

DETERMINISTIC FLUID MODELS OF CONGESTION CONTROL IN HIGH-SPEED NETWORKS

Sanjay Shakkottai
R. Srikant

Coordinated Science Laboratory
University of Illinois at Urbana-Champaign, U.S.A

ABSTRACT

Congestion control algorithms, such as TCP or the closely-related additive increase-multiplicative decrease algorithms, are extremely difficult to simulate on a large scale. The reasons for this include the complexity of the actual implementation of the algorithm and the randomness introduced in the packet arrival and service processes due to many factors such as arrivals and departures of sources and uncontrollable short flows in the network. To make the simulation tractable, often deterministic fluid approximations of these algorithms are used. These fluid approximations are in the form of deterministic delay differential equations. In this paper, we ignore the complexity introduced by the window-based implementation of such algorithms and focus on the randomness in the network. We justify the use of deterministic models for proportionally-fair congestion controllers under a limiting regime where the number of sources in a network is large.

1 INTRODUCTION AND MODEL

There has been a lot of recent work on decentralized end-to-end congestion control algorithms for the Internet. These are based on ECN marking with the goal of building a low-loss, low-queueing-delay network. The control algorithms are designed on the premise that each user has a utility function, which the user is trying to maximize, while the network is simultaneously trying to maintain some sort of fairness amongst various users. In the algorithms proposed, the network tries to achieve its goal by *marking* packets during congestion (see Ramakrishnan and Floyd (1999), Kunniyur and Srikant (2000)). The notion of fairness (from the network's point of view) which has been used is weighted *proportional fairness* (see Kelly (1997)). Through appropriate choice of the weights, other fairness criteria such as minimum potential delay fairness can be realized. If we interpret the utility function of the user as the users' willingness to pay for bandwidth, and suppose

that the price paid by the user is proportional to the number of marks received, then, a weighted proportionally fair scheme leads to *same price per unit bandwidth* paid by any user for utilizing some resource in the network. The algorithms proposed have a *decentralized* implementation to achieve the network and user objectives simultaneously. In this paper, we focus on the case where all users have the same utility function of $\log(x)$, leading to a proportionally fair sharing of the bandwidth.

There have been various delay differential equation models for Internet congestion control. The generic model of such a system consists of a source, which sends data at rate $x(t)$, a router which signal congestion to sources by marking packets, and a receiver which detects the marks and informs the source to increase or decrease its transmission rate. Associated with the router is a marking function, which marks the a fraction of the flow as a function of the total arrival rate. The larger the fraction is, the more aggressively the source "backs off". Such scenarios have been studied (Kelly (1997), Kelly (1999), Kunniyur and Srikant (2000), Kunniyur and Srikant (2001)), and in the absence of delays, differential equation models of such systems have been shown to converge. With delays, bounds have been derived on the behavior of the flow (Shakkottai, Srikant and Meyn (2001)). However, in a realistic scenario, we have short flows which do not adapt, thus causing "noise" at the router. Further, the marking function could base its decision on more than the instantaneous arrival rate. A formal justification of how delay differential equations models correspond to "real" systems does not seem to exist in current literature.

In this paper, we study "noisy" congestion control algorithms and show that the deterministic differential equations that have been studied earlier in the literature are appropriate limits in a many flows regime.

Related work includes that of Vojnovic, Le Boudec and Boutremans (2000), where a stochastic approximations (see Kushner and Yin (1997)) based model is considered. The authors study the rate process for additive increase, multiplicative decrease (AIMD) algorithms in the rare negative

feedback regime. Under the assumption of small gains (Δ and β defined in Section 1.1 are small), they show that an asynchronous implementation of a generic AIMD converges to an ordinary differential equation. However, the small gain assumption is not valid in practice. Further, this approach leads to a fluid limit which does not capture the oscillations due to delayed feedback. In fact, it is known that the rate control algorithms do not always converge when delays are present (see Johari and Tan (2000), Massoulié (2000), Shakkottai, Srikant and Meyn (2001), Vinnicombe (2001)). We believe that a justification of the delay-differential model does not come from the previous approach, but rather, from a many-sources approximation, where we scale the capacity of the system along with the number of flows. We will show that in this regime, the delay-differential equations are suitable limits of such a system. Our approach can be thought of as a functional-differential equation analog of “averaging” used for ordinary differential equations as in Kurtz (1970) and Khalil (1996).

We next describe a simple rate based marking model marking model for a single link and all flows having the same round trip delay. In a longer version of this paper, available from the authors Shakkottai and Srikant (2001), we have studied variants of these models, such as networks with a general topology, different round-trip delays, “random” marking functions, queue based marking, etc.

1.1 RATE BASED MARKING WITH NOISE

We consider a sequence of systems (indexed by n), where the first system is the following. A single best-effort flow accesses a link of capacity c . We first fix some (integer) time $T > 0$, and the system is assumed to evolve in discrete time-steps. At each time $i = 0, 1, \dots, T$, the user adapts its transmission rate x_i depending on the feedback it receives from the router. The router marks some amount of the flow it receives, and this amount is proportional to the user transmission rate. In practice, for a packet based system, such marking could be implemented using ECN marks (see Ramakrishnan and Floyd (1999)). For a fluid model such as ours, we assume that some volume of the fluid is marked (see Kelly (1997), Kelly (1999), Kunniyur and Srikant (2000)). The fraction of fluid marked is determined by means of a marking function, about which we will discuss more later. We assume that there is a round trip delay between the source and router of $d \in \mathcal{Z}$. Thus, the rate at time $i + 1$ depends on the amount marked at the router half a round trip back, which in-turn, depends on the user transmission rate a further half round-trip time back. Thus, we can describe the evolution of the user transmission rate by

$$x_{i+1} = (x_i + \Delta - \beta x_{i-d} p(x_{i-d} + \hat{e}_{i-d}))^+$$

where Δ, β are positive constants which determine the rate at which source increases or decreases its transmission rate; $p(\cdot)$ is the marking function, and \hat{e}_i is a “noise” process.

In this case, we assume that two flows access the router: the first is the user’s data flow, which is represented by $\{x_i\}$ and the other flow is an uncontrolled flow, possibly generated by some other short-duration flows, popularly known as web-mice, passing through the link, which is represented by the sequence $\{\hat{e}_i\}$. The sequence $\{\hat{e}_i\}$ is assumed to be a stochastic process, with $E(\hat{e}_1) = a > 0$. Let $\tilde{e}_i = \hat{e}_i - a$. We assume that the “noise” process \tilde{e}_i is a bounded, stationary-ergodic zero-mean stochastic process for $i \geq 0$; and for $i \leq 0$, $\tilde{e}_i = \tilde{e}_0$. The mean initial conditions (i.e., for $t \leq 0$) is given by sampling $\theta(t)$, $-T \leq t \leq 0$, where $\theta(t)$ is a non-negative, bounded, Lipschitz continuous trajectory. Thus, the initial conditions for above system is given by $\theta(-i) + \tilde{e}_0$. Then, we have the dynamics of the system governed by

$$x_{i+1} = (x_i + \Delta - \beta x_{i-d} p(x_{i-d} + a + \tilde{e}_{i-d}))^+$$

with $x_{-i} = \theta(-i) + \tilde{e}_0$.

Finally, we comment on the marking function itself. This function is based on the total data rate accessing the router and determines the fraction of flow to be marked, and satisfies the following criteria:

Assumption 1.1. We assume that $p(\cdot)$ satisfies

- (i) $0 \leq p(x) \leq 1$
- (ii) $p(x) = 0$ for $x < 0$. Further, there exists a δ satisfying $0 \leq \delta \ll \Delta$ such that $|\beta p(x)| < \delta/a$ for $x < a$.
- (iii) $p(x)$ is an increasing function.
- (iv) $p(x)$ is Lipschitz continuous.

The first property is obvious, as the marking function represents the fraction of flow marked. To understand the second property, we first note that a , the mean arrival rate of the uncontrolled flows, is typically less than 25% of the link capacity. Thus, condition (ii) expresses the intuitive reasoning that, if the total arrival rate (i.e., sum of arrival rates of uncontrolled and controlled rates) at a link is less than 25%, no congestion indication should be provided. The third property is again clear: the larger the arrival rate is, the greater is the fraction marked. Finally, the last condition is a technical condition, which says that the function is “smooth”. As an example, a possible rate based marking function is of the form

$$p(x) = \frac{(x - \bar{c})^+}{x}$$

In a deterministic fluid model, this has the interpretation of the fraction of fluid lost when the arrival rate exceeds the

“virtual” capacity $N\tilde{c}$. It has been noted in Kunniyur and Srikant (2000) that this function is the marking function derived by taking the limit of the $M/M/1/B$ loss formula, when we scale the arrival rate, target arrival rate (Nc) and buffer size simultaneously. It is designed such that in the absence of delays and noise (i.e., there is no fluctuation in the uncontrolled flows), the flow rate x_i converges to $(c - a)$. Thus, $p(\cdot)$ satisfies

$$\Delta = \beta(c - a)p(c).$$

In the example marking function considered earlier, suppose we wanted the source rate to be $x = c - a$, then \tilde{c} should be chosen to satisfy

$$\tilde{c} = c - \frac{\Delta c}{\beta(c - a)} \tag{1}$$

and should also satisfy $\tilde{c} \geq a$. Our goal in this study is to determine when the above stochastic system can be approximated by the following deterministic system:

$$\dot{x} = \Delta - \beta x(t - d)p(x(t - d) + a) \tag{2}$$

with initial conditions given by $\theta(t), 0 \geq t \geq -T$. The congestion control algorithm corresponds to a resource allocation problem where all flows have $\log(\cdot)$ utility functions (see Kelly, Maulloo and Tan (1998)). As an aside, we note that the local stability of this equation (i.e., whether $x(t) \rightarrow c - a$ for a linearized system) has been studied in Johari and Tan (2000) and Massoulié (2000), while global boundedness has been studied in Shakkottai, Srikant and Meyn (2001).

Our objective here is to show that when we scale the system such that many flows access the node and the capacity of the link is nc , the above delay-differential equation is a close approximation of the stochastic delay-difference equation. We have so far described the system model when a single flow accesses the link. We now describe how the model scales in n . First, there are n flows, with link capacity nc . For every time step in the first model, we assume that there are n time-steps in the n th model. This represents the fact that we need to increase the time resolution to study n processes. In practice, we can view each time slot at a measurement interval over which rates are measured in the system and control actions by the routers and sources are updated. Typically, this measurement interval is measured in terms of the number of packets that can be processed by a typical router. For example, the time-step could be “100 packets long.” By scaling both the time-step and the capacity, we maintain a constant time-step, as measured in packets. The flows are now represented by $\{y_i^k\}, k = 1, 2, \dots, n$ where the subscript $i \in \{0, 1, \dots, nT\}$ represents the time-index, and the superscript k represents the flow index. As

the delay seen by a user does not change with n but the time resolution increases in n , it follows that the delay (measured in index steps) scales in n . Thus, (ignoring non-negativity constraints, which is discussed later) each flow y_i^k evolves according to

$$y_{i+1}^k = y_i^k + \frac{\Delta}{n} - \frac{\beta}{n} y_{i-nd}^k \tag{3}$$

$$p \left(\frac{1}{n} \sum_{j=1}^n (y_{i-nd}^j + \tilde{e}_{i-nd}^j + a) \right).$$

As in the single flow case, the above equations represent n flows, each which has an additive increase factor of and the backoff for each flow is proportional to the delayed transmission rate, which the fraction marked being a function of the *total arrival rate* to the router. We remark that the above system can be interpreted as a decentralized means of achieving a proportionally fair allocation (see Kelly (1997)) of bandwidth among the n users.

We note three things in the above equation. First, the gain constants Δ and β are scaled by n . This is because the time step in now $\frac{1}{n}$ of the first system. Thus, the gains are also scaled to maintain the same gain for each flow over the original time-step.

Second, there are now n uncontrolled flows (i.e., $\{e_i^k\}, k = 1, 2, \dots, n, i \geq 0$) accessing the system, with each flow having a bounded rate, i.e., $|e_i^k| < K$. We assume that these flows are iid (across flows, although these could be correlated in time), and are stationary and ergodic. For $i < 0$, we assume a constant (random) initial condition, i.e., $e_{-i}^k = e_0^k$. Then, we have for any $\epsilon > 0$,

$$\Pr \left(\sup_{-nT \leq i \leq nT} \left| \frac{1}{n} \sum_{j=1}^n \tilde{e}_i^j \right| > \epsilon \right)$$

$$\leq (nT + 1) \Pr \left(\left| \frac{1}{n} \sum_{j=1}^n \tilde{e}_0^j \right| > \epsilon \right)$$

$$\leq (nT + 1) e^{-na}$$

for some $a = a(\epsilon) > 0$. The first step follows from the union bound, and the fact that the flows are stationary. The second step (i.e., existence of $a > 0$) follows from Chernoff’s bound (it trivially follows that the MGF exists, as the process is bounded). Thus, it follows that

$$\Pr \left(\sup_{-nT \leq i \leq nT} \left| \frac{1}{n} \sum_{j=1}^n \tilde{e}_i^j \right| > \epsilon \right) \xrightarrow{n \rightarrow \infty} 0$$

exponentially fast. Now, applying Borel-Cantelli Lemma, it follows that

$$\sup_{-nT \leq i \leq nT} \left| \frac{1}{n} \sum_{j=1}^n \tilde{e}_i^j \right| \rightarrow 0 \text{ a.s.} \quad (4)$$

We remark that the above model allows each “noise” process to be long-range dependent. We only need the flows to be bounded and iid.

Third, the marking function is seen to scale its argument by n (the marking function acts on the *average* arrival rate as opposed to the *total* arrival rate), i.e., if the marking function for the n th system were to be represented by $p^n(x)$, then, we would have

$$p^n(nx) = p(x).$$

This is done so that for each n , if the centered error processes $\{\tilde{e}_i^k\}, k = 1, \dots, n$ were identically zero, the delay-difference equations (if they converge) would have a steady state value of $\lim_{k \rightarrow \infty} y_i^k = c - a$. We remark that *adaptive* marking functions have been proposed (see Kunniyur and Srikant (2001)) which automatically scale in n as described above, without explicitly having knowledge of n .

Now, let x_i^n represent the average rate at time i , and e_i^n represent the average (centered) noise at time i , i.e.,

$$x_i^n = \frac{1}{n} \sum_{j=1}^n y_i^j$$

$$e_i^n = \frac{1}{n} \sum_{j=1}^n \tilde{e}_i^j.$$

Then, (by adding the various equations for y_i^j), we have that x^n satisfies the following stochastic delay-difference equation:

$$x_{i+1}^n = x_i^n + \frac{\Delta}{n} - \frac{\beta}{n} x_{i-nd}^n p(x_{i-nd}^n + a + e_{i-nd}^n)$$

Now, we embed the above equation in “continuous-time”, i.e., we study the above process over an interval of time $[0, T]$ (without loss of generality, assume that $T = kd$, for some $k > 1$). For $nt \in \mathcal{Z}$, we let

$$x^n(t) = x_{nt}^n$$

$$e^n(t) = e_{nt}^n$$

and use a straight line approximation to interpolate between the times $t = \frac{i}{n}$. Thus, we see that the above equation can

be represented by the equation

$$\dot{x}^n(t) = \Delta - \beta x^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right) p\left(x^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right) + a + e^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right)\right). \quad (5)$$

Further, assume that the initial condition is given by

$$x^n(t) = \theta(t) + e^n(0) \quad -T \leq t \leq 0, \quad nt \in \mathcal{Z}$$

and a straight line interpolation is used for $nt \notin \mathcal{Z}$. This means that each flow $\{y^k\}, k = 1, 2, \dots, n$ could have a different initial condition, but the nominal initial condition is given by $\theta(t)$. Also, note that from (4), it follows that $e^n(t) \rightarrow 0$ uniformly in $[-T, T]$. We note that the above differential equation is to be interpreted as a shorthand notation to represent the process $x^n(t)$ given by the unique trajectory solving the integral equation

$$x^n(t) = x^n(0) + \int_{s=0}^t \Delta - \beta x^n\left(\frac{\lfloor n(s-d) \rfloor}{n}\right) p\left(x^n\left(\frac{\lfloor n(s-d) \rfloor}{n}\right) + a + e^n\left(\frac{\lfloor n(s-d) \rfloor}{n}\right)\right) ds$$

In this paper, unless otherwise stated, any differential equation is to be interpreted as a representation of the corresponding integral equation.

We note that as this model has delays, the above model allows $x^n(t)$ to be negative as well. From a practical standpoint, we know that the rates are non-negative. For now, we will ignore this constraint. In Shakkottai and Srikant (2001), we show that for n sufficiently large, the rates $x^n(t)$ driven by the above equation will be non-negative (under suitable initial conditions). Thus, the constraint $x^n(t) \geq 0$ will be redundant, and hence, in the many-flows regime, the above system is a valid model of reality.

Our objective is to show that the trajectory generated by (5) and that in (2) are “close” in some suitable sense. In the next section, we will state a general convergence result for functional differential equations. We will then show that the system we have described fits into this framework, and the desired results can be proved.

2 CONVERGENCE OF FUNCTIONAL DIFFERENTIAL EQUATIONS

Let \mathcal{R}^2 be endowed with the \mathcal{L}_1 norm, and $C([0, T], \mathcal{R}^2)$ be the space of continuous, \mathcal{R}^2 valued functions on $[0, T]$ with the supremum norm. We denote any element of $C([0, T], \mathcal{R}^2)$ by the tuple (ϕ, ψ) , and $\|(\phi, \psi)\| = \sup_{t \in [0, T]} (|\phi(t)| + |\psi(t)|)$. Now, consider a sequence of functionals $b_n : C([0, T], \mathcal{R}^2) \mapsto \mathcal{R}$, such that $\{b_n\}$ are

Lipschitz continuous and bounded with parameters L and M respectively, i.e.,

Assumption 2.1.

$$\begin{aligned} |b_n(\bar{\phi}, \bar{\psi})| &\leq M \\ |b_n(\bar{\phi}_1, \bar{\psi}_1) - b_n(\bar{\phi}_2, \bar{\psi}_2)| &\leq L \|(\bar{\phi}_1, \bar{\psi}_1) - (\bar{\phi}_2, \bar{\psi}_2)\| \end{aligned}$$

Let $C([0, T], \mathcal{R})$ be the space of continuous, real valued functions on $[0, T]$, and endowed with the sup topology. Fix $a, b, c > 0$ and let $\tilde{C}[0, T]_{a,b,c} = \{\bar{x}_t \in C([0, T], \mathcal{R}) : \bar{x}_t \text{ is Lipschitz continuous with parameter } a, x(t) \geq -c \text{ and } x(t) \leq b\}$. As an aside, it can be shown that $\tilde{C}[0, T]_{a,b,c}$ is a compact subset of $C([0, T], \mathcal{R})$.

Now, we consider the following stochastic functional differential equations (SFDEs). Let $x^n(t)$ be the unique, continuous solution of the following FDE. By the conditions imposed on b_n , there exists such a solution Hale (1977) (the Caratheodory conditions). For $t \in [0, T]$, consider

$$\dot{x}^n(t) = b_n\left(\bar{x}_{\lfloor nt \rfloor / n}^n, \bar{\mathbf{0}}\right)$$

where $\bar{x}_t^n \triangleq [x^n(s), s \in [t - T, t]] \in C([0, T], \mathcal{R})$, $\bar{\mathbf{0}}$ is a process identically equal to zero in $[-T, T]$, and the initial condition is given by

$$x^n(t) = \theta(t) \quad -T \leq t \leq 0, \quad nt \in \mathcal{Z}$$

and a straight line interpolation is used for $nt \notin \mathcal{Z}$.

Assume $|\theta(\cdot)| < A$, and that $\theta(t)$ is Lipschitz continuous with parameter M . Then, it follows $x(t)$ is Lipschitz continuous with Lipschitz constant M and bounded (with bounds $\pm(A + MT)$).

Next, consider for $t \in [0, T]$, the SFDE

$$\dot{\mathbf{x}}^n(t) = b_n\left(\bar{\mathbf{x}}_{\lfloor nt \rfloor / n}^n, \bar{\mathbf{e}}_{\lfloor nt \rfloor / n}^n\right)$$

here $\bar{\mathbf{x}}_t^n \in C[0, T]$, is given by $\bar{\mathbf{x}}_t^n = [\mathbf{x}^n(s), s \in [t - T, t]]$, $\bar{\mathbf{e}}_t^n \in C([0, T], \mathcal{R})$ is given by $\bar{\mathbf{e}}_t^n = [\mathbf{e}^n(s), s \in [t - T, t]]$, and the initial condition is given by

$$x^n(t) = \theta(t) + \mathbf{e}^n(t) \quad -T \leq t \leq 0, \quad nt \in \mathcal{Z}$$

and a straight line interpolation is used for $nt \notin \mathcal{Z}$.

$[\mathbf{e}^n(t), -T \leq t \leq T]$ is a zero mean “error” process, satisfying the following condition.

Assumption 2.2. Assume that $\exists K > 0$ such that $|\mathbf{e}^n(t)| < K$ uniformly in n , and satisfying the following condition:

$$\sup_{t \in [-T, T]} |\mathbf{e}^n(t)| \xrightarrow{n \rightarrow \infty} 0 \quad \text{a.s.}$$

Now, let $b : C([0, T], \mathcal{R}^2) \mapsto \mathcal{R}$ satisfy the following condition.

Assumption 2.3. Assume that $b(\cdot, \cdot)$ is a Lipschitz continuous and bounded function with parameters L, M respectively (without loss of generality, assume $L \geq 1$), satisfying

$$\sup_{\bar{x}_t \in \tilde{C}[0, T]_{M, A+MT, -(A+MT)}} |b_n(\bar{x}_t, \bar{\mathbf{0}}) - b(\bar{x}_t, \bar{\mathbf{0}})| \xrightarrow{n \rightarrow \infty} 0$$

Finally, for $t \in [0, T]$, consider the FDE

$$\dot{x}(t) = b(\bar{x}_t, \bar{\mathbf{0}})$$

where $\bar{x}_t \triangleq [x(s), s \in [t - T, t]] \in C[0, T]$, and the initial condition is given by

$$x(t) = \theta(t) \quad -T \leq t \leq 0.$$

For n large, it seems reasonable to believe that the trajectories of $x(t)$, $x^n(t)$ and $\mathbf{x}^n(t)$ are “close”. The following result has been proved in Shakkottai and Srikant (2001).

Theorem 2.1.

$$\sup_{t \in [0, T]} |x(t) - \mathbf{x}^n(t)| \xrightarrow{n \rightarrow \infty} 0 \quad \text{a.s.}$$

3 RATE BASED MARKING AND THE MANY FLOWS LIMIT

In this section, we show that in the many-flows regime, the trajectories for the stochastic and deterministic rate-based marking differential equations introduced earlier are “close”. We recall from the model description in Section 1.1 that we currently ignore the non-negativity constraint on the flows. This issue is handled in Shakkottai and Srikant (2001), where we show that for “reasonable” initial conditions, the trajectories remain non-negative.

Formally, we consider the delay-differential equations

$$\begin{aligned} \dot{x}^n(t) &= \Delta - \beta x^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right) \\ & p\left(x^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right) + a + e^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right)\right) \end{aligned} \tag{6}$$

with initial conditions given by

$$x^n(t) = \theta(t) + e^n(0) \quad -T \leq t \leq 0, \quad nt \in \mathcal{Z}$$

and a straight line interpolation is used for $nt \notin \mathcal{Z}$. The candidate limiting system is described by

$$\dot{x}(t) = \Delta - \beta x(t-d)p(x(t-d) + a)$$

with initial conditions given by $\theta(t)$. Using Theorem 2.1, we show the following result.

Lemma 3.1. *For the systems described above, we have*

$$\sup_{t \in [0, T]} |x^n(t) - x(t)| \rightarrow 0 \text{ a.s.}$$

as $n \rightarrow \infty$.

Proof. We prove some properties of (6). First, we note that as $|e_i^n| < K$, we have that if $a + x^n(\frac{\lfloor n(t-d) \rfloor}{n}) < -K$, $p\left(x^n(\frac{\lfloor n(t-d) \rfloor}{n}) + a + e^n(\frac{\lfloor n(t-d) \rfloor}{n})\right) = 0$ (as $p(z) = 0$ for $z < 0$). Thus, we have that

$$x^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right) p\left(x^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right) + a + e^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right)\right) \geq -(K+a)$$

and

$$\dot{x}^n(t) \leq \Delta + \beta(K+a).$$

As we are studying the process over the interval $[0, T]$, and $|x^n(0)| < K + \theta(0) \triangleq A$, we can uniformly upper bound (in n) $x^n(t)$ by $A + (\Delta + \beta(a+K))T$. Next, we lower bound $x^n(t)$. As $p(z) \leq 1$, and using our upper bound, we have

$$\dot{x}^n(t) \geq -\beta[A + (\Delta + \beta(a+K))T]$$

Thus, we have

$$x^n(t) \geq \theta(0) - K - T\beta[A + (\Delta + \beta(a+K))T].$$

Next, we define

$$M_x = \max\left(\frac{|\theta(0) - K - T\beta[A + (\Delta + \beta(a+K))T]|}{A + (\Delta + \beta(a+K))T}, \right)$$

$$L_x = \max(\beta[A + (\Delta + \beta(a+K))T], \Delta + \beta(K+a)).$$

We observe that $|x^n(t)| \leq M_x$ and is Lipschitz continuous with parameter L_x . Let $M = \max(K, M_x)$. For any $z \in \mathcal{R}$, define $\lfloor z \rfloor = (-M) \vee (z \wedge M)$. Define

$$b(x_t, e_t) = \Delta - \beta[x(t-d)] p([x(t-d)] + [e(t-d)] + a).$$

As we note that in (5), $|x^n(t)| \leq M$ and $|e^n(t)| \leq M$, we can rewrite it as

$$\begin{aligned} \dot{x}^n(t) &= \Delta - \beta\left[x^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right)\right] \\ &\quad p\left(\left[x^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right)\right] + a + \left[e^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right)\right]\right) \\ &= b\left(x_{\lfloor nt \rfloor}^n, e_{\lfloor nt \rfloor}^n\right). \end{aligned} \tag{7}$$

Now, we show that $b(\cdot, \cdot)$ is bounded and Lipschitz continuous. First, as $0 \leq p(\cdot) \leq 1$, we have

$$\begin{aligned} |b(x_t, e_t)| &= |\Delta - \beta[x(t-d)] \\ &\quad p([x(t-d)] + [e(t-d)] + a)| \\ &\leq \Delta + \beta M \end{aligned}$$

Next, we have

$$\begin{aligned} |b(x_t, e_t) - b(y_t, r_t)| &= |\beta[y(t-d)] p([y(t-d)] + [r(t-d)] + a) \\ &\quad - \beta[x(t-d)] p([x(t-d)] + [e(t-d)] + a)| \\ &\leq \beta|[y(t-d)] p([y(t-d)] + [r(t-d)] + a) \\ &\quad - [x(t-d)] p([y(t-d)] + [r(t-d)] + a) \\ &\quad + [x(t-d)] p([y(t-d)] + [r(t-d)] + a) \\ &\quad - [x(t-d)] p([x(t-d)] + [e(t-d)] + a)|. \end{aligned}$$

As $p(\cdot)$ is Lipschitz continuous with parameter P , and $|x| \leq M$, we have

$$\begin{aligned} |b(x_t, e_t) - b(y_t, r_t)| &\leq \beta|[y(t-d)] - [x(t-d)]| \\ &\quad + \beta PM(|[y(t-d)] - [x(t-d)]| \\ &\quad + |[e(t-d)] - [r(t-d)]|) \\ &\leq \beta(1 + PM)|x(t-d) - y(t-d)| \\ &\quad + \beta PM|e(t-d) - r(t-d)| \\ &\leq \beta(1 + PM)||(x_t, e_t) - (y_t, r_t)||. \end{aligned}$$

Thus, the conditions for Theorem 2.1 are satisfied, and we have the desired result, i.e., as $n \rightarrow \infty$, we have

$$\sup_{t \in [0, T]} |x^n(t) - x(t)| \rightarrow 0 \text{ a.s.}$$

◇

Next, we study a similar limit for *each flow* as opposed to the aggregate flow. Recall that the rate of flow i (when there are n flows in the system) adapts according to (3), which can be represented by

$$\begin{aligned} \dot{y}^{i,n}(t) &= \Delta - \beta y^{i,n}\left(\frac{\lfloor n(t-d) \rfloor}{n}\right) \\ &\quad p\left(x^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right) + a + e^n\left(\frac{\lfloor n(t-d) \rfloor}{n}\right)\right) \end{aligned} \tag{8}$$

with the initial conditions given by $\theta(t) + \tilde{e}_0^i$, sampled appropriately and interpolated. It can be shown using an

analysis similar to that carried out for the aggregate flow that as $n \rightarrow \infty$, the flow trajectory approaches the trajectory of the following delay-differential equation

$$\dot{y}^i(t) = \Delta - \beta y^i(t-d)p(x(t-d)+a) \quad (9)$$

and with *random* initial conditions given by $\theta(t) + \tilde{e}_0^i$, with $|\tilde{e}_0^i| < K$.

In Shakkottai and Srikant (2001), we address the issue of non-negativity of the trajectories of $y^{i,n}(t)$, $x^n(t)$ and $x(t)$. We show that under reasonable initial conditions, for n large enough, the trajectories will remain non-negative. We also deal with other marking schemes, such as those based on queue-length. A discussion of general network topologies and multiple delays is also available there.

REFERENCES

- Hale, J. 1977. *Theory of Functional Differential Equations*, New York: Springer-Verlag.
- Johari, R. and D. Tan. 2000. End-to-end congestion control for the internet: delays and stability. preprint. A shorter version was presented at the *Workshop on Stochastic Networks*, Madison, WI.
- Kelly, F.P. 1997. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, vol. 8, pp. 33–37.
- Kelly, F. P. 1999. Mathematical modeling of the Internet. *Proc. of the Fourth International Congress on Industrial and Applied Mathematics*, Edinburgh.
- Kelly, F. P. , A. Maulloo, and D. Tan. 1998. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252.
- Khalil, H. 1996. *Nonlinear Systems*. Prentice Hall.
- Kunniyur, S. and R. Srikant. 200. End-to-end congestion control: utility functions, random losses and ECN marks. *Proc. IEEE Infocom*, Tel Aviv, Israel.
- Kunniyur, S. and R. Srikant. 2001. A time-scale decomposition approach to adaptive ECN marking. *Proc. IEEE Infocom*, Anchorage, AK.
- Kurtz, T.G. 1970. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, vol. 7, pp. 49–58.
- Kushner, H. and G. Yin. 1997. *Stochastic Approximation Algorithms and Applications*. New York: Springer-Verlag.
- Massoulié, L. 2000. Stability of distributed congestion control with heterogeneous feedback delays. *Microsoft Research Technical Report*.
- Ramakrishnan, K.K. and S. Floyd. 1999. A proposal to add explicit congestion notification (ECN) to IP. *RFC 2481*.
- Shakkottai, S., R. Srikant and S. Meyn. 2001. Bounds for utility function based congestion controllers in the presence of delay. *CSL Technical Report*, University of Illinois. Available online via <<http://www.comm.csl.uiuc.edu/~srikant/pub.html>>.
- Shakkottai, S. and R. Srikant. 2001. Mean FDE models for Internet congestion control under a many-flows regime. *CSL Technical Report*, University of Illinois. Available online via <<http://www.comm.csl.uiuc.edu/~srikant/pub.html>>.
- Vinnicombe, G. 2001. On the stability of end-to-end congestion control for the Internet. *University of Cambridge Technical Report*.
- Vojnovic, M., J. Le Boudec and C. Boutremans. 2000. Global fairness of additive-increase and multiplicative-decrease with heterogeneous round-trip times. *Proc. IEEE Infocom*, Tel Aviv, Israel.

AUTHOR BIOGRAPHIES

SANJAY SHAKKOTTAI obtained his B.E. in Electronics and Communications Engineering from U.V.C.E., Bangalore, and an M.E. in Communications Engineering from the Indian Institute of Science, Bangalore. He is currently pursuing a Ph.D. in Electrical Engineering at the University of Illinois at Urbana-Champaign. shakkott@uiuc.edu

R. SRIKANT received his B. Tech. from the Indian Institute of Technology, Madras, India in 1985, and the M.S. and Ph. D. degrees from the University of Illinois in 1988 and 1991, respectively, all in Electrical Engineering. He was a Member of Technical Staff in AT&T Bell Laboratories from 1991 to 1995. Since August 1995, he has been with the University of Illinois where he is currently an Associate Professor in the Department of General Engineering and Coordinated Science Laboratory, and an Affiliate in the Department of Electrical and Computer Engineering. His research interests include communication networks, queueing theory and stochastic control. He received an NSF CAREER award in 1997. rsrikant@uiuc.edu