11),

$$\begin{aligned} \frac{dr}{d\theta} &= \frac{D \sin((D+1)\theta) \cos(D\theta) - (D+1) \cos((D+1)\theta) \sin(D\theta)}{\sin^2((D+1)\theta)} \\ &> 0, \end{aligned}$$

since $\sin(D\theta) > \sin((D+1)\theta)$ and $-\cos((D+1)\theta) + \cos(D\theta) > 0$ for $\theta \in [\theta^*, \frac{\pi}{D+1}]$. So r increases from 1 to ∞ for $\theta \in [\theta^*, \frac{\pi}{D+1}]$.

Differentiating the RHS of (10) with respect to r gives:

$$\frac{d}{dr}(10) = \frac{2(r^2 - 1) + (4D + 2)a^2 r^{-2D}}{4r^2},$$

which shows that the RHS of (10) is increasing in r for $r \geq 1$, tending to $r/2$ when r is large. Thus a root of modulus greater than unity exists for which $\theta \in [\theta^*, \frac{\pi}{D+1}]$ when $a > 2 \sin(\theta^*/2)$. This proves the second part of the theorem. \square

2 Rate of convergence: one route, one resource

For the simplified case where we have only one resource, we may also study rate of convergence to the stable point. In this section, we consider this problem, via the theory of differential-difference equations. For convenience, we restate here a result due to Hayes [7] taken from Bellman and Cooke ([1], Theorem 13.8).

Lemma 4 (Hayes) *All the roots of $be^\lambda + c - \lambda e^\lambda = 0$, where b and c are real, have negative real parts if and only if: (1) $b < 1$; and (2) $b < -c < \sqrt{a_1^2 + b^2}$, where a_1 is the root of $a = b \tan a$ such that $0 < a < \pi$. If $b = 0$, we take $a_1 = \pi/2$.*

We will study the rate of convergence of the system in Figure 1, through the following differential-difference equation:

$$\frac{d}{dt}x(t) = \kappa[w - x(t - D)p(x(t - D))], \quad (12)$$

where $D = d_1 + d_2$. This is just the continuous-time analog of the discrete-time equation (7). We study the stability of this simple system via its linearized version ([1], Theorem 11.2). Let the fixed point be (x, p) where $p = p(x)$ and $w = xp$. Then linearizing with $x[t] = x + y[t]$, we obtain:

$$\frac{d}{dt}y(t) = -\kappa(p + xp')y(t - D),$$

neglecting higher order terms. The corresponding characteristic equation, obtained by substituting $y = e^{st}$, is thus:

$$s = -\kappa(p + xp')e^{-sD},$$

which after substituting $\lambda = sD$, reduces to:

$$-\kappa(p + xp')D - \lambda e^\lambda = 0. \quad (13)$$

The fixed point is locally stable if all roots of the above equation have negative real part ([12], Section 9.4). For each D , we are interested in the maximum value of κ such that the system is locally stable.

Theorem 5 *The system (12) is locally stable if:*

$$\kappa(p + xp') < \frac{\pi}{2D},$$

and unstable if:

$$\kappa(p + xp') > \frac{\pi}{2D}.$$

Proof. A direct application of Lemma 4 to equation (13) with $b = 0$ and $c = -\kappa D(p + xp')$. \square

Notice that the result from the differential-difference equation approach is just the corresponding discrete-time Theorem 3 with $2 \sin(\pi/(4D + 2))$ replaced by $\pi/2D$. This is not surprising since the discrete-time system (7) is an approximation to the differential-difference equation (12). Suppose we rewrite equation (7) in the following form:

$$x[t + 1/2] = x[t - 1/2] + \kappa(w - x[(t - 1/2) - (D - 1/2)]\mu[(t - 1/2) - (d_2 - 1/2)]),$$

and let t and D take values in $\{n/2 : n \text{ odd}\}$. Then, setting $t' = t - 1/2$ and replacing D by $D - 1/2$, the bound for Theorem 3 becomes $2 \sin(\pi/4D)$, which tends to $\pi/2D$ as $D \rightarrow \infty$.

Define $a = \kappa(p + xp')$; then, from above, the linearization of (12) is:

$$\frac{d}{dt}y(t) = ay(t - D). \quad (14)$$

Suppose the system (14) is stable. Let λ^* be the root of equation (13) with the smallest (in modulus) negative real part. Then the rate of convergence to the stable point is equal to $|\operatorname{Re} \lambda^*|D^{-1}$. We will investigate how the rate of convergence depends on a in the region of stability $(0, \pi(2D)^{-1})$.

Lemma 6 *The rate of convergence of the system (14) is monotonic decreasing from D^{-1} to 0 for $a \in [(eD)^{-1}, \pi(2D)^{-1}]$.*

Proof. Let the root be $\lambda = -\gamma + \delta i$ where $\gamma > 0$. Then from equation (13):

$$\gamma e^{-\gamma} \cos \delta + \delta e^{-\gamma} \sin \delta = aD; \quad (15)$$

$$\gamma e^{-\gamma} \sin \delta - \delta e^{-\gamma} \cos \delta = 0. \quad (16)$$

Notice that if (γ, δ) satisfies the above equations, then so does $(\gamma, -\delta)$. Hence, without loss of generality, we may assume that $\delta \geq 0$. Equation (15) gives:

$$\gamma = \frac{\delta}{\tan \delta}, \quad (17)$$

which on substitution into (16) yields:

$$\frac{\delta}{\sin \delta} e^{-\delta/\tan \delta} = aD. \quad (18)$$

Observe that $\delta \in [2n\pi, (2n + \frac{1}{2})\pi]$ where $n \in \mathbb{Z}_+$ since $\tan \delta \geq 0$ and $\sin \delta \geq 0$ from equations (17)-(18). Also, from (17) γ is monotonic decreasing in δ from 1 to 0 for $\delta \in [0, \pi/2]$. Then the LHS of (18) is monotonic increasing in δ from e^{-1} to $\pi/2$ for $\delta \in [0, \pi/2]$. Hence, a root of equations (17)-(18) exists for $\delta \in [0, \pi/2]$ which is decreasing in γ as a increases from $(eD)^{-1}$ to $\pi(2D)^{-1}$. It remains to show that this is the root with the smallest γ . Notice that if $\delta/\sin \delta$ is larger, then γ has to be larger for equation (18) to be satisfied. But $\delta/\sin \delta \leq \pi/2$ for $n = 0$ and $\delta/\sin \delta > \pi/2$ for $n > 0$, so taking $n = 0$ yields the root with the smallest γ . Then, since the rate of convergence is γD^{-1} , we have the required result. \square

Lemma 7 *The rate of convergence of the system (14) is monotonic increasing and convex from 0 to D^{-1} , and the convergence is nonoscillatory, for $a \in (0, (eD)^{-1}]$.*

Proof. Consider the function $-\lambda e^\lambda$ when λ is real and negative. Suppose we maximize the function with respect to λ . We have:

$$\frac{d}{d\lambda}(-\lambda e^\lambda) = -e^\lambda - \lambda e^\lambda,$$

which equals zero when $\lambda = -1$, giving a maximum of e^{-1} . Moreover, $-\lambda e^\lambda$ is increasing in λ for the interval $(-\infty, -1]$ and concave decreasing for $[-1, 0]$. The function is shown in Figure 2, from which it is clear that equation (13) has 2 negative roots when $a < (eD)^{-1}$ and 1 negative root when $a = (eD)^{-1}$. Moreover, the modulus of the smaller root (in modulus) increases from 0 to 1 as a increases from 0 to $(eD)^{-1}$, and is clearly a convex function of a .

It remains to show that taking λ to be real yields the root with the smallest (in modulus) negative real part; this will also show that the convergence is nonoscillatory. First note that the modulus of the smaller real root is not greater than 1. Suppose there are complex roots for $a \leq (eD)^{-1}$. Then using the same argument as in the proof of Lemma 6, we see that there are no roots for $n = 0$, $\delta \neq 0$, since the LHS of equation (18) is greater than e^{-1} . Also, for $n > 0$, $\delta/\sin \delta > \pi/2$, so the complex root, if any, will have $\gamma > 1$. This completes the proof. \square

We combine the above lemmas into the following theorem.

Theorem 8 *The maximum rate of convergence for the system (14) is D^{-1} when $a = (eD)^{-1}$. The rate of convergence is monotonic increasing and convex from 0 to D^{-1} , and the convergence is nonoscillatory, for $a \in (0, (eD)^{-1}]$. The rate of convergence is monotonic decreasing from D^{-1} to 0 for $a \in [(eD)^{-1}, \pi(2D)^{-1})$.*

In particular, the maximum rate of convergence to the equilibrium point x of (12) is achieved if and only if:

$$\kappa(p + xp') = \frac{1}{eD},$$

and in this case, the equilibrium is nonoscillatory.

Figure 3 illustrates the result of Theorem 8, showing how the rate of convergence for the system (14) varies with $a \in (0, \pi(2D)^{-1})$, the region of stability. Using the analogy from ordinary differential equations, we see that the system is over-damped for $0 < a < (eD)^{-1}$ and under-damped for $(eD)^{-1} < a < \pi(2D)^{-1}$. The system is critically damped when $a = (eD)^{-1}$, and this also gives the optimal rate of convergence of D^{-1} or “time constant” of D . Intuitively, this makes sense since D is the roundtrip delay.

The result of Theorem 8 has also been shown by Chong et al. [2] using essentially Lemma 4. However, their claim that the rate of convergence is a concave function of gain differs from what is depicted in Figure 3; in fact, by Lemma 7, the rate of convergence is convex for $a \in (0, (eD)^{-1}]$.

3 Stability criteria: networks with constant roundtrip delay

We wish now to formulate results analagous to those of Section 1 for general networks, i.e., systems described by the equations (4)-(5). The theory for multidimensional delayed difference equations is significantly more complex than in the one-dimensional case, as we shall discover; nonetheless, it is possible to prove a result analogous to Theorem 3 for networks, in the special case where $D_r = D$ for all $r \in R$.

We begin, as before, by linearizing the system about the stable point $x = (x_r, r \in R)$, given by (3). For all $j \in J$, define $p_j = p_j(\sum_{s \in S} x_s)$; assume that at this point p_j is differentiable, and let $p'_j = p'_j(\sum_{s \in S} x_s)$. Define $y_r[t]$ by $x_r[t] = x_r + \kappa_r^{1/2} x_r^{1/2} y_r[t]$; then since $\sum_{j \in J} A_{jr} p_j = w_r x_r^{-1}$, the linearization of (4)-(5) yields:

$$\begin{aligned} y_r[t+1] &= y_r[t] - \kappa_r w_r x_r^{-1} y_r[t - D_r] \\ &\quad - \sum_{j \in J} \sum_{s \in R} A_{jr} \kappa_r^{1/2} x_r^{1/2} A_{js} \kappa_s^{1/2} x_s^{1/2} p'_j y_s[t - d_1(j,s) - d_2(j,r)], \quad r \in R, \end{aligned} \quad (19)$$

neglecting higher order terms. Fix $\lambda \in \mathbb{C}$. We now ask the question: does there exist a (possibly complex) vector $\alpha = (\alpha_r, r \in R)$, such that $y[t] = \alpha \lambda^t$ is a solution to this system of equations? If such an α exists, the solution is a normal mode. As before, local stability results if all normal modes satisfy $|\lambda| < 1$.

Our approach will be to find the conditions under which α exists, such that $\alpha \lambda^t$ is a normal mode. Suppose such an α exists; then, substituting into (19):

$$\begin{aligned} \alpha_r \lambda^{t+1} &= \alpha_r \lambda^t - \kappa_r w_r x_r^{-1} \alpha_r \lambda^{t-D_r} \\ &\quad - \sum_{j \in J} \sum_{s \in R} A_{jr} \kappa_r^{1/2} x_r^{1/2} A_{js} \kappa_s^{1/2} x_s^{1/2} p'_j \alpha_s \lambda^{t-d_1(j,s)-d_2(j,r)}, \quad r \in R. \end{aligned}$$

After cancelling λ^t throughout, multiplying by λ^{D_r} , and applying the definition of D_r , we have:

$$(\lambda^{D_r+1} - \lambda^{D_r})\alpha_r + \kappa_r w_r x_r^{-1} \alpha_r - \sum_{j \in J} \sum_{s \in R} A_{jr} \kappa_r^{1/2} x_r^{1/2} \lambda^{d_1(j,r)} A_{js} \kappa_s^{1/2} x_s^{1/2} \lambda^{-d_1(j,s)} p'_j \alpha_s = 0, \quad r \in R. \quad (20)$$

Define the following matrices: $\kappa = \text{diag}(\kappa_r, r \in R)$; $W = \text{diag}(w_r, r \in R)$; $X = \text{diag}(x_r, r \in R)$; $P' = \text{diag}(p'_j, j \in J)$; and $A(\lambda) = (A_{jr} \lambda^{-d_1(j,r)}, j \in J, r \in R)$. Then we can express (20) as the following matrix equation:

$$\left(\text{diag}(\lambda^{D_r+1} - \lambda^{D_r}, r \in R) + \kappa W X^{-1} + \kappa^{1/2} X^{1/2} A(\lambda^{-1})^T P' A(\lambda) X^{1/2} \kappa^{1/2} \right) \alpha = 0.$$

But this is possible if and only if the determinant of the matrix premultiplying α is zero, i.e., if and only if the following equation holds:

$$\det \left(\text{diag}(\lambda^{D_r+1} - \lambda^{D_r}, r \in R) + \kappa W X^{-1} + \kappa^{1/2} X^{1/2} A(\lambda^{-1})^T P' A(\lambda) X^{1/2} \kappa^{1/2} \right) = 0. \quad (21)$$

The equation (21) is the *characteristic equation* for the system (19). A normal mode exists only for those λ which are roots of the characteristic equation. If all roots of (21) have modulus less than unity, then the system (19) will be asymptotically stable, and hence the stable point x will be locally stable for the system (4)-(5).

We now investigate conditions to control the maximum modulus of the roots of the characteristic equation. We will write $p(\lambda, \kappa)$ for the left hand side of (21), to emphasize the changes in the characteristic equation as κ varies. We also make the following definition:

$$C(\lambda, \kappa) = \kappa W X^{-1} + \kappa^{1/2} X^{1/2} A(\lambda^{-1})^T P' A(\lambda) X^{1/2} \kappa^{1/2}.$$

We have the following theorem, which gives stability criteria in the situation where $D_r = D$ for all $r \in R$.

Theorem 9 *Suppose $D_r = D$ for all $r \in R$. The system (4)-(5) is locally stable if the following condition is satisfied for all $r \in R$:*

$$\kappa_r \left(\sum_{j \in r} p_j + \sum_{j \in r} p'_j \sum_{s: j \in s} x_s \right) < 2 \sin \left(\frac{\pi}{2(2D+1)} \right). \quad (22)$$

Proof. We will divide the proof into five steps.

Step 1. If $0 < a < 2 \sin \frac{\pi}{2(2D+1)}$, then no roots of $\lambda^{D+1} - \lambda^D + a = 0$ have modulus equal to unity. This deduction follows from the proof of Theorem 3.

Step 2. The maximum modulus of the roots λ of $p(\lambda, \kappa) = 0$ is continuous in κ . This follows from Lemma 1. In fact, a small point of subtlety arises: $p(\lambda, \kappa) = 0$ is not immediately a polynomial equation, as it involves terms which are powers of λ^{-1} . However, this is easily rectified by multiplying through with a suitably large power of λ .

Step 3. Show that for any κ satisfying the hypotheses of the theorem, $p(\lambda, \kappa) = 0$ has no roots of modulus equal to unity. Let κ satisfy the hypotheses of the theorem, and suppose there exists $\lambda = e^{i\theta}$, $0 \leq \theta \leq 2\pi$, such that $p(\lambda, \kappa) = 0$. Since $w_r x_r^{-1} = \sum_{j \in R} A_{jr} p_j$, $|\lambda| = 1$, and all κ_r , w_r , x_r , and p_j are nonnegative, the hypotheses of the theorem yield:

$$\begin{aligned} & \left| \kappa_r w_r x_r^{-1} + \sum_{j \in J} A_{jr} \kappa_r x_r p'_j \right| + \sum_{s \neq r} \left| \kappa_r \sum_{j \in J} A_{jr} \lambda^{d_1(j,r)} A_{js} \lambda^{-d_1(j,s)} x_s p'_j \right| \\ & \leq \kappa_r \left(\sum_{j \in R} A_{jr} p_j + \sum_{s \in R} \sum_{j \in J} |A_{jr} A_{js} \lambda^{d_1(j,r) - d_1(j,s)} x_s p'_j| \right) \\ & = \kappa_r \left(\sum_{j \in r} p_j + \sum_{j \in r} \sum_{s: j \in s} x_s p'_j \right) \\ & < 2 \sin \left(\frac{\pi}{2(2D+1)} \right), \quad r \in R. \end{aligned}$$

Although the first line of this computation seems awkward, it is in fact the “absolute row sum” of row r of the matrix $\kappa(WX^{-1} + A(\lambda^{-1})^T P' A(\lambda) X)$. Since the spectral radius of any square matrix is bounded by its maximum absolute row sum ([8], Chapter 8), we have the following bound (where $\|\cdot\|$ denotes spectral radius):

$$\begin{aligned} \|C(\lambda, \kappa)\| &= \left\| \kappa W X^{-1} + \kappa^{1/2} X^{1/2} A(\lambda^{-1})^T P' A(\lambda) X^{1/2} \kappa^{1/2} \right\| \\ &= \left\| \kappa W X^{-1} + \kappa A(\lambda^{-1})^T P' A(\lambda) X \right\| \\ &< 2 \sin \left(\frac{\pi}{2(2D+1)} \right). \end{aligned}$$

Thus, we can bound the spectral radius of $C(\lambda, \kappa)$ in (21). We now make the following observation: if $D_r = D$, then the characteristic equation can be written more simply as:

$$\det((\lambda^{D+1} - \lambda^D)I + C(\lambda, \kappa)) = 0. \quad (23)$$

Further, as $\lambda = e^{i\theta}$, the matrix $C(\lambda, \kappa)$ is Hermitian. It is also positive definite, so we can write $\Gamma \Phi \Gamma^T = C(\lambda, \kappa)$, where Γ is unitary and $\Phi = \text{diag}(\phi_r, r \in R)$ is the diagonal matrix of strictly positive eigenvalues of $C(\lambda, \kappa)$. (Note that both Γ and Φ are dependent on λ and κ , though this has been suppressed in the notation.) By the bound computed above, we know that $\phi_r < 2 \sin \frac{\pi}{2(2D+1)}$ for all $r \in R$. Since Γ is unitary, we can factor it out of the characteristic equation:

$$\det(\text{diag}(\lambda^{D+1} - \lambda^D + \phi_r, r \in R)) = 0.$$

But we are now taking the determinant of a diagonal matrix; and so, if λ is a root of unit modulus such that $p(\lambda, \kappa) = 0$, we must have:

$$\lambda^{D+1} - \lambda^D + \phi_r = 0$$

for ϕ_r satisfying $0 < \phi_r < 2 \sin \frac{\pi}{2(2D+1)}$. This contradicts Step 1, and so Step 3 is proven.

Step 4. Show there exists a κ satisfying the hypotheses of the theorem, such that all roots λ of $p(\lambda, \kappa) = 0$ have modulus less than unity. For convenience, we assume that $R = \{1, 2, \dots, N\}$. Define $R_n = \{1, 2, \dots, n\}$. We denote by $p_n(\lambda, \kappa) = 0$ the characteristic equation defined by the subnetwork of routes in R_n ; mathematically, this corresponds to replacing $C(\lambda, \kappa)$ with the submatrix $C(\lambda, \kappa) = ([C(\lambda, \kappa)]_{rs}, r, s \in R_n)$ in (23). (Although the number of components in κ may change depending on n , we will suppress this dependence and write κ for the vector of gain parameters, with the number of components understood from context.) The result is proven inductively on n .

First suppose that $n = 1$. We have shown in Theorem 3 that there exists a κ_1 satisfying the hypotheses of the theorem for which all roots of $p_1(\lambda, \kappa) = 0$ have modulus less than unity.

Now, inductively, assume there exist parameters $\kappa_1, \dots, \kappa_{n-1}$ such that all roots λ of $p_{n-1}(\lambda, \kappa) = 0$ have modulus less than unity. We consider roots of the equation $p_n(\lambda, \kappa) = 0$. Let $k = (\kappa_1, \dots, \kappa_{n-1}, 0)$. Then note the following relation:

$$p_n(\lambda, k) = (\lambda^{D+1} - \lambda^D)p_{n-1}(\lambda, \kappa).$$

This follows by decomposing the determinant which defines p_n . Thus, there are a total of nD roots: $nD - 1$ of these roots have modulus less than unity (including $\lambda = 0$ of multiplicity D), and $\lambda = 1$ has multiplicity 1. If we can ensure that a small increase in κ_n from zero reduces the magnitude of the root $\lambda = 1$, then an argument exactly analogous to the case $n = 1$ would show that κ_n can be chosen such that $p_n(\lambda, \kappa) = 0$ has all roots of modulus less than unity for gain parameters $\kappa_1, \dots, \kappa_n$.

We proceed by using the Implicit Function Theorem, as before. By expanding the determinant which defines p_n , we can write:

$$p_n(\lambda, \kappa) = (\lambda^{D+1} - \lambda^D)p_{n-1}(\lambda, \kappa) + \kappa_n q_{n-1}(\lambda, \kappa),$$

where q_{n-1} is defined by:

$$q_{n-1}(\lambda, \kappa) = \det \left(\begin{array}{c} (\lambda^{D+1} - \lambda^D) \begin{pmatrix} 1 & \dots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \dots & 1 & 0 \\ 0 & \dots & 0 & 0 \end{pmatrix} + C_n(\lambda, \kappa_1, \dots, \kappa_{n-1}, 1) \end{array} \right).$$

First we differentiate p_n with respect to λ , evaluating at $(1, k)$:

$$\left. \frac{\partial}{\partial \lambda} p_n(\lambda, \kappa) \right|_{(1, k)} = p_{n-1}(1, \kappa).$$

Now, since we know that p_{n-1} has all roots of modulus less than unity at $(\kappa_1, \dots, \kappa_{n-1})$, we can conclude that $p_{n-1}(1, \kappa) \neq 0$. So, applying the Implicit Function Theorem, there exists a differentiable complex-valued function $h(\kappa_n)$ and $\varepsilon_n > 0$ such that $h(0) = 1$, and $p_n(h(\kappa_n), \kappa_n) = 0$ if $-\varepsilon_n < \kappa_n < \varepsilon_n$. As before, we proceed to compute the derivative of $h(\kappa_n)$ using the implicit definition in p_n . Since $h(0) = 1$, and:

$$(h(\kappa_n)^{D+1} - h(\kappa_n)^D)p_{n-1}(h(\kappa_n), \kappa) + \kappa_n q_{n-1}(h(\kappa_n), \kappa) = 0,$$

we conclude:

$$h'(0)p_{n-1}(1, \kappa) + q_{n-1}(1, \kappa) = 0.$$

But $p_{n-1}(1, \kappa)$ and $q_{n-1}(1, \kappa)$ are the following determinants:

$$\begin{aligned} p_{n-1}(1, \kappa) &= \det(C_{n-1}(1, \kappa)), \text{ and} \\ q_{n-1}(1, \kappa) &= \det(C_n(1, \kappa_1, \dots, \kappa_{n-1}, 1)). \end{aligned}$$

In both cases, these determinants are strictly positive, as they are the determinants of positive definite matrices. Thus, we have $h'(0) = -q_{n-1}(1, \kappa)/p_{n-1}(1, \kappa) = -a_n$, where $a_n > 0$. This leads to the following Taylor expansion:

$$h'(\kappa_n) = 1 - \kappa_n a_n + o(\kappa_n).$$

So as κ_n increases away from zero, the root $\lambda = 1$ decreases in magnitude. By the same argument as in the proof of Theorem 3, this is sufficient to complete the inductive step: namely, we can choose a κ_n satisfying the hypotheses of the theorem, and such that all roots λ of $p_n(\lambda, \kappa) = 0$ have modulus less than unity. Q.E.D.

Now, if we take $n = N$, then $p_n = p$, and so we can find a vector $\kappa = (\kappa_1, \dots, \kappa_N)$ satisfying the hypotheses of the theorem such that all roots λ of $p(\lambda, \kappa) = 0$ have modulus less than unity; we will denote this κ by κ^* . This completes Step 4.

Step 5. Completion of proof. Now suppose that, for some κ satisfying the hypotheses of the theorem, $p(\lambda, \kappa) = 0$ has a root of modulus greater than unity. Consider the path $\kappa(t) = t\kappa^* + (1-t)\kappa$, for $0 \leq t \leq 1$. All roots of $p(\lambda, \kappa(1)) = 0$ have modulus less than unity by Step 4; so, by Step 2 (continuity of maximum modulus of roots), there exists t such that $p(\lambda, \kappa(t)) = 0$ has a root λ of modulus unity, i.e., $|\lambda| = 1$. But, since $\kappa(t)$ satisfies the hypotheses of the theorem (both κ^* and κ satisfy the hypotheses of the theorem, as does any convex combination), this is a contradiction to Step 3. So we conclude that no such κ exists; i.e., for all κ satisfying the hypotheses of the theorem, $p(\lambda, \kappa) = 0$ has all roots of modulus less than unity. \square

The proof of Theorem 9 fails at Step 3 if all D_r are not equal; at this point, the determinant defining (21) may not decompose into a product form. The generality in the result given here is that the delays $d_1(j, r)$ and $d_2(j, r)$ may be chosen arbitrarily, subject to the constraint that $d_1(j, r) + d_2(j, r) = D$ for all $j \in J, r \in R$. Notice that the decomposition of Step 3 works because $C(\lambda, \kappa)$ is a Hermitian matrix when $|\lambda| = 1$; this property is a direct result of the fact that $d_1(j, r) + d_2(j, r) = D_r$ for all $j \in J, r \in R$, emphasizing the importance of the roundtrip delay in this analysis.

3.1 Nonadaptive users

The simplest extension to the model (4)-(5) is to add users who act at intervals of time significantly longer than the roundtrip delays D_r ; in fact, so slowly that their rates appear constant on the timescale of the D_r . Intuitively, this is the description of a user with an ‘‘infinite roundtrip delay’’: they set their rate, then update it on a time horizon far longer than the update times of most users in the system.

We may formalize this model as follows. Suppose that $R = R_1 \cup R_2$, and consider the following

model:

$$x_r[t+1] = x_r[t] + \kappa_r \left(w_r - x_r[t - D_r] \sum_{j \in r} \mu_j[t - d_2(j, r)] \right), \quad r \in R_1; \quad (24)$$

$$x_u[t] = x_u, \quad u \in R_2; \quad (25)$$

where $x_u > 0$ for all $u \in R_2$, and:

$$\mu_j[t] = p_j \left(\sum_{s \in R_1: j \in s} x_s[t - d_1(j, s)] + \sum_{v \in R_2: j \in v} x_v \right), \quad j \in J. \quad (26)$$

The routes in R_2 are *nonadaptive*: they do not respond to congestion indication signals from the resource. We will show that if $D_r = D$ for all $r \in R_1$, Theorem 9 provides a sufficient stability condition for the system (24)-(26).

As before, let $x = (x_r, r \in R)$ be the fixed point, with $p_j = p_j(\sum_{s \in R: j \in s} x_s)$ and $p'_j = (\sum_{s \in R: j \in s} x_s)$. Note that these definitions include the routes in both R_1 and R_2 . We have the following theorem.

Theorem 10 *Suppose $D_r = D$ for all $r \in R_1$. The system (24)-(26) is locally stable if the following condition is satisfied for all $r \in R_1$:*

$$\kappa_r \left(\sum_{j \in r} p_j + \sum_{j \in r} p'_j \sum_{s \in R: j \in s} x_s \right) < 2 \sin \left(\frac{\pi}{2(2D+1)} \right). \quad (27)$$

Proof. The linearization of the system (24)-(26) is given by (19) for all $r \in R_2$, with p_j and p'_j defined as above. Thus, the exact same proof as used for Theorem 9 goes through to give the result. \square

3.2 Instantly adaptive users

We now consider the situation opposite to Section 3.1: users updating at time intervals significantly shorter than the roundtrip delays D_r . We may model such users by assuming they reach equilibrium quickly—indeed, instantaneously. Intuitively, this is the description of a user with “zero roundtrip delay”: equilibrium may be reached instantaneously because there is negligible lag in communication between the network and the user.

We formalize the model as follows. Again, let $R = R_1 \cup R_2$, and model the set of users in R_2 by:

$$x_u[t] = \frac{w_u}{\sum_{j \in u} \mu_j[t]}, \quad u \in R_2. \quad (28)$$

The set of users in R_1 is modeled as before:

$$x_r[t+1] = x_r[t] + \kappa_r \left(w_r - x_r[t - D_r] \sum_{j \in r} \mu_j[t - d_2(j, r)] \right), \quad r \in R_1. \quad (29)$$

Finally, $\mu_j[t]$ is defined as before, but with the modification that routes in $x_u[t]$ suffer no delay:

$$\mu_j[t] = p_j \left(\sum_{v \in R_2: j \in v} x_v[t] + \sum_{s \in R_1: j \in s} x_s[t - d_1(j, s)] \right), \quad j \in J. \quad (30)$$

The existence and uniqueness of the equilibrium (3), demonstrated in [11], shows the system (28)-(30) is well defined. The routes in R_2 are *instantly adaptive*: they instantaneously achieve equilibrium given the current state of the network. Thus, they are users who update rates much more quickly than the users in R_1 .

As before, the system has a unique equilibrium point $x = (x_r, r \in R)$. Then, defining $y_r[t] = \kappa_r^{-1/2} x_r^{-1/2} (x_r[t] - x_r)$ for $r \in R_1$, and $y_u[t] = x_u[t] - x_u$ for $u \in R_2$, we can linearize to obtain:

$$\begin{aligned} y_r[t+1] &= y_r[t] - \kappa_r w_r x_r^{-1} y_r[t - D_r] \\ &\quad - \sum_{v \in R_2} \sum_{j \in J} A_{jr} \kappa_r^{1/2} x_r^{1/2} A_{jv} p'_j z_v[t - d_2(j, r)] \\ &\quad - \sum_{s \in R_1} \sum_{j \in J} A_{jr} \kappa_r^{1/2} x_r^{1/2} A_{js} \kappa_s^{1/2} x_s^{1/2} p'_j y_s[t - d_1(j, s) - d_2(j, r)], \quad r \in R_1; \\ y_u[t] &= -x_u^2 w_u^{-1} \left(\sum_{v \in R_2} \sum_{j \in J} A_{jv} A_{ju} p'_j y_v[t] \right. \\ &\quad \left. + \sum_{s \in R_1} \sum_{j \in J} A_{ju} A_{js} \kappa_s^{1/2} x_s^{1/2} p'_j y_s[t - d_1(j, s)] \right), \quad u \in R_2. \end{aligned}$$

Let $A_{R_1}(\lambda) = (A_{jr} \lambda^{-d_1(j, r)}, j \in J, r \in R_1)$, and let $A_{R_2} = (A_{ju}, j \in J, u \in R_2)$. Let $X_{R_1} = \text{diag}(x_r, r \in R_1)$; similarly, define W_{R_1} , κ_{R_1} , X_{R_2} , and W_{R_2} . We try a solution of the form $\alpha \lambda^t$. Letting $\alpha_{R_1} = (\alpha_r, r \in R_1)$, $\alpha_{R_2} = (\alpha_u, u \in R_2)$, we have:

$$(A_{R_2}^T P' A_{R_2} + X_{R_2}^{-2} W_{R_2}) \alpha_{R_2} = -A_{R_2}^T P' A_{R_1}(\lambda) X_{R_1}^{1/2} \kappa_{R_1}^{1/2} \alpha_{R_1},$$

and

$$\begin{aligned} & \left(\text{diag}(\lambda^{D_r+1} - \lambda^{D_r}, r \in R_2) + \kappa_{R_1} W_{R_1} X_{R_1}^{-1} \right. \\ & \quad \left. + \kappa_{R_1}^{1/2} X_{R_1}^{1/2} A_{R_1}(\lambda^{-1})^T P' A_{R_1}(\lambda) X_{R_1}^{1/2} \kappa_{R_1}^{1/2} \right) \alpha_{R_1} + \kappa_{R_1}^{1/2} X_{R_1}^{1/2} A_{R_1}(\lambda^{-1})^T P' A_{R_2} \alpha_{R_2} = 0. \end{aligned}$$

After substituting the first equation into the second, we obtain the characteristic equation for (28)-(30):

$$\begin{aligned} & \det(\text{diag}(\lambda^{D_r+1} - \lambda^{D_r}, r \in R_1) + \kappa_{R_1} W_{R_1} X_{R_1}^{-1} + \kappa_{R_1}^{1/2} X_{R_1}^{1/2} A_{R_1}(\lambda^{-1})^T P' A_{R_1}(\lambda) X_{R_1}^{1/2} \kappa_{R_1}^{1/2} \\ & \quad - \kappa_{R_1}^{1/2} X_{R_1}^{1/2} A_{R_1}(\lambda^{-1})^T P' A_{R_2} (A_{R_2}^T P' A_{R_2} + X_{R_2}^{-2} W_{R_2})^{-1} A_{R_2}^T P' A_{R_1}(\lambda) X_{R_1}^{1/2} \kappa_{R_1}^{1/2}) = 0. \quad (31) \end{aligned}$$

We make the following definitions:

$$\begin{aligned} M &= (P')^{1/2} A_{R_2} X_{R_2} W_{R_2}^{-1/2}, \quad \text{and} \\ D(\lambda, \kappa_{R_1}) &= \kappa_{R_1} W_{R_1} X_{R_1}^{-1} + \\ & \quad \kappa_{R_1}^{1/2} X_{R_1}^{1/2} A_{R_1}(\lambda^{-1})^T (P')^{1/2} (I - M(I + M^T M)^{-1} M^T) (P')^{1/2} A_{R_1}(\lambda) X_{R_1}^{1/2} \kappa_{R_1}^{1/2}. \end{aligned}$$

Then the characteristic equation reduces to:

$$\det(\text{diag}(\lambda^{D_r+1} - \lambda^{D_r}, r \in R_1) + D(\lambda, \kappa_{R_1})) = 0.$$

Notice the similar structure of $D(\lambda, \kappa_{R_1})$ and $C(\lambda, \kappa)$; this allows us to prove the following theorem.

Theorem 11 *Suppose $D_r = D$ for all $r \in R_1$. The system (28)-(30) is locally stable if the condition (22) is satisfied for all $r \in R_1$.*

Proof. $D(\lambda, \kappa_{R_1})$ is identical in structure to $C(\lambda, \kappa)$, but with P' replaced by $(P')^{1/2}(I - M(I + M^T M)^{-1}M^T)(P')^{1/2}$ (note the latter matrix is independent of λ). Hence, we may use exactly the same proof as used for Theorem 9; the only difficulty arises in Step 3, where we showed that for κ satisfying the hypotheses of the theorem and $\lambda = e^{i\theta}$, $\|C(\lambda, \kappa)\| < 2 \sin(\frac{\pi}{2(2D+1)})$. We therefore proceed to show that for κ_{R_1} satisfying the hypotheses of the theorem, and λ such that $|\lambda| = 1$, $\|D(\lambda, \kappa_{R_1})\| < 2 \sin(\frac{\pi}{2(2D+1)})$.

We first show $\|I - M(I + M^T M)^{-1}M^T\| \leq 1$. Note M is $|J| \times |R|$; so write $M = U\Sigma V^T$, where U is the $|J| \times |J|$ orthogonal matrix of eigenvectors of MM^T , V is the $|R| \times |R|$ orthogonal matrix of eigenvectors of $M^T M$, and Σ is the $|J| \times |R|$ matrix of singular values of M : $\Sigma_{ij} = 0$ for $i \neq j$, and $\Sigma_{11} \geq \dots \geq \Sigma_{qq}$, where $q = \min(|J|, |R|)$. This is the singular value decomposition (SVD) of M [8]. Substituting, and using $V^T V = U^T U = VV^T = UU^T = I$, we have:

$$\begin{aligned} \|I - M(I + M^T M)^{-1}M^T\| &= \|I - (U\Sigma V^T)(I + V\Sigma^T \Sigma V^T)^{-1}V\Sigma^T U^T\| \\ &= \|U(I - \Sigma V^T V(I + \Sigma^T \Sigma)^{-1}V^T V \Sigma)U^T\| \\ &= \|I - \Sigma(I + \Sigma^T \Sigma)^{-1}\Sigma^T\|. \end{aligned}$$

The last expression is the spectral radius of a diagonal matrix, where each entry is of the form $1 - \Sigma_{ii}^2 / (1 + \Sigma_{ii}^2) \leq 1$; hence $\|I - M(I + M^T M)^{-1}M^T\| \leq 1$.

The argument above also shows that $I - M(I + M^T M)^{-1}M^T$ is positive semidefinite; further, it is symmetric. If $|\lambda| = 1$, then from the definition of $D(\lambda, \kappa_{R_1})$, we may write $D(\lambda, \kappa_{R_1}) = F + G^*HG$, where F is diagonal and positive definite, H is real diagonal and positive definite with $\|H\| \leq 1$, and $C(\lambda, \kappa_{R_1}) = F + G^*G$. Using the bound for $\|C(\lambda, \kappa_{R_1})\|$ from our proof of Theorem 9:

$$\begin{aligned} \|D(\lambda, \kappa_{R_1})\| &= \|F + G^*HG\| \\ &= \max_{v:|v|=1} \langle (F + G^*HG)v, v \rangle \\ &= \max_{v:|v|=1} \langle Fv, v \rangle + \langle HGv, Gv \rangle \\ &\leq \max_{v:|v|=1} \langle Fv, v \rangle + \langle Gv, Gv \rangle \\ &= \max_{v:|v|=1} \langle Fv, v \rangle + \langle G^*Gv, v \rangle \\ &= \|C(\lambda, \kappa_{R_1})\| \\ &< 2 \sin\left(\frac{\pi}{2(2D+1)}\right). \end{aligned}$$

This completes the proof. □

The results of this subsection and the previous subsection add two extremes of users to the basic result of Theorem 9: those who act much more slowly, and much more quickly, than users with roundtrip delays D_r . The result of Theorem 11 is particularly revealing. By showing that $\|D(\lambda, \kappa_{R_1})\| \leq \|C(\lambda, \kappa_{R_1})\|$, notice that the system with instantly adaptive users is, loosely, “more stable” than the basic system (4)-(5). This is to be expected, as instantly adaptive users damp oscillations in the system; in other words, the network is more stable if there are users who back off quickly when congestion increases.

4 Stability criteria: networks with diverse roundtrip delays

As was noted above, the condition that $D_r = D$ for all $r \in R$ was essential to the proof of Theorem 9. We now wish to investigate extensions to Theorem 9, which remove the condition of constant roundtrip delay. Given the form of the bounds in Theorems 3 and 9, we may make the following conjecture:

Conjecture 12 *The system (4)-(5) is locally stable if the following condition is satisfied for all $r \in R$:*

$$\kappa_r \left(\sum_{j \in r} p_j + \sum_{j \in r} p'_j \sum_{s: j \in s} x_s \right) < 2 \sin \left(\frac{\pi}{2(2D_r + 1)} \right). \quad (32)$$

This is just Theorem 9, without the stipulation that $D_r = D$ for all $r \in R$.

An immediate question arises: does there exist a counterexample to Conjecture 12? To attempt to answer this question, we generated 10,000 “random” networks, as follows. A random number of routes, from 1 to 5, and a random number of resources, from 1 to 5, were chosen. A random A matrix was then generated. For each j and r , D_r and $d_1(j, r)$ were randomly generated, with D_r between 1 and 15; we then set $d_2(j, r) = d_1(j, r) - D_r$. Finally, we chose vectors x , p , and p' randomly, and fixed w according to equation (3). This allowed us to compute the critical value of κ_r for each $r \in R$ as predicted by equation (32). We randomly chose κ_r to lie in the open interval between 0 and this critical value.

For each of these 10,000 networks, we computed the maximum modulus of the roots of the characteristic equation (21). *In every instance, the maximum modulus was strictly less than 1, so the network was locally stable.* This experiment certainly suggests that Conjecture 12 might be true, and an interesting future research problem concerns finding a formal proof of Conjecture 12.

5 Simulations

If we are to implement the stability criteria given here in actual networks, then an obvious question is how well our model of a network approximates that of a real system. Our model uses a *fluid flow approximation* instead of discrete data packets. Also, queues at the resources are not modeled at all. In this chapter, we will address these concerns through simulations, and compare with our earlier theoretical results.

Our simulations involve one resource, modeled as a discrete-time queue, with several users. We will investigate stability of this stochastic system under various roundtrip delays and gain

parameters, and determine what it means for queuing delays to become negligible. We will then look more closely at instability in the system, revealing the phenomenon of *phase-locking*. Finally, we will investigate stability when users with different roundtrip delays interact with each other.

5.1 A discrete-time queue with threshold marking

We will consider a stochastic model analogous to (4)-(5), with one resource and several users. We wish to model queuing effects, but on a discrete-time scale. The dynamics of such a queuing resource are modeled as:

$$Q_t = (Q_{t-1} + Y_t - 1)^+, \quad (33)$$

where Q_t is the queue length at time t , and Y_t is the number of packets which arrive at the queue at time t . The interpretation of the model is as follows. Packets arrive at the queue as a (discrete-time) stochastic process Y_1, Y_2, \dots . At time t , the resource first adds Y_t packets (if any) to the queue. If the queue is now nonempty, exactly one packet is served. Notice the similarity of these equations with the analysis of a (continuous-time) M/D/1 queue, examined at time points where individuals leave the queue; in that case, the queue length equation is $Q_t = (Q_{t-1} - 1)^+ + Y_t$ [6].

We will suppose that, for each t , Y_t is a Poisson random variable of parameter $y < 1$, independent of the past history of the process. This makes Q_t a discrete time Markov chain. We will try to compute an invariant distribution $\pi = (\pi_n, n \geq 0)$ for this Markov chain. Letting $p_n = \mathbb{P}(Y_t = n) = e^{-y} y^n / n!$, we know that π must satisfy:

$$\begin{aligned} \pi_0 &= (p_0 + p_1)\pi_0 + p_0\pi_1; \\ \pi_1 &= p_2\pi_0 + p_1\pi_1 + p_0\pi_2; \\ &\vdots \end{aligned}$$

We may use these equations to calculate the generating function of π , $H(s)$:

$$H(s) = \sum_{n \geq 0} \pi_n s^n = \frac{\pi_0 e^{-y} (1-s)}{e^{-y(1-s)} - s}.$$

Since we know that $\lim_{s \uparrow 1} H(s) = 1$, we conclude that $\pi_0 = e^y (1-y) > 0$. This last result also implies that $\pi_n > 0$ for all $n \geq 0$. Thus, Q possesses an invariant distribution π , given by the coefficients of the power series of $H(s)$.

We may interpret the parameter y as the mean number of packets per unit time which arrive at the queue. A simple choice for a congestion indication mechanism is a *threshold marking* scheme: we mark all arriving packets whenever the queue is at or above a threshold B . The user interprets the number of marks received as an indication of the level of congestion at the resource, and slows the sending rate accordingly.

The corresponding congestion indication function $p_B(y)$ in our fluid flow model is just the probability that an arriving packet receives a mark. Suppose the queue is in equilibrium; then we wish to calculate the probability the queue is at or above B , given that a packet arrives:

$$p_B(y) = \mathbb{P}(Q_t \geq B | Y_t > 0) = \mathbb{P}(Q_t \geq B) = 1 - \sum_{n=0}^{B-1} \pi_n,$$

by independence of Q_t and Y_t . For a fixed value of B , we can compute a closed form expression for $p_B(y)$ by finding the first B terms of the power series for $H(s)$. Note that, in the simulations which follow, the queue will not necessarily be at equilibrium; thus, $p_B(y)$ is an approximation to the true marking probability.

5.2 The simulation model

We are now ready to describe the model used for simulations. We will consider a set of users R sharing one resource, modeled as a discrete-time queue with threshold marking; let $m_r[t]$ be the number of marks received by user r at time t . We will suppose that each user maintains a sending “rate” $x_r[t]$, updated as follows:

$$x_r[t+1] = x_r[t] + \kappa_r (w_r - m_r[t]). \quad (34)$$

At time t , user r generates packets according to a Poisson distribution with mean $x_r[t]$. These packets suffer a propagation delay of $d_1(r)$ before arriving at the resource.

At the resource, packets arriving at time t are added to the queue in random order, ensuring that no user receives priority at the queue. The resource then serves exactly one packet, if the queue is nonempty. Each packet returns to its original sender; a packet returning to user r suffers a propagation delay of $d_2(r)$ in transmission from the resource to the user. Thus, the marked packets received by user r at time t , $m_r[t]$, left the resource $d_2(r)$ units of time earlier. We still define $D_r = d_1(r) + d_2(r)$; however, notice that now the observed roundtrip delay consists of the propagation delay D_r plus a random queueing delay.

The last paragraph makes clear the analogy between our simulation and equations (4)-(5): in the fluid flow approximation, $m_r[t]$ is given by $x_r[t - D_r]\mu[t - d_2(r)]$, which we had described as the number of marks received by user r at time t . We are now explicitly modeling, at the packet level, the congestion indication behavior described by (4)-(5)—but we are also adding to this model queueing at the resource.

In Theorems 3 and 9, the stability criteria are expressed in terms of the stable point (x, p, p') . For our simulations, we can fix the stable point in advance. We begin by choosing x_r for each user $r \in R$; then, given the threshold B , we can explicitly calculate $p = p_B(\sum x_r)$ from the theory in the last section. Similarly, we can calculate $p' = p'_B(\sum x_r)$. (Note this is only sensible as long as $\sum x_r < 1$, so the arrival rate does not exceed service capacity.) We then expect the rates to converge to x_r at equilibrium, as long as $w_r = x_r p$ for each user. Further, we can use the triple (x, p, p') to determine the critical gain κ_r for each user, from (32).

5.3 Simulation 1: Variance of rates

Recall that Theorems 3 and 9 present sufficient conditions for *local stability* by means of the linearized system (19). In our simulations, however, we have implemented a stochastic model of the nonlinear system (4)-(5). When the gain becomes large, we know the linearized system may not be stable—and hence the nonlinear system is not guaranteed to be locally stable.

However, the behavior of the stochastic model used for the simulations is significantly more complex. First, because $p(y) \rightarrow 1$ as $y \rightarrow 1$, we do not expect the rates to become arbitrarily large when the system is locally unstable. Instead, local instability of the nonlinear system (4)-(5)

is typically visible as large oscillations in the rates. These oscillations are due to the instability caused by the propagation delay D_r .

Because the simulations are stochastic, another significant type of “instability” may be present: a high variance of the rates at equilibrium. It has previously been shown that, for small κ , the variance of the rates scales approximately linearly with κ ([11], Equation 19). For a larger gain κ_r , therefore, we expect stochastic fluctuations to increase the spread of the rates at equilibrium. The stochastic effects of modeling queueing at the resource contribute to random variation in the rates.

Notice that both unstable behaviors will lead to a high sample variance of the rates $x_r[t]$. In the first case, large oscillations in the rates will lead to a large variance. In the second case, stochastic fluctuations in the rates will lead to a large variance. Our goal is to investigate the relative importance of these two types of effects for different roundtrip delays D , and therefore better understand the relationship between propagation delay and queueing delay.

We simulate 10 users with the same roundtrip delay. We let $R = \{1, \dots, 10\}$, and fix $D_r = D$ for all users. We set $x_r = 0.075$, and the threshold $B = 1$. For this threshold, $p(y) = 1 - e^y(1 - y)$, and $p'(y) = ye^y$. We calculate p and p' as described in the last subsection, and let $w_r = x_r p$ for all users. Define $\kappa^*(D)$ to be the critical gain as given by (32):

$$\kappa^*(D) = \frac{2 \sin\left(\frac{\pi}{2(2D+1)}\right)}{p + p'(\sum_{r \in R} x_r)}.$$

We will investigate stability at six different roundtrip delays: $D = 100, 500, 1000, 5000, 10000, 50000$, with $d_1(r) = d_2(r) = D/2$ for all $r \in R$. For each roundtrip delay, we ran simulations where, for all $r \in R$, $\kappa_r = 0.005\kappa^*(D)$, $\kappa_r = 0.05\kappa^*(D)$, $\kappa_r = 0.1\kappa^*(D)$, \dots , $\kappa_r = 1.25\kappa^*(D)$. Thus, there were a total of 26 simulations for each roundtrip delay D .

For each $r \in R$, we initialized the rate to zero: $x_r[0] = 0$. Each simulation was run for $100D$ iterations. Over the last $10D$ iterations, we calculated the sample variance of each rate $x_r[t]$, σ_r^2 . In Figure 4, we have plotted the average variance for each simulation against the gain parameter κ_r , where κ_r is expressed as a multiple of $\kappa^*(D)$.

If the system behaved exactly according to the predictions made in Theorem 9, then we would expect variance to be near zero as long as $\kappa_r < \kappa^*(D)$. However, for $D = 100$, this is obviously not the case. As D increases, we move closer to the theoretical prediction, until at $D = 50000$ the theoretical prediction is fairly accurate. Recall, however, that our fluid flow approximation presumes queueing delays are negligible. In Figure 4 we are observing the effects of modeling both queueing delays and propagation delays.

At a short propagation delay D , i.e., $D = 100$, the critical gain $\kappa^*(D)$ is relatively large. As discussed above, this leads to a high stochastic spread at equilibrium—the random queueing delay is not negligible. As κ_r increases away from zero, the variance rapidly increases. At first, this increase is primarily due to stochastic fluctuations in the rates. For larger κ , instability due to the propagation delay adds to the variance as well.

For larger delays D , however, the critical gain $\kappa^*(D)$ is relatively small. The smaller gain parameter now averages the queueing behavior over a much longer time interval, and this reduces stochastic fluctuations in the rates. For $\kappa_r < \kappa^*(D)$, therefore, neither the queueing delay nor the propagation delay are causing a large variance in the rates. As the gain κ_r increases past the critical gain $\kappa^*(D)$, large oscillations in the rates result. Local instability due to the propagation delay is observed.

Based on the preceding discussion, notice that Figure 4 gives a qualitative explanation of what it means for queueing delays to become negligible. Both propagation delays and queueing delays contribute to a large variance in the rates, but in different ways; queueing delays are negligible when stability of the stochastic simulation is predicted accurately by stability of the deterministic system (4)-(5). In Figure 4, this occurs when (approximately) $D \geq 10000$.

We conclude with a simple example. The United Kingdom academic network, JANET, currently employs two 155 Mbps OC-3 links to the United States. Assuming that packets are 512 bytes, and that the transatlantic roundtrip delay is approximately 200 ms, we find that $D \approx 15000$. In the context of our simulations, for $D = 15000$, queueing delays are already nearly negligible relative to propagation delay. This type of calculation supports the claim made in the Introduction: in the future, we expect queueing delays to become small relative to propagation delays.

5.4 Simulation 2: Instability and phase-locking

As discussed in the last subsection, when local instability results due to the propagation delay, we expect large oscillations; however, as we will discover, these oscillations are not independent of one another. Instead, when all users share the same roundtrip delay, their rates tend to become “phase-locked” [4]: all rates oscillate exactly in phase with each other.

One such example is presented in Figure 5. The parameters for the simulation are $D = 5000$, and $\kappa_r = \kappa^*(D)$ for all $r \in R$. All other parameters are set exactly as in the last subsection, with the exception of the initial rates: we set $x_r[t] = 0.01r$ for each $r \in R$. Notice that, although the users start at different initial rates, they quickly synchronize with each other. The rates become phase-locked, and large oscillations result. This is exactly like resonance in physical systems: when all roundtrip delays are equal, instability takes a very structured form.

It is interesting to note that phase-locking is a phenomenon that is only observed for longer roundtrip delays. At $D = 100$ and $D = 500$, for example, no discernible phase-locking was observed; at $D = 1000$, the effect begins to be observed, but only when $D = 5000$ does phase locking become distinctly visible. For smaller D , queueing delays are not negligible relative to propagation delays. The significant stochastic effects caused by queueing at the resource perturb the system away from a phase-locked equilibrium.

5.5 Simulation 3: Diverse roundtrip delays

Recall that in Section 4, numerical experiments suggested we might be able to extend our results to networks with diverse roundtrip delays (as expressed in Conjecture 12). In our third simulation, we wished to test the impact of sharing a resource between users with different roundtrip delays.

We again shared a resource between ten users, as in Simulation 1. All parameters for the resource are exactly as in Simulation 1. Each user again has $x_r = 0.075$, and $w_r = x_r p$. For each roundtrip delay $D = 100, 500, 1000, 5000, 10000$, we created two users. For each user, we chose κ_r as the largest gain parameter possible such that the standard deviation of the rates from Simulation 1 remained below $0.05x_r$ (see Figure 4). This meant that, for example, $\kappa_r = 0.05\kappa^*(D)$ for $D = 100$, but that $\kappa_r = 0.9\kappa^*(D)$ for $D = 10000$. We then ran the simulation for 100000 iterations.

The presence of multiple roundtrip delays did not significantly affect the behavior of the system. In all cases, the users converged (approximately) to the equilibrium value of $x_r = 0.075$. The

standard deviation of the rates did not exceed 5.8% of x_r for any of the users, suggesting that the predictions of Simulation 1 remain robust even in the presence of diverse delays.

The results of this simulation are important when we consider the implementation of our theoretical results in future networks. The simulations in this section and numerical experiments in Section 4 support the claim that, ultimately, stability can be governed by local control of gain parameters. As long as stability conditions are known for a given roundtrip delay, this simulation suggests those conditions may be applied to ensure stability with diverse roundtrip delays.

6 Conclusion

This paper has considered stability of network rate control in the presence of communication delays. Stability conditions are given for both a single resource and a large network with constant roundtrip delay, and under the addition of instantly adaptive and completely nonadaptive users. Rate of convergence is completely characterized for the single resource case. Finally, simulation results contrast the effects of queueing delays and propagation delays, and also reveal the phenomenon of phase-locking.

Notice that Theorems 9-11 allow for a decentralized implementation by route. If we define $y_j = \sum_{s:j \in s} x_s$, and assume that $\sum_{j \in r} p_j \leq \alpha$ for all $r \in R$ and $y_j p'_j / p_j \leq \beta$ for all $j \in J$, then the LHS of (22) is $\leq \kappa_r \alpha (1 + \beta)$ no matter how many resources lie on each route r . In other words, the results we have found are independent of routing and scale well as the number of resources grows—especially important for a constantly evolving network on a global scale, such as the Internet.

With the implementation described in the last paragraph, the condition (6) becomes even simpler: the “route-dependent constant” on the right-hand side becomes route-independent. Thus, to remain stable, each user will only be asked to keep their gain inversely proportional to the roundtrip delay of their route, with the constant of proportionality *specified in advance*.

If we interpret $\sum_{j \in r} p_j$ to be the equilibrium probability a packet on route r is marked, then $\sum_{j \in r} p_j \leq 1$, so we may set $\alpha = 1$. However, we must consider more closely the uniform bound $y_j p'_j / p_j \leq \beta$ for every resource in the network. Although we can have a decentralized implementation, there is a design requirement for the feedback function at every resource. As an example, consider a resource j which behaves as an M/M/1 queue, with service times exponentially distributed with mean C_j (so C_j is the “capacity” of the resource). Suppose the resource is in equilibrium, and the workload arriving at the resource is Poisson of rate y_j . Further, suppose that the resource uses a threshold marking scheme: a packet is marked if it arrives to find more than B_j packets in the queue. The function p_j is given by the probability that an arriving packet is marked. By the PASTA property (Poisson Arrivals See Time Averages), we can easily calculate $p_j(y_j)$ from the stationary distribution of an M/M/1 queue:

$$p_j(y_j) = \sum_{i \geq B_j} \left(\frac{y_j}{C_j} \right)^i \left(1 - \frac{y_j}{C_j} \right) = \left(\frac{y_j}{C_j} \right)^{B_j}.$$

Thus, we have the result that $y_j p'_j / p_j = B_j$, and the condition $y_j p'_j / p_j \leq \beta$ becomes simply $B_j \leq \beta$. Implementation of this constraint on the threshold level B_j requires no knowledge of user flows, as summarized by y_j . The preceding discussion highlights the importance of the key results Theorems 9-11: design of a network can be simplified by removing central control. Instead, all

users are given a constraint ($\kappa_r D_r < \text{constant}$); and each resource is given a constraint ($y_j p'_j / p_j < \text{constant}$). When both these design conditions are met, local stability of the systems studied can be guaranteed.

Under the assumption that propagation delays will become more significant than queueing delays, the results of this paper suggest that network stability can be guaranteed by consideration of a delayed fluid flow model. The evolution of telecommunications networks in the next decade remains a contested and unresolved issue. However, as networks expand to the size and complexity of the Internet, we may safely conclude theoretical predictions such as those given here will become increasingly important to guarantee robust behavior.

Acknowledgements

Ramesh Johari was supported by a British Marshall Scholarship. David Tan was supported by St. John's College, Cambridge. We are grateful to Frank Kelly for his support, enthusiasm, and guidance; we also wish to acknowledge helpful conversations with Damon Wischik and Andrew Neitzke.

References

- [1] Bellman, R., and Cooke, K.L. (1963). *Differential-Difference Equations*. Academic Press: New York.
- [2] Chong, S., Nagarajan, R., and Wang, Y.T. (1998). Designing stable ABR flow control with rate feedback and open loop control: first-order control case. *Performance Evaluation* **34**: 189-206.
- [3] Chong, S., Nagarajan, R., and Wang, Y.T. (1998). First-order rate-based flow control with dynamic queue threshold for high-speed wide-area ATM networks. *Computer Networks and ISDN Systems* **29**: 2201-2212.
- [4] Floyd, S., and Jacobson, V. (1992). On traffic phase effects in packet-switched gateways. *Internetworking: Research and Experience* **3**, no. 3: 115-156.
- [5] Gibbens, R.J., and Kelly, F.P. (1999). Resource pricing and the evolution of congestion control. *Automatica* **35**: 1969-1985.
- [6] Gross, D., and Harris, C.M. (1985). *Fundamentals of Queueing Theory*. Wiley: New York.
- [7] Hayes, N.D. (1950). Roots of the transcendental equation associated with a certain differential-difference equation. *Journal of the London Mathematical Society* **25**: 226-232.
- [8] Horn, R.A., and Johnson, C.R. (1985). *Matrix Analysis*. Cambridge University Press: Cambridge, U.K.
- [9] Jacobson, V. (1988). Congestion avoidance and control. In: *Proceedings of ACM SIGCOMM '88*, pp. 314-329.
- [10] Kelly, F.P. (2000). Models for a self-managed Internet. To appear in: *Philosophical Transactions of the Royal Society of London A 2000*.
- [11] Kelly, F.P., Maulloo, A.K., and Tan, D.K.H. (1998). Rate control in communication networks: shadow prices, proportional fairness, and stability. *Journal of the Operational Research Society* **49**: 237-252.
- [12] Luenberger, D.G. (1979). *Introduction to Dynamic Systems: Theory, Models, and Applications*. Wiley: New York.
- [13] Rice, J.A. (1995). *Mathematical Statistics and Data Analysis*. Duxbury Press: Belmont, CA.

PSfrag replacements

$x[t]$




Figure 1: Single resource and single user.




Figure 2: Plot of the function $y = -\lambda e^\lambda$.




Figure 3: Rate of convergence versus gain.




Figure 4: Variance of rates versus gain, for different roundtrip delays.

g replacements




Figure 5: Phase-locking of rates.