

A Non-Equilibrium Analysis & Control Framework for Active Queue Management

Tansu Alpcan¹, Paul Wang², Prashant G. Mehta², and Tamer Başar³
tansu.alpcan@telekom.de, (paulwang, mehtapg)@uiuc.edu, tbasar@control.csl.uiuc.edu

Abstract— We present a non-equilibrium analysis and control approach for the Active Queue Management (AQM) problem in communication networks. Using simplified fluid models, we carry out a bifurcation study of the complex dynamic queue behavior to show that non-equilibrium methods are essential for analysis and optimization in the AQM problem. We investigate an ergodic theoretic framework for stochastic modeling of the non-equilibrium behavior in deterministic models and use it to identify parameters of a fluid model from packet level simulations. For computational tractability, we use set-oriented numerical methods to construct finite-dimensional Markov models including control Markov chains and hidden Markov models. Subsequently, we develop and analyze an example AQM algorithm using a Markov Decision Process (MDP) based control framework. The control scheme developed is optimal with respect to a reward function defined over the queue size and aggregate flow rate. We implement and simulate our illustrative AQM algorithm in the *ns-2* network simulator. The results obtained confirm the theoretical analysis and exhibit promising performance when compared with well-known alternative schemes under persistent non-equilibrium queue behavior.

I. INTRODUCTION

Communication networks such as the Internet exhibit a wide variety of complex dynamic behavior. Examples of such complex behavior include user flow rate oscillations in the presence of delays [1], dynamic synchronization of the flows passing through the same bottleneck link [2], and chaotic behavior of user flows and queues at the routers [3].

The control of complex networks has been a focus of much recent research interest. Much of this line of research, however, has focused on a single-point equilibrium solution and analysis of its stability properties. We note Kelly’s framework for network capacity optimization [4], [5] and game theoretic approaches for network control and optimization [6]–[8]. Each of these approaches lead to a static optimal equilibrium solution. Lyapunov functions [9], [10] and linear control theoretic methods [11]–[14] are then used to ensure its dynamic stability.

This paper is concerned with the analysis and control of complex behavior, referred to as non-equilibrium dynamics, in communication networks. Neither static optimization nor linear control theoretic approaches are suitable for analysis or

control in non-equilibrium settings. For illustrative purposes, we consider here the Active Queue Management (AQM) problem. The AQM provides a mechanism by which a link (router) sends advanced congestion notification to the users. In particular, an AQM algorithm uses the queue length information to either mark or drop packets. The latter is the case in the widely-used droptail algorithm. Random Early Detection (RED) [15] and its variations such as AVQ [16], REM [17], BLUE [18], and E-RED [19] are other well-known examples of AQM algorithms [9], [20] with different characteristics, which have been proposed and studied by the research community.

In [11], a fluid-flow model of TCP interacting with AQM schemes is linearized around the equilibrium. Then, the AQM analysis and design is formulated as a linear control problem whose stability properties are investigated. Linear stability properties of networks with TCP-RED interaction has been studied in [13]. A mean-field model based on N-particle Markov process and for the congestion windows of multiple TCP sources multiplexed through a buffer implementing RED has been presented in [21]. Through simulation studies and an asymptotic analysis the applicability of the model is established. In addition, stability and robustness properties of the resulting system with respect to time delays have been identified. A nonlinear and bifurcation analysis of RED on a TCP network has been conducted in [22] where a discrete-time dynamical model is used to analyze the TCP-RED operating point and its stability with respect to various RED controller and system parameters. In a subsequent study [23], a stochastic model of a bottleneck RED gateway under a large number of heterogeneous TCP flows has been proposed and the asymptotic behavior of the system has been investigated. More recently, the interaction between additive-increase multiplicative-decrease (AIMD) congestion control and droptail AQM schemes is investigated through models utilizing nonnegative matrix theory [24].

Although much of control-oriented analysis and design has appeared within linear settings, there is significant evidence in literature for complex and chaotic queue behavior in Internet-type networks and their models [3], [22], [25], [26]. Non-equilibrium fluctuations in queue behavior has been observed both experimentally and numerically [25], [27]. For appropriate parameter values, the simplified fluid models also exhibit persistent non-equilibrium behavior [3], [22], [28]. The non-equilibrium behavior in queues may be due to random noise or could arise as self-excited “chaotic oscillations” and there are suggestions for both in the literature [27]. Papers concerned

⁰An earlier version of this paper was presented at the 7th IFAC Symposium on Nonlinear Control Systems (NOLCOS 2007, Pretoria, South Africa, August 22-24, 2007).

¹Deutsche Telekom Laboratories, Ernst-Reuter-Platz 7, 10587 Germany.

²Dept. of Mechanical Science & Engineering and the Coordinated Science Laboratory, University of Illinois, 1206 W. Green Street, Urbana IL 61801.

³Coordinated Science Laboratory, University of Illinois, 1308 West Main Street, Urbana IL 61801.

with deriving mean or limit models [23], [29], [30], obtaining AQM performance with linear control methods [11], [31], or with carrying out describing function based analysis [32] typically assume a random noise. The papers concerned with local instability/bifurcation analysis [22], [33]–[36], and (global) numerical investigations [1], [3], [22] using deterministic fluid models show these oscillations as self-excited. With simple fluid models, analytical methods from bifurcation theory have been used to show that these self-excited oscillations can arise as a result of supercritical Hopf bifurcation [33], [35], [37] and of period doubling and border-collision bifurcations [22].

Even though the methods of this paper are relevant to both noise-driven and self-excited cases, we are primarily motivated by the analysis and control issues for the latter case, i.e., where the non-equilibrium queue behavior arises as a result of nonlinear dynamics and not random noise. In such a case, there is a gap between available methods that focus on static optimization, and simulations that show persistent non-equilibrium behavior that does not need any noise. From a practical viewpoint, explicit analysis and control of non-equilibrium behavior could play an important role in the performance of the overall congestion control scheme.

In the remainder of this section, we discuss the main elements of the proposed approach. We represent, for modeling and control, the dynamic variables by their stochastic counterparts. Even though the models are deterministic, the analysis and control approach is stochastic. This enables us to go beyond the frameworks based on a single equilibrium-point. The modeling approach is based upon the methods of Ergodic theory for representing complex behavior in nonlinear dynamical systems. In particular, we replace the dynamical models by their stochastic counterparts - the so-called Perron-Frobenius (P-F) operator [38]–[40]. While, the dynamical model propagates the initial condition, the Perron-Frobenius operator propagates uncertainty in initial condition. The main advantage is that it is generally easier to represent the complex asymptotic dynamic behavior as invariant probability measures of the P-F operator. In the context of this paper, we do this to represent and model the queue behavior.

For computational purposes, we use set-oriented numerical methods for discretization of the dynamical systems; cf. [41], [42]. Our goal is to use these simulation based methods to construct finite-dimensional Markov chains from the dynamic model. These Markov chains are then used to carry out, numerically, dynamic analysis and control design. A somewhat different Markov modeling of communication networks has additionally been considered in [43], [44]. In [43], a Hidden Markov Model for a communication channel has been studied where the channel switches between different states. Each state corresponds to the probability that a packet sent over the channel will be lost. In [44], a Markovian Model for the RED algorithm has been proposed where the states are composed of the (average) queue size and some flags. Using this model, the impact of RED on the mean delay and loss rate has been analyzed. The basic idea of our work – use of stochastic approaches for analysis and control of non-equilibrium dynamic behavior in deterministic network settings – is very different in nature and novel.

The stochastic modeling approach we take enables us to carry out:

Bifurcation analysis: Although the methods of local bifurcation theory can be used to analyze the instability and (local) onset of bifurcation as in [33], [35], these methods are less relevant to the AQM problem because of the global nature of the queue oscillation and presence of discontinuities such as saturations in these models. For instance, the chaotic oscillations have been typically studied using numerical simulations as in [3], [22]. Stochastic methods, on the other hand, enable analysis of the asymptotic chaotic dynamics in terms of invariant measures [38]. Bifurcations in chaotic regime can be understood via analysis of the spectrum of the P-F operator or the support of the invariant measure [45]. In particular, we use these results to understand qualitative changes in queue behavior.

Identification of the fluid model parameters from packet level *ns2* [46] simulations: In the non-equilibrium regime, the results of both the fluid-model as well as the *ns-2* simulation display rich time-series behavior. As a result, the question of how well fluid models approximate the reality or even *ns-2* simulations is not straight-forward [47]. We validate the fluid-approximation based results with the ones of *ns-2* simulations by comparing the two invariant measures using the L^1 -norm as metric. We use this comparison to justify the choice of the parameters for the fluid model.

Control synthesis for shaping non-equilibrium behavior: We propose a Markov Decision Process (MDP) based framework for optimization of the asymptotic dynamics. Even though the control framework considered is quite general, we use without any loss of generality an AQM control structure similar to RED for the purposes of this paper. We pose and solve the control problem as a MDP where the full state – user and queue behavior – is observed, and as control of Hidden Markov Models (HMMs) where only the queue size is observed. Both the analysis and control are verified using simulations in MATLAB.

The outline for the rest of the paper is as follows: In Section II, a well-known network model of user and queue behavior along with an equilibrium and stability analysis of the Droptail scheme is presented. In Section III, a stochastic modeling of the network together with its discrete approximation as finite-dimensional Markov models is described and used to carry out bifurcation analysis as well as identification of network model parameters using the *ns-2* simulations. In Section V, an MDP-based framework for optimization and control of these models is summarized. The Matlab and *ns-2* simulation results demonstrating control of non-equilibrium behavior under the AQM algorithm developed are described in Section VI. The paper ends with the concluding remarks of Section VII.

II. DETERMINISTIC FLUID MODEL

A. Single Bottleneck Link with Symmetric Users

We consider a single bottleneck link of a network with fixed capacity C shared by M users. Instead of conducting a packet level analysis of the network, we adopt a network model

based on fluid approximations [6], [9]. Each user is associated with a unique connection for simplicity and transmits with a nonnegative flow rate x_i over this bottleneck link. For $x_i \in \mathbb{R}^+ \doteq [0, \infty)$, the i^{th} user is assumed to follow a transfer control protocol (TCP)-like additive-increase multiplicative-decrease flow control scheme,

$$\dot{x}_i(t) = \kappa \left(\frac{1}{d} - \beta x_i(t)^2 p(t) \right), \quad (1)$$

where $0 \leq p \leq 1$ is the observed rate of marking (or depending on the implementation, dropping) of its packets, κ denotes the step-size, and d and β denote the (symmetric) rate-increase and decrease parameters, respectively. For a prescribed $p(t)$, the ODE (1) has a well-defined solution in \mathbb{R}^+ for all time because the right hand side is Lipschitz in x_i and \mathbb{R}^+ is a positively invariant set with respect to (1). This is because $\dot{x}_i(t) = \frac{\kappa}{d} > 0$ at the boundary $x_i = 0$. We use the underline notation $\underline{x} \doteq (x_1, \dots, x_M)$ to denote the vector of user flow rates, where $\underline{x} \in \mathbb{R}^{+M}$.

The packet marking occurs at the link whose dynamics are next described. If the aggregate sending rate of users exceeds the capacity C of the link, then the arriving packets are queued in the buffer q of the link. The non-negative queue size evolves according to the ODE

$$\dot{q}(t) = \begin{cases} \sum_{i=1}^M x_i(t) - C & q \in (0, B), \\ \min(0, \sum_{i=1}^M x_i(t) - C) & q = B, \\ \max(0, \sum_{i=1}^M x_i(t) - C) & q = 0. \end{cases} \quad (2)$$

where we assume a maximum buffer size of B at which the queue saturates and any incoming packet after this point is dropped; cf. [11]. $p(\cdot)$ in (1) is set by the AQM control and takes the general form $p = F(q)$. As an example, packet marking for the widely used droptail AQM scheme is described by

$$p = \begin{cases} 0 & , \text{ if } q < B \\ 1 & , \text{ otherwise.} \end{cases} \quad (3)$$

It is the objective of this paper to discuss questions pertaining to 1) (non-equilibrium) dynamic analysis for a given AQM scheme F and 2) control synthesis of the optimal F .

With a large number of users M , a detailed non-equilibrium analysis of the multi-user model (1) is infeasible. In order to simplify the analysis, we note that the equations are equivariant with respect to the permutation group with the group action $x_i \rightarrow x_j$ for $i, j \in \{1, \dots, M\}$. As a result of this symmetry, the linear subspace

$$S = \{ \underline{x} \in \mathbb{R}^{+M} : x_i = x_1 \} \quad (4)$$

is a fixed-point space; cf., [48]. We will refer to S as the synchrony subspace. In particular, the subspace S is positively invariant with respect to dynamics of (1). In the following section, we show that the subspace S is also stable with respect to arbitrary initial conditions of user flow rate, i.e., the user flow rates synchronize after a period of transients. As a result, we will analyze the non-equilibrium dynamics of (1)-(2) for only the symmetric fixed-point space where all the users have the same flow-rate.

B. Synchronization to symmetric fixed-point space

Theorem II.1. Consider the multi-user setup of (1)-(2), where $p = F(q)$ and F satisfies the condition that $F(B) > 0$ then the synchrony subspace S is asymptotically stable, i.e., as $t \rightarrow \infty$, $x_i(t) = x_1(t)$ for all $i = 1, \dots, M$.

Proof. To show asymptotic stability, we use the Lasalle's invariance theorem [49]. The steps in the proof are 1) we propose a Lyapunov function $V(\underline{x})$, 2) set $E \doteq \{ \underline{x} \in \Omega \subset \mathbb{R}^{+M} | \dot{V}(t) = 0 \}$, where Ω is a compact positively invariant with respect to (1), and 3) show that the largest invariant set in E , denoted by M , lies in $S \cap \Omega$. We outline the three steps below:

- 1) The Lyapunov function is taken to be

$$V(\underline{x}) = \frac{1}{2} \sum_{i=2}^M (x_i - x_1)^2. \quad (5)$$

Using (1), the time derivative of (5) is

$$\dot{V}(\underline{x}) = -\kappa\beta \sum_{i=2}^M (x_i - x_1)^2 \cdot (x_i + x_1) \cdot p. \quad (6)$$

Since $p \in [0, 1]$, $\dot{V}(t) \leq 0$.

- 2) Set

$$\Omega = \left\{ \underline{x} \in \mathbb{R}^{+M} : x_i \leq \max \left(C + B, \frac{1}{\sqrt{\beta d}} \right) \right\}. \quad (7)$$

To see that Ω is positively invariant, note that $\dot{x}_i \leq 0$ whenever $x_i > \max \left(C + B, \frac{1}{\sqrt{\beta d}} \right)$. This shows that the trajectories are bounded within set Ω and thus there exists a (largest such) compact invariant set M that contains all the limit points.

- 3) Finally, we show that $M \subset S \cap \Omega$. Using (6), first note that if $x \in E$ then either (a) $x_i = x_1$, or (b) $x_1 = x_i = 0$ for all $i = 2, M$, or (c) $p = 0$. We consider case (c) first. Suppose $p \equiv 0$ over a trajectory then using (1),

$$x_i(t) = x_i(0) + \frac{\kappa}{d} t, \quad (8)$$

and there exists a finite time $T < \frac{d(C+1)}{\kappa M} + B$ such that $q(T) = B$ and $p = F(B) > 0$. As a result, any set of points with $p \equiv 0$ is not an invariant set and the case (a) or (b) applies. In both these cases, $\underline{x} \in S$. Since S is a fixed-point space, it contains its invariant set and $M \subset S \cap \Omega$ as desired. \square

We denote the symmetric user flow rate as x . In the fixed-point space, the system dynamics are

$$\begin{aligned} \dot{x}(t) &= \kappa \left(\frac{1}{d} - \beta x(t)^2 p(t) \right), \\ \dot{q}(t) &= \begin{cases} Mx(t) - C & q \in (0, B), \\ \min(0, Mx(t) - C) & q = B, \\ \max(0, Mx(t) - C) & q = 0. \end{cases} \end{aligned} \quad (9)$$

The number of the users M is now a system parameter and we will investigate bifurcations in dynamic behavior with respect to this and other parameters. We note that the important effect of delay has been ignored in this model.

C. Equilibrium and Stability Analysis of Droptail

In this section, we carry out a preliminary stability and bifurcation analysis of the multi user fluid model (1)-(2) with a droptail AQM (3). The analysis is analytically tractable because the equilibrium dynamics arise entirely in the fixed-point space. The result is summarized with the aid of the following Theorem:

Theorem II.2. Consider the model (1)-(2) with a droptail AQM (3). Using $\beta^* \doteq \frac{1}{d} \left(\frac{C}{M}\right)^{-2}$ to represent the critical value of β , we have the following conclusion regarding the equilibrium solution:

- 1) For values of $\beta < \beta^*$, there exists a unique equilibrium solution with user flow rates $x_i = \bar{x} \doteq \frac{1}{\sqrt{d\beta}}$ and saturated value of queue $q = \bar{q} \doteq B$. This equilibrium solution, denoted as \bar{x} , lies in the subspace S and is stable.
- 2) For values of $\beta > \beta^*$, no equilibrium solution can exist. The asymptotic dynamics are non-equilibrium but lie in the fixed-point space S .

Proof. For an equilibrium solution to exist, there can only be two possibilities: either $p \equiv 0$ or $p \equiv 1$. The first possibility is trivially ruled out because in the absence of feedback from the router, the user flow rate x_i will increase without bound. The latter case is more interesting. Suppose $p \equiv 1$. Then, the only possible equilibrium of (1) – if it exists – is given by

$$x_i = \bar{x} = 1/\sqrt{d\beta} \quad (10)$$

and is symmetric. Substituting this solution in (9) leads to two possible cases: **(a)** if $\bar{x} > \frac{C}{M}$ then q increases and saturates to Q_{\max} . According to droptail, this then leads to $p = 1$ validating our assumption. This equilibrium point is valid, and provided it is stable, can be observed within a simulation. **(b)** If $\bar{x} < \frac{C}{M}$ then q decreases and empties to 0. With droptail, this then leads to $p = 0$ thereby invalidating our assumption and resulting in no equilibrium.

Now, the condition $\bar{x} > \frac{C}{M}$ is equivalent to $\beta < \beta^*$. The argument above thus shows that an equilibrium solution cannot exist for the range of values $\beta > \beta^*$. To complete the proof, we note that the equilibrium $\bar{x} > \frac{C}{M}$ is locally stable due to the stability of its linearization:

$$\begin{aligned} \dot{\delta x} &= -2\kappa\sqrt{\frac{\beta}{d}}\delta x, \\ \dot{\delta q} &= 0. \end{aligned} \quad (11)$$

□

This theorem shows that for sufficiently small feedback gain β , the queue is full (regardless of B) and the packets are always being dropped. Such an equilibrium, even though it is stable, is clearly not desirable. As the feedback gain β is increased, one reaches the critical value

$$\beta^* = \frac{1}{d} \left(\frac{C}{M}\right)^{-2} \quad (12)$$

beyond which the assumption $p \equiv 1$ is violated. An equilibrium solution cannot exist for the range of values $\beta > \beta^*$.

These results are visualized in Figure 1: either an equilibrium solution does not exist or when it exists, the user flow rate is greater than capacity and the queue is always full.

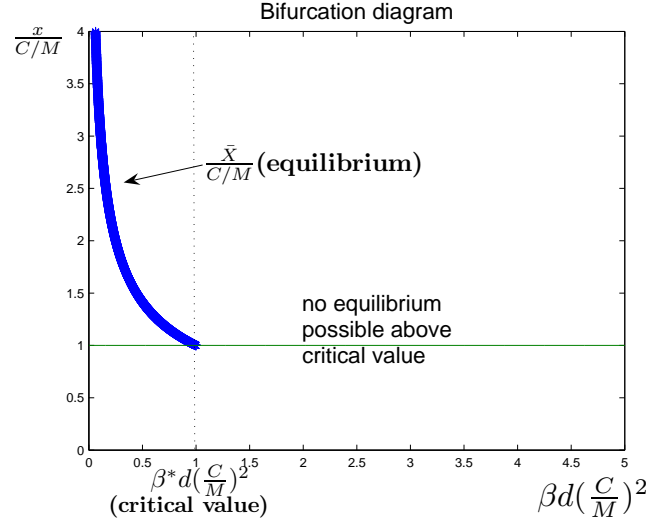


Fig. 1. Bifurcation diagram for the droptail AQM.

For the system (9) under the droptail AQM scheme the queue begins “oscillating” about its upper-limit B for values of β greater than the critical value. Furthermore, even the small oscillations at the onset are not periodic. Figure 2(a) depicts an typical time-series of this system and the Figure 2(b) shows the largest incursion of these oscillations as a function of the parameter β . As observed in Figure 2(a), the queue behavior

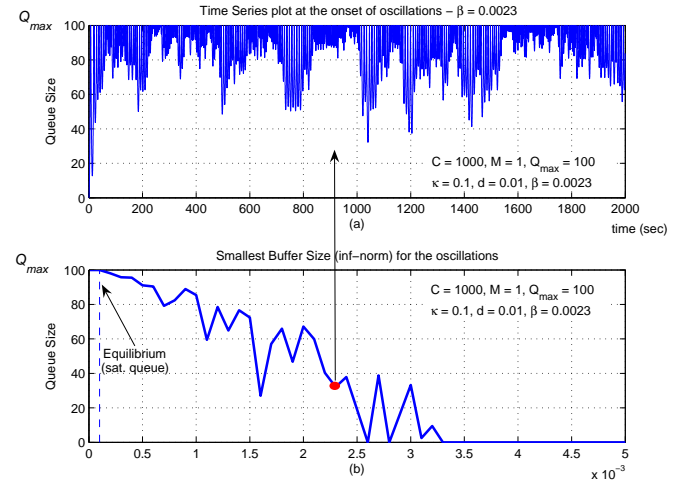


Fig. 2. (a) Time-series of incipient oscillations for $\beta = \beta^*+$ and (b) a numerically determined bifurcation diagram of the non-equilibrium queue behavior.

is both complex (non-periodic) and global in the phase space. The non-periodic nature of oscillations arise because of the discontinuity in p . However, even with a somewhat smoother version of the function $p(\cdot)$, the local methods of bifurcation theory are perhaps not best-suited for global analysis with large queue oscillations.

TABLE I
MODEL PARAMETERS

Link capacity	$C = 1000$
# of users	$M = 10, M^* = 25$
Queue bounds	$Q_{min} = 0, Q_{max} = 100$
User parameters	$\kappa = 0.05, d = 0.01, \beta = 0.0625$

The results obtained from the fluid model also shed light to more realistic simulations such as the ones in ns-2. In the TCP, the “parameter β ” is implicit and fixed. Using (12), the critical behavior in ns-2 simulations arises as a function of either C or M . In particular, with a fixed C there exists a critical value of M , denoted by M^* , such that the queue saturates for $M > M^*$ and oscillates for $M < M^*$.

D. Parameter Identification

For the fluid-approximation (9), the increase-rate parameter $\frac{\kappa}{d}$ and the decrease-rate parameter $\kappa \cdot \beta$ are identified to match the ns-2 simulations. Specifically, $\frac{\kappa}{d}$ is identified from the average slope of the TCP additive increase phase. In order to identify the parameter $\kappa\beta$, multiple ns-2 simulations for different number of users M were carried out. From these simulations, M^* was determined as the critical number of users for which the queue starts to oscillate. The queue is saturated for $M > M^*$ and oscillates for $M < M^*$. Using (12) the parameter

$$\kappa\beta = \frac{\kappa}{d} \left(\frac{C}{M^*} \right)^{-2}. \quad (13)$$

Figure 3 depicts the probability distributions (histogram) obtained from the ns2 time-series data (9). Table II-D depicts the identified nominal model parameters for the fluid approximation (9). For analysis and control synthesis, we will use $M = 10$ number of users. Note that this corresponds to a non-equilibrium queue behavior with droptail AQM.

Before presenting the analysis, we make a few remarks regarding approximations implicit in this identification. With ns2 simulations, the time-series data shows several features which are not captured by fluid approximation models. Vis-a-vis parameter identification, two approximations had to be made. The first approximation was in the determination of the critical M^* . ns2 simulations show a large range of values of M with small oscillations about the saturated queue level. To accommodate this, M^* was identified to be the number of users at the onset of “larger” oscillations; see Fig. 3. The second and more significant approximation was in ignoring the delay. Both stochastic and spectral analysis of the time-series data from ns2 simulations indicate the important effect of delay. Although, the effect of delay was found to be important in establishing a reasonable match to ns2 simulations, it will not be considered in this paper.

III. STOCHASTIC MODEL AND BIFURCATION ANALYSIS

For analysis and control design, we employ stochastic (Markovian) representations of non-equilibrium dynamics. These representations aid the analysis of asymptotic aspects

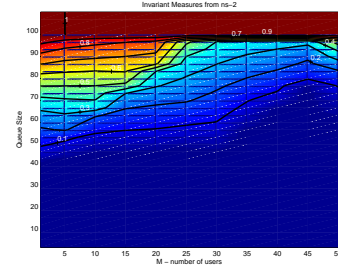


Fig. 3. Bifurcation diagram in terms of the invariant measure: ns-2 dynamics as the parameter M is increased with the droptail AQM.

of these dynamics in terms of invariant measures and their numerical approximations. We begin by considering the dynamical system $T : X \times Q \rightarrow \mathbb{R}^{+2}$ obtained by sampling the solutions of the ODE. In particular, denote $\phi(t; x_0, q_0)$ to be the solution operator for the fluid approximation ODE (9) and set $T(x_0, q_0) = \phi(\Delta t; x_0, q_0)$. The sampling time Δt is taken to be of the same order as the delay and $(x_0, q_0) \in X \times Q$ denote the initial condition. The AQM policy will be tested on the sampled data system with sampling time of ΔT . This policy is consistent with the optimization problem that will be posed and solved in the following section. In effect, this reflects the nature of AQM where a time-period of order round trip time is used for the AQM policy update. $X \subset \mathbb{R}^+$ denotes the compact state-space for $x(\cdot)$, $Q = [0, B] \subset \mathbb{R}^+$ denotes the compact state-space for $q(\cdot)$ and $S \doteq X \times Q$.

In stochastic settings, the basic object of interest is the Perron-Frobenius (P-F) operator \mathbb{P} corresponding to the dynamical system T . It is given by $\mathbb{P}[\nu](A) = \nu(T^{-1}(A))$, where $A \subset \mathcal{B}(S)$, the Borel σ -algebra of S and $\nu \in \mathcal{M}(S)$, the measure space on S . $T^{-1}(A)$ denotes the pre-image set, i.e., $T^{-1}(A) \doteq \{x \in S : T(x) \in A\}$. While the dynamical system T describes the nonlinear evolution of an initial condition, the P-F operator \mathbb{P} describes the linear evolution of the uncertainty (probability density function) in initial conditions. The advantage of using a stochastic framework is that asymptotic dynamics of T can be interpreted as invariant measures of the stochastic operator \mathbb{P} . The invariant measure is a probability measure that is also a fixed-point of the P-F operator \mathbb{P} , i.e., $\mathbb{P}\nu_1 = 1 \cdot \nu_1$. From Ergodic theory, an invariant measure is always known to exist under the assumption that the mapping $T : S \rightarrow S$ is at least continuous and S is compact; cf., [38].

The set-oriented numerical methods have recently been employed for constructing efficient finite-dimensional approximations of the P-F operator; cf. [41], [42]. The approximation arises as a Markov matrix defined with respect to a finite partition $\mathcal{S}_L \doteq \{D_1, \dots, D_L\}$ of the phase space S . Instead of a Borel σ -algebra $\mathcal{B}(S)$, consider now a σ -algebra of the all possible subsets of \mathcal{S}_L . A real-valued measure ν_j is defined by ascribing to each element D_j a real number. Thus, one identifies the associated measure space with a finite-dimensional real vector space \mathbb{R}^L . Using Galerkin approximations, the discrete P-F approximation arises as a matrix

$$P_{ij} = \frac{m(T^{-1}(D_j) \cap D_i)}{m(D_i)} \quad (14)$$

on the “measure space” \mathbb{R}^L ; m is the Lebesgue measure [42], [47]. The resulting matrix is non-negative and if $T : D_i \rightarrow S$, $\sum_{j=1}^L P_{ij} = 1$, i.e., P is a Markov or a row-stochastic matrix. P is interpreted as an approximation of \mathbb{P} obtained by considering a certain random perturbation of the dynamical system T . P converges to \mathbb{P} in L^2 as the partition gets finer and finer [39].

The partition \mathcal{S}_L for the stochastic approximation of the network model is constructed by taking a quantization for the user flow-rates (in X) and queue size (in Q) between a lower and upper bound. The lower bounds are taken to be 0 because of the non-negativity of these quantities. The upper bounds are taken to be suitable multiples of link capacity and maximum queue buffer size. On account of computational constraints, we chose $O(10)$ quantization levels for the user flow rate and the bottleneck link queue size. The two quantized partitions are denoted as $\mathcal{X} = [X_1, \dots, X_{22}]$ and $\mathcal{Q} = [Q_1, \dots, Q_{22}]$ and the partition size $L = 484$. We denote $\mathcal{S}_L \doteq \mathcal{X} \times \mathcal{Q}$, where the states $s \in \mathcal{S}_L$ are indexed as $\{(X_1, Q_1), (X_1, Q_2), \dots, (X_2, Q_1), (X_2, Q_2), \dots\}$. Table II tabulates the quantization values used for constructing the cells in \mathcal{X} and \mathcal{Q} . The sub-script L is dropped for the remainder of the paper to simplify the notation.

TABLE II
QUANTIZATION OF STATES (\mathcal{X} AND \mathcal{Q})

X_1	$0 - 0.05C/M$	Q_1	$0 - 5$
X_2	$0.05C/M - 0.1C/M$	Q_2	$5 - 10$
\vdots	\vdots	\vdots	\vdots
X_{20}	$0.95C/M - C/M$	Q_{20}	$95 - 100$
X_{21}	$C/M - 1.05C/M$	Q_{21}	$100 - 105$
X_{22}	$1.05C/M - 1.1C/M$	Q_{22}	$105 - 110$

For the network dynamical system (9), a Markov Model (MM) consists of a Markov chain with states in \mathcal{S} and transition probabilities (entries of P in Eq. (14)) between these states. The entry P_{ij} denotes the transition probability of the next state being in D_j conditioned on the current state being in D_i . The state evolution associated with the nonlinear dynamical system T (Eq. (9)) is replaced its stochastic approximation,

$$\nu(n+1) = \nu(n)P, \quad (15)$$

where $\nu(\cdot) \in \mathbb{R}^L$ is the row probability vector.

One approach for numerically approximating P in Eq. (14) is to use a Monte Carlo algorithm with several short term simulations using the dynamical system T . Here, N uniformly distributed random samples $n_i^{in} = (x, q)_i$ $i = 1, \dots, N$ in \mathcal{S} are used as initial conditions for the dynamical system T in Eq. (9). Denote $n_i^{out} = (x, q)_i$ $i = 1, \dots, N$, as the image of these points after one iterate of the dynamical system. After identifying the input and output samples n_i^{in} and n_i^{out} with the states $s \in \mathcal{S}$ of the MM, the transition probability from state i to j is estimated as

$$\hat{P}_{ij} = \frac{\sum_{[k:n_k^{out} \in j]} 1}{\sum_{[l:n_l^{in} \in i]} 1},$$

where $\sum_{[k:n_k^{out} \in j]}$ denotes the number of points k such that $n_k^{out} \in j$. The algorithms for constructing these approximations and their numerical convergence properties appear in [50]. As one takes finer partitions, the invariant measure

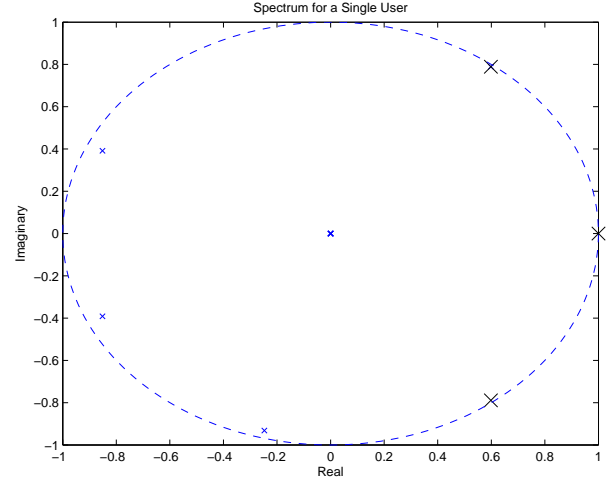


Fig. 4. Eigenvalues of the Markov chain for 10 users ($M = 10$).

of Markov matrix P converges to a weak limit ν^* that approximates the invariant measure of the P-F operator \mathbb{P} ; see Theorem 3.1 in [42]. In typical situations, the support of the invariant measure is the attractor set. Thus, the stochastic MM provides a description of the original dynamical system’s asymptotic behavior.

Figure 4 depicts the spectrum of the Markov chain P . There is a unique eigenvalue at 1 and the corresponding eigenvector gives the invariant measure, denoted as ν_1 . Figure 5 compares this invariant measure with the ergodic averages computed using time-domain simulation; cf., [47]. As shown in the figure, the MM is fairly accurate in describing the asymptotic behavior of the system in terms of probabilities.

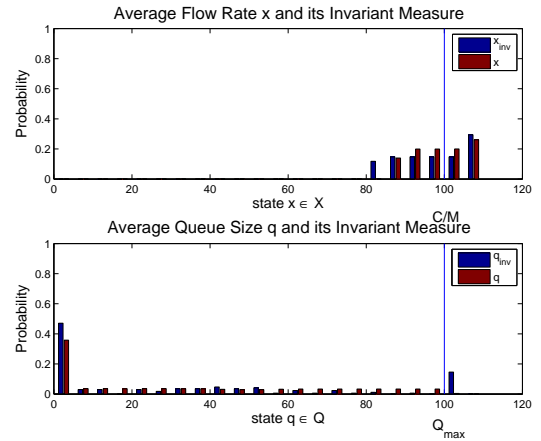


Fig. 5. The time-averaged dynamics and invariant measures from the MM for (top) the user flow rate x and (bottom) the queue size q ; 1-10 are the 10 quantization bins \mathcal{X} and \mathcal{Q}

The presence of complex spectrum close to the unit circle

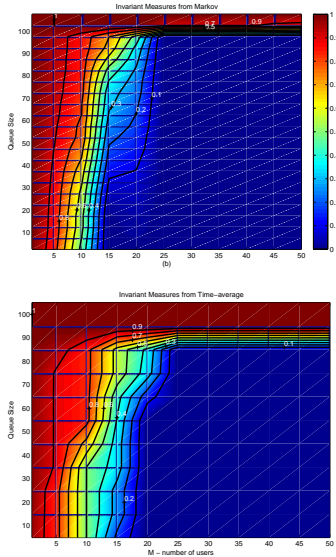


Fig. 6. Bifurcation diagram in terms of the invariant measure: (left) MM dynamics, (right) the time-averaged dynamics as the parameter M is increased with the droptail AQM.

also suggests oscillations in the asymptotic dynamics; cf., [39], [40], [47]. This is indeed consistent with the results of the time-domain simulation which exhibit oscillations in the non-equilibrium regime. A comparison between the frequency obtained using spectrum of Markov chain and the Discrete Fourier transform (DFT) of time-series data is summarized as part of the bifurcation analysis.

Having obtained the Markov model, we carried out a bifurcation analysis with the number of users M serving as a parameter. At the critical value M^* , the support of the invariant measure $\bar{\mu}$ changes in a qualitative sense. Figure 6 compares the invariant measure $\bar{\mu}$ for the parameterized MM with the ergodic averages computed using time-domain simulations. The time-domain results with both the ns-2 simulator and the fluid approximation are given. For quantitative comparison, Figure 7 depicts the L^1 distance between the invariant measure computed using the MM and the ergodic average with time-series data. The L^1 norm between two distributions p_{ts} (with time-series) and p_{mm} (with Markov model) is given in Fig. 7. It is defined as

$$\|p_{ts} - p_{mm}\|_1 := \frac{1}{2} \sum_i |p_{ts}(i) - p_{mm}(i)|. \quad (16)$$

The L^1 -error is small indicating that the MM is fairly accurate in capturing the asymptotic behavior with the time-domain simulation over a range of values of M . Finally, Figure 8 compares the frequencies of the oscillations obtained from Markov model and the DFT of the time-series data. We refer the reader to our paper [47] for details on the spectral analysis for chaotic systems.

IV. CONTROL MARKOV CHAINS AND HMMS

Given the user and queue dynamical system in (9), we define the AQM control structure using ten separate packet marking (or dropping) schemes p_u , $u = 1, \dots, 10$. \mathcal{U} denotes

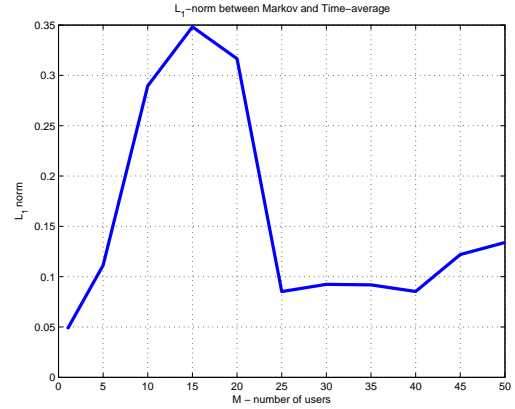


Fig. 7. L^1 -norm between time-averages and MM

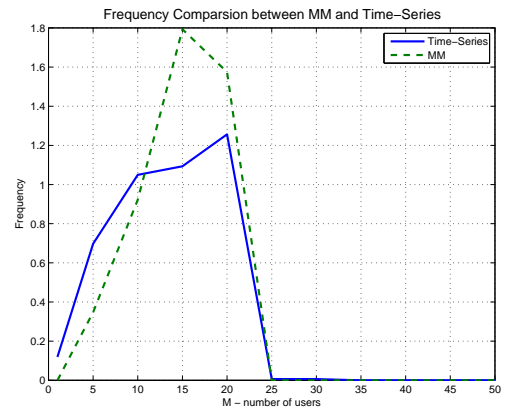


Fig. 8. Frequency of oscillation of time-averages and MM

the set with these finite number of control schemes, i.e., $\mathcal{U} = \{p_u\}_{u=1}^{10}$. These schemes, implemented at the bottleneck link, set the value of $p(\cdot) = p_u$ in the user equation. The AQM scheme p_1 corresponds to well-known ‘‘droptail’’ behavior whereas others can be interpreted as variants of RED, where the packet marking probability is constant instead of linearly increasing.

The queued packets are not marked (corresponds to $p_u = 0$), if the queue size is less than a certain lower threshold Q_{min} . The packets are always marked if the queue size is larger than a upper threshold Q_{max} (corresponds to $p_u = 1$). Q_{max} is set to the buffer size B . If the current queue size is between the two thresholds ($Q_{min} < q < Q_{max}$), the packets are marked according to $p_u = (u - 1) \times 0.1$ for $u = 1, \dots, 10$. Table III summarizes the packet marking schemes, where $Q_{min} = 0$ and $Q_{max} = 100$. The number and properties of are chosen for simplicity and illustrative purposes. Our analysis can be extended to a more complex control structure in a straightforward manner.

The AQM control structure modifies the dynamical system (9) and leads to control Markov chains P^u , where the superscript u corresponds to the choice of the (fixed) control scheme p_u . The element P_{ij}^u denotes the probability of the next

TABLE III
CONTROL POLICIES

Policy	$q \leq Q_{min}$	$Q_{min} < q < Q_{max}$	$q \geq Q_{max}$
p_1	0	0	1
$p_u (u = 2 - 10)$	0	$(u - 1) \times 0.1$	1

state being in D_j conditioned on the current state being in D_i and control being p_u . The control Markov chain corresponds to the approximation of the Perron-Frobenius operator of the control dynamical system and as such is a straightforward extension of the discussion in Section II.

It is not realistic to assume knowledge of both the user flow rates and the queue sizes for control design. At the bottleneck link, one would typically know only the queue size and not the the user flow rates. Therefore, we consider a hidden Markov model (HMM) for describing the dynamical system's behavior in the presence of partial observations. For the network model, *only* the set of queue states \mathcal{Q} is assumed to be observed. The *emission matrix* E maps the set of states \mathcal{S} of the MM to the set of queue states \mathcal{Q} and has the structure

$$E := \begin{bmatrix} 1 & 0 & \cdots & 0 & 1 & 0 & \cdots \\ 0 & 1 & & & 0 & 1 & \cdots \\ \vdots & & \ddots & \vdots & \vdots & & \ddots \\ 0 & 0 & \cdots & 1 & 0 & 0 & \cdots \end{bmatrix}_{\mathcal{Q} \times \mathcal{X}},$$

where E_{ij} denotes the probability of the current queue in cell Q_i conditioned on the state in cell D_j .

Remark IV.1. Note that all three of the AQM schemes in Section IV do not mark (or drop) packets if the queue size is less than Q_{min} .

V. OPTIMIZATION, CONTROL AND ESTIMATION

Now that we have a Markov model on a finite state space \mathcal{S} with finitely many control actions \mathcal{U} , we pose the control problem as a Markov Decision Process (MDP). In particular, given a state $s \in \mathcal{S}$, we would like to determine the optimal policy $s \rightarrow \pi^*(s)$ such that a certain expected reward is maximized over a time horizon. The policy refers to any rule for choosing control. For optimization purposes, the reward per stage $g(s)$ is defined over the states in \mathcal{S} and is depicted in Figure 9. It is chosen to be the largest for a queue size between 10 and 20. The expectation is that a positive but moderate queue size will ensure maximum capacity utilization while preventing large queue fluctuations and resultant buffer overflows. For large values of queue, the states with large values of user flow rate are penalized more than those with smaller values. For very small queue size, lower values of user flow rate are penalized.

Assuming reward per stage $g(s)$ and a policy π , the infinite horizon discounted reward is given by

$$J_\pi(s) := E_\pi \left[\sum_{n=0}^{\infty} \alpha^n g(s(n)) \mid s(0) = s \right],$$

where E_π represents the conditional expectation given that the policy π is employed, and $\alpha < 1$ is the discount factor. By choosing the discount factor lower, the resulting policies can

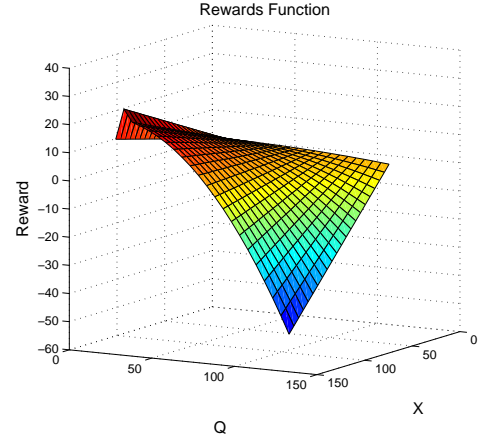


Fig. 9. The reward function on $\mathcal{X} \times \mathcal{Q}$.

be made more myopic (i.e. heavily discounting future rewards) which can be useful when the system at hand is not fully stationary. The optimal reward function is defined by

$$J_\alpha(s) := \max_{\pi} J_\pi(s), \quad s \in \mathcal{S}.$$

A policy $\pi^*(s)$ is optimal if $J_{\pi^*}(s) = J_\alpha(s)$ for all states s . From the theory of MDP (see Chapter 7 in [51]), the optimal reward function satisfies the Bellman operator equation

$$J_\alpha(s) = \max_u \left\{ g(s) + \alpha \sum_j P_{ij}^u J_\alpha(j) \right\}, \quad (17)$$

and the optimal policy $\pi^* : \mathcal{S} \rightarrow \mathcal{U}$ is a stationary policy, i.e., the control at time n depends only upon the state at time n .

For finite state space, as in our case, there are numerous approaches for obtaining the optimal stationary policy. In the paper, we choose the *value iteration* algorithm where the solution to (17) is obtained using the recursion

$$J_{k+1}(s) = \max_u \left\{ g(s) + \alpha \sum_j P_{ij}^u J_k(j) \right\}. \quad (18)$$

This recursion leads to the optimal stationary policy π^* .

Using value of $\alpha = 0.3$, the reward function $g(s)$ given in Fig. 9 and the choice of discrete control schemes in Table III, we used (18) to synthesize the optimal stationary AQM policy solution π^* on $\mathcal{X} \times \mathcal{Q}$. Figure 10 depicts this policy referred to as NEQM (Non-Equilibrium Queue Management) in the remainder of this paper. We note that the NEQM policy, while being more general, is not very different from the existing AQM schemes such as RED in terms of dependence of $p(\cdot)$ on queue size. In particular, the policy drops (or marks) packets aggressively for larger queue sizes. As such, it provides an optimization-based explanation of why RED is a reasonable AQM scheme in practice.

Using the NEQM policy μ^* , we also computed the invariant probability measures of the controlled Markov matrix P^{μ^*} for user flow rates (in \mathcal{X}) and queue size (in \mathcal{Q}). These measures,

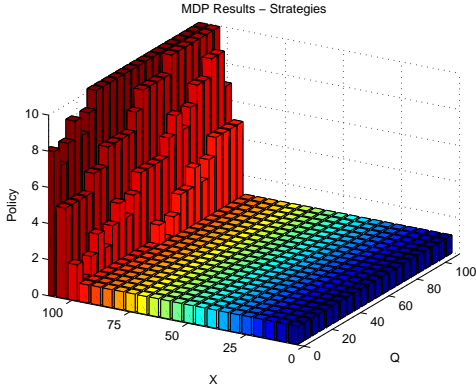


Fig. 10. The optimal policy μ^* on $\mathcal{X} \times \mathcal{Q}$ obtained by solving the MDP.

shown in Figure 11, suggest the possibility of non-equilibrium queue dynamics. This was indeed verified using numerical simulations as discussed in the following section.

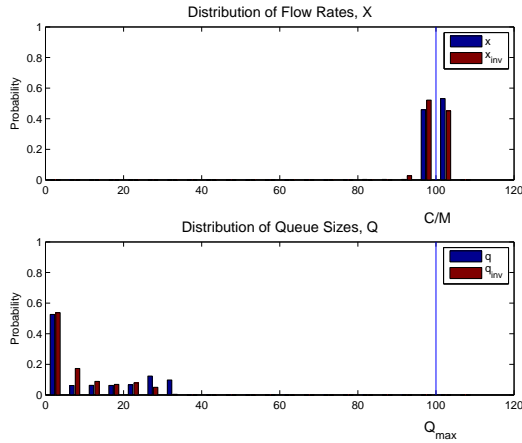


Fig. 11. The asymptotic stationary probabilities of the states in (top) \mathcal{X} (flow rate) and (bottom) \mathcal{Q} (queue size) under optimal AQM control obtained from MM.

Before describing the numerical simulations, we briefly discuss the implementation of NEQM policy for the practical situation where the user flow rate is hidden. For such a case, we applied the Viterbi algorithm [52] to estimate the state (in \mathcal{S}) from the queue time-series. In particular, given the HMM consisting of E, P^u ($u = 1, \dots, 10$) and given a sequence of m observations at the bottleneck link $o := [o_1, \dots, o_m]$, $o_i \in \mathcal{Q}$, we estimate the single best state sequence (path) $\hat{s} = [\hat{s}_1, \dots, \hat{s}_m]$ of the real state sequence $s = [s_1, \dots, s_m]$, $s \in \mathcal{S}$. The chosen estimation criterion was to maximize $P(\hat{s} | o, P^u, E)$ given the observations o . Figure 12 depicts the typical estimation results obtained with a window based implementation of the Viterbi algorithm.

VI. SIMULATIONS

A. Numerical Analysis in Matlab

We carried out Matlab based simulations with multi-user fluid approximation models and parameter values consistent

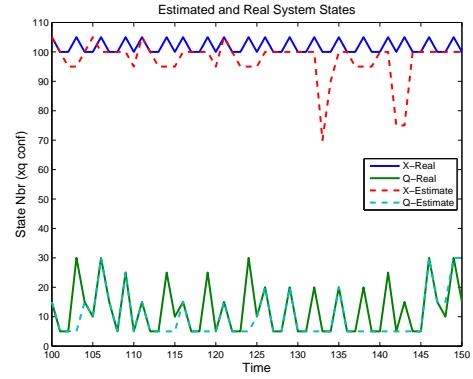


Fig. 12. Estimated \hat{s} and actual system states $s \in \mathcal{S} = \mathcal{X} \times \mathcal{Q}$ under AQM control.

with the preceding analysis. Apart from the symmetric user case, we also simulated a case with asymmetric users in order to test the robustness of the AQM scheme. Table IV summarizes the simulation parameters for both the symmetric and the asymmetric user cases.

TABLE IV
SIMULATION MODEL PARAMETERS

Link capacity	$C = 1000$
# of users	$M = 10$
Queue bounds	$q_{min} = 10, q_{max} = 100$
Symmetric user parameters	$\kappa = 0.05, d = 0.01, \beta = 0.0625$
Asymmetric user parameters	$d_i \in [0.008, 0.015], \beta_i \in [0.08, 0.15]$

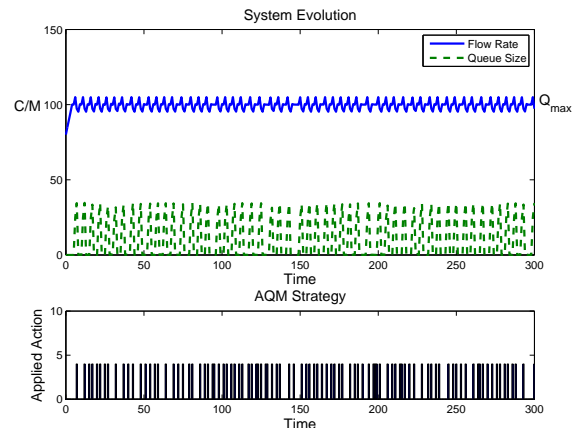


Fig. 13. Results of the simulation with a state-dependent NEQM policy: (top) The evolutions of user flow rate x , link queue size q , and (bottom) the AQM scheme deployed versus time are shown.

We begin by describing the results of simulations carried out using full state feedback. Using the NEQM policy described in Section V, the individual user flow rates synchronize after a short period of transients. Figure 13 depicts the non-equilibrium evolution of the controlled system – x (averaged over 10 users) and the queue q – as a function of time.

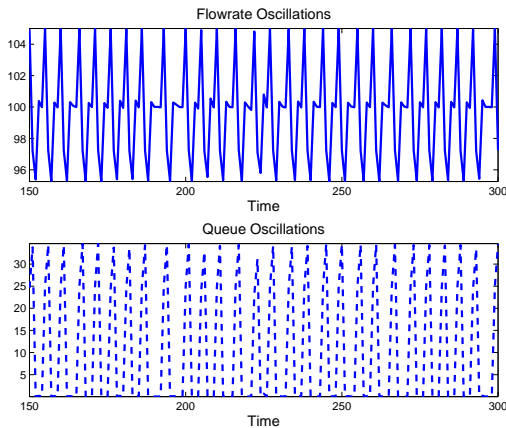


Fig. 14. Results of the simulation with a state-dependent NEQM policy: (top) The evolutions of user flow rate x and (bottom) link queue size q are shown.

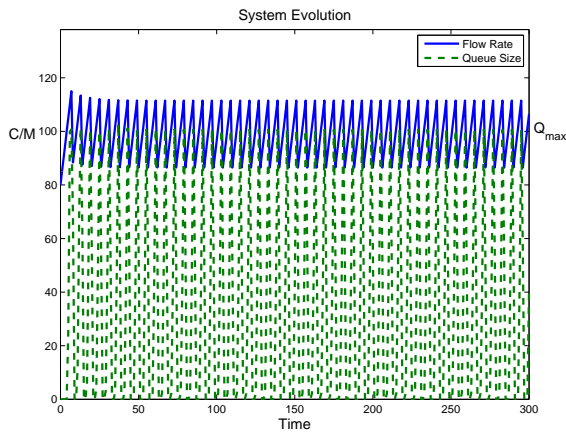


Fig. 15. Results of the simulation with the drop-tail AQM scheme: the evolutions of user flow rate x , link queue size q , the state s of the MM versus time are shown.

Note that on account of synchronization, x also represents the individual user flow rates. The AQM accomplishes near capacity utilization while maintaining a small average queue size. Averages of the non-equilibrium solution are consistent with the choice of reward $g(s)$ used in MDP (see Fig. 9). Figure 13 also depicts the sequence of control actions, i.e., the specific AQM scheme p_u implemented at any given time instance. Note that consistent with the MDP based optimization, the stationary AQM policy was applied using an update of control action every Δt seconds together with a ZOH. For better illustrating the non-equilibrium aspects of closed-loop dynamics, Figure 11 depicts the invariant measures of the asymptotic dynamics and Figure 14 depicts the individual time-series for x and q but now at a greater resolution.

In order to compare these results against the standard AQM, we next simulated the multi-user system with the AQM scheme 1 (droptail). The results for droptail are shown in Figure 15. It is evident from the two figures that while either of the two solutions are non-equilibrium, the MDP based solution is clearly better because it shows a larger user

capacity utilization, smaller queue fluctuations, and averages close to the requirement with respect to the reward function. In summary, the MDP based solution uses the state information to better anticipate the congestion and adjust the packet marking accordingly. That it is able to do so with non-equilibrium queue behavior that still achieves very close to maximum capacity utilization (see Figure 13) is notable.

Next, we describe the results for the more practical case where only the instantaneous queue size q was assumed to be known. For estimation purposes, we used the Viterbi algorithm as discussed in Section V. We formally assumed certainty equivalence to hold and treated the estimate as the state. The NEQM was thus applied according to the estimates. Figure 16 depicts the results of this observer based control. The estimation errors had a noticeable but small effect on the performance of the algorithm. Analysis of this will be the subject of future investigations.

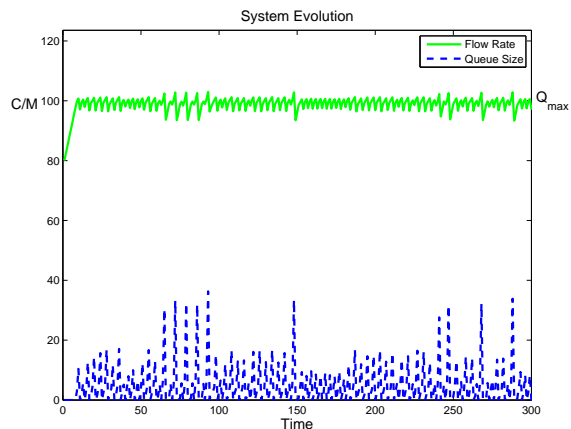


Fig. 16. Results of the simulation with a state estimation based feedback control: (top) the evolutions of user flow rate x and link queue size q ; (bottom) the real s vs. estimated \hat{s} state of the HMM and the AQM scheme deployed vs. time are shown.

Finally, we discuss the results of simulations carried out with asymmetric users; see Table IV for the simulation parameters. For the asymmetric case, the parameters d_i and β_i for i^{th} -user were picked from a uniform distribution whose range is indicated in the table.

We again used the optimal NEQM policy together with an estimation based on the reduced order symmetric Markov model. With Markov model, the symmetric user flow rate x was *formally* replaced by the spatially averaged flow rate for the M asymmetric users ($\frac{1}{M} \sum_{i=1}^M x_i$). As a result, there are modeling errors introduced because of the reduction in dimension (M states to a single state) and spatial averaging. Note that the time-domain simulations were carried out with the $M + 1$ -dimensional asymmetric multi-user dynamical system (1) and (2). Figure 17 depicts the results of this simulation. Remarkably, even though the dynamic behavior now is more complex (see for example, the queue trajectory), the performance shows fair average capacity utilization (for x) and queue size q .




Fig. 17. Results of the simulation with a state estimation based feedback control for the asymmetric multi-user case: (top) the evolutions of average flow rate x and link queue size q ; (bottom) the real s and estimated state \hat{s} of the HMM and the AQM scheme deployed versus time are shown.

B. Simulations in Ns2

We now simulate a multi-user, single bottleneck link scenario implemented in the network simulator ns-2. We send 25 TCP flows over this duplex link of capacity 8 Mbps corresponding to 100 packets/second for an average packet size of 1000 bytes. The maximum queue size is chosen to be $Q_{max} = 100$ packets and the propagation delay to be 10 ms in the forward direction.

The NEQM is implemented in ns-2 similar to the other AQM schemes by extending the C++ queue class. To determine the state s at a given time the current queue size and flow rate passing through the link are measured periodically. Notice that, for real life deployments the current queue size is readily available (to a router). Furthermore, a network measurement tool such as *abing* [53] can be run continuously in the background or variations in the queue size in conjunction with the link capacity can be utilized to estimate the aggregate flow rate. Once the state is known, the packets are dropped according to the drop probabilities determined by the specific AQM scheme dictated by NEQM. For example, if the policy p_2 is active at a time instance then incoming packets are not dropped if $q < q_{min}$, they are dropped with probability 0.3 if $q_{min} < q < q_{max}$, and are always dropped if $q > q_{max}$. For these simulations a similar set of parameters are used as the ones in Table IV. However, the number of AQM policies (Table III) are reduced from ten to four in order to simplify the implementation complexity in the simulation environment.

Figure 18 (a) and (b) compares the queue and the user behavior obtained with RED, and the NEQM. We observe that ns-2 results verify the numerical analysis of the reduced order model using Matlab. There is better queue utilization with the NEQM policy in terms of size and variation as compared to the droptail as well as RED. We note that, similar to RED, the mean flow rate shows consistent capacity utilization for NEQM flows. On the other hand, we observe in Figure 19 that there is a strong contrast in the mean user flow rate between NEQM and the droptail scheme, which exhibits heavy

oscillations.




Fig. 18. Results of the ns-2 simulations (for the parameters in Table II) with (a) RED and (b) the NEQM schemes.




Fig. 19. Results of the ns-2 simulations (for the parameters in Table II) with (a) Droptail and (b) the NEQM schemes.

We observe a correspondence between the results of ns-2 simulations and reduced order model analysis. However, there are also quantitative differences such as the frequency content of the queue oscillations. In particular, numerical analysis of the reduced order model shows high frequency dynamics. Such a difference was also observed during the parameter identification and discussed therein. Compared to the results of the reduced order model, the user flow rates with ns-2 show more stochasticity. We conjecture this to be the effect of noise in ns-2 simulations resulting from TCP windowing effects and quantization of the flows via packets. Finally, when NEQM and RED are compared we observe that their behavior is similar. This is not surprising as the four individual schemes used by NEQM are quite close to RED.

In order to investigate more realistic network settings, we also study the effect of disturbances due to varying background traffic and other factors. Disturbance analysis was carried out by injecting a constant bit rate (CBR) flow, which randomly decreases the link capacity. This flow is randomly injected at intervals of 0.1 seconds as a uniform disturbance of between 0 and 250 packets/second corresponding up to 25% of the link capacity. Figure 20 depicts the user flow-rates and the queue behavior in the presence of disturbance. We observe a visible degradation in RED while the NEQM scheme performs almost as good as before, which illustrates the robustness property of the NEQM. This result is not very surprising when taking into account the fact that NEQM makes use of both the flow rate and queue size during its operation.

VII. CONCLUSION

In this paper, we have presented an ergodic theoretic framework for stochastic modeling, model computations, analysis, identification, and MDP based control of nonlinear dynamics




Fig. 20. Results of the ns-2 simulations in the presence of disturbance with (a) RED, and (b) the NEQM scheme.

in communication networks. The framework was demonstrated using both ns-2 simulations and simplified fluid models for the AQM problem. Using equilibrium and local bifurcation analysis as well as numerical simulations with simplified fluid models, we illustrated the fact that methods for analysis and optimization of non-equilibrium queue behavior are *essential* for a solution to this problem. Next, we presented an identification methodology for identifying the parameters of the nonlinear dynamical system model which explicitly take the non-equilibrium behavior into account. Using stochastic approximations of this model, an optimization problem was defined and solved via standard MDP methods to obtain a specific optimal stationary AQM policy, NEQM, as an example. The NEQM scheme developed provides first such explanation of AQM schemes such as RED and was shown here to perform better than the droptail scheme using ns-2 simulations.

We view the analysis and optimization framework presented here as the first step to better understand the AQM schemes mentioned in the introduction. As we have discussed in Section II, a study of the non-equilibrium dynamics is crucial for the analysis of non-equilibrium behavior we observe in networked systems. Likewise, the optimization framework outlined above and the NEQM scheme as its illustrative example is a first step towards understanding the role of non-equilibrium control in optimizing queue behavior and system performance. In order to do this rigorously, we will need to understand the role of various complexity mitigation approaches including

- 1) symmetry reduction,
- 2) rigorous spatial averaging for the multi-user case,
- 3) numerical methods used for approximations of stochastic operators, and
- 4) selection of a finite number of policies for tractable computation of optimal control.

This is a focus of continuing investigations.

ACKNOWLEDGMENT

The authors would like to thank Uday Shanbhag and Umesh Vaidya for helpful discussions.

REFERENCES

[1] S. Shakkottai, R. Srikant, and S. P. Meyn, "Bounds on the throughput of congestion controllers in the presence of feedback delay," *IEEE/ACM Transactions on Networking*, vol. 11, no. 6, pp. 972–981, December 2003.

[2] H. Han, C. Hollot, D. Towsley, and Y. Chait, "Synchronization of tcp flows in networks with small droptail buffers," in *Proc. of the 44th IEEE Conference on Decision and Control*, Seville, Spain, December 2005, pp. 6762–67.

[3] A. Veres and M. Boda, "The chaotic nature of TCP congestion control," in *Proc. IEEE Infocom*, vol. 3, March 2000, pp. 1715–1723.

[4] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: Shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.

[5] F. P. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, vol. 8, pp. 33–37, January 1997.

[6] T. Alpcan and T. Başar, "A utility-based congestion control scheme for Internet-style networks with delay," *IEEE Transactions on Networking*, vol. 13, no. 6, pp. 1261–1274, December 2005.

[7] T. Alpcan, T. Başar, and R. Tempo, "Randomized algorithms for stability and robustness analysis of high-speed communication networks," *IEEE Transactions on Neural Networks*, vol. 16, no. 5, pp. 1229–1241, September 2005.

[8] T. Alpcan and T. Başar, "Global stability analysis of an end-to-end congestion control scheme for general topology networks with delay," in *Proc. of the 42nd IEEE Conference on Decision and Control*, Maui, HI, December 2003, pp. 1092 – 1097.

[9] R. Srikant, *The Mathematics of Internet Congestion Control*, ser. Systems & Control: Foundations & Applications. Boston, MA: Birkhauser, 2004.

[10] S. Deb and R. Srikant, "Global stability of congestion controllers for the Internet," *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 1055–1060, June 2003.

[11] C. V. Hollot, V. Misra, D. Towsley, and W. Gong, "Analysis and design of controllers for AQM routers supporting TCP flows," *IEEE Transactions on Automatic Control*, pp. 945–959, June 2002.

[12] G. Vinnicombe, "On the stability of networks operating TCP-like congestion control," in *Proc. 15th IFAC World Congress on Automatic Control*, Barcelona, Spain, July 2002.

[13] S. H. Low, F. Paganini, J. Wang, and J. C. Doyle, "Linear stability of TCP/RED and a scalable control," *Computer Networks Journal*, vol. 43, no. 5, pp. 633–647, December 2003.

[14] R. Johari and D. Tan, "End-to-end congestion control for the Internet: Delays and stability," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 818–832, December 2001.

[15] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397–413, August 1993.

[16] S. Kunniyur and R. Srikant, "Analysis and design of an adaptive virtual queue (AVQ) algorithm for active queue management," in *Proc. of ACM SIGCOMM*, August 2001, pp. 123–134.

[17] S. Athuraliya, V. H. Li, S. H. Low, and Q. Yin, "Rem: Active queue management," *IEEE Network*, vol. 15, no. 3, pp. 48–53, May/June 2001.

[18] F. Wu-chang, K. Shin, D. Kandlur, and D. Saha, "The BLUE active queue management algorithms," *IEEE/ACM Transactions on Networking*, vol. 10, no. 4, pp. 513–528, August 2002.

[19] S. Liu, T. Başar, and R. Srikant, "Exponential-RED: a stabilizing AQM scheme for low- and high-speed TCP protocols," *IEEE/ACM Transactions on Networking*, vol. 13, no. 5, pp. 1068–1081, October 2005.

[20] S. Kunniyur and R. Srikant, "End-to-end congestion control schemes: Utility functions, random losses and ECN marks," in *Proc. of the IEEE Infocom*, 2000, pp. 1323–1332. [Online]. Available: citeseer.nj.nec.com/358188.html

[21] F. Baccelli, D. R. McDonald, and J. Reynier, "A mean-field model for multiple TCP connections through a buffer implementing RED," *Performance Evaluation*, vol. 49, no. 1–4, pp. 77–97, 2002.

[22] P. Ranjan, E. Abed, and R. J. La, "Nonlinear instabilities in TCP-RED," *IEEE/ACM Transactions on Networking*, vol. 12, no. 6, pp. 1079–1092, December 2004.

[23] P. Tinnakornsrisuphap and R. J. La, "Asymptotic behavior of heterogeneous TCP flows and RED gateway," *IEEE/ACM Trans. on Networking*, vol. 14, no. 1, pp. 108–120, February 2006.

[24] R. Shorten, C. King, F. Wirth, and D. Leith, "Modelling TCP congestion control dynamics in drop-tail environments," *Automatica*, vol. 43, no. 3, pp. 441–449, March 2007.

[25] J. Chung and M. Claypool, "Analysis of active queue management," in *Proc. of 2nd IEEE International Symposium on Network Computing and Applications (NCA)*, Cambridge, Massachusetts, April 2003.

[26] V. Firoiu and M. Borden, "A study of active queue management for congestion control," in *Proc. IEEE Infocom*, Tel Aviv, Israel, March 2000.

- [27] D. Wischik, *Queueing theory for TCP*, University of Illinois at Urbana-Champaign, 2006, 7th Intl. Conf. on Stochastic Networks.
- [28] T. Alpcan, P. G. Mehta, T. Başar, and U. Vaidya, "Control of non-equilibrium dynamics in communication networks," in *Proc. of Conference of Decision and Control*, San Diego, CA, 2006. [Online]. Available: <http://decision.csl.uiuc.edu/alpcan>
- [29] S. Deb and R. Srikant, "Rate-based versus queue-based models of congestion control," *IEEE Transactions on Automatic Control*, vol. 51, no. 4, pp. 606–619, April 2006.
- [30] F. Baccelli, D. R. McDonald, and J. Reynier, "A mean-field model for multiple TCP connections through a buffer implementing RED," *Performance Evaluation*, vol. 49, pp. 77–97, 2002.
- [31] K. B. Kim, "Design of feedback controls supporting TCP based on the state-space approach," *IEEE Transactions on Automatic Control*, vol. 51, no. 7, pp. 1086–99, July 2006.
- [32] E. Korkut, "Limit cycling in TCP Networks," Master's thesis, University of Massachusetts, Amherst, MA, 2006.
- [33] G. Raina, "Local bifurcation analysis of some dual congestion control algorithms," *IEEE Transactions on Automatic Control*, vol. 50, no. 8, pp. 1135–1146, August 2005.
- [34] G. Raina and D. Wischik, "Buffer sizes for large multiplexers: TCP queueing theory and instability analysis," in *EuroNGI Conf. on Next Generation Internet Networks*, 2005.
- [35] C. Li, G. Chen, X. Liao, and Y. Juebang, "Hopf bifurcation in an internet congestion control model," *Chaos, Solitons Fractals*, vol. 19, pp. 853–862, 2004.
- [36] X. F. Wang, "Controlling bifurcation and chaos in internet congestion control system," in *Procs. of the 4th World Congress on Intelligent Control and Automation*, Shanghai, China, June 2002, pp. 573–576.
- [37] H. Yin, P. Wang, T. Alpcan, and P. G. Mehta, "Hopf bifurcation and oscillation in communication networks with delays," in *Proc. of 17th IFAC World Congress*, 2008.
- [38] A. Lasota and M. C. Mackey, *Chaos, Fractals, and Noise: Stochastic Aspects of Dynamics*. New York: Springer-Verlag, 1994.
- [39] M. Dellnitz and O. Junge, "On the approximation of complicated dynamical behavior," *SIAM Journal on Numerical Analysis*, vol. 36, pp. 491–515, 1999.
- [40] I. Mezic and A. Banaszuk, "Comparison of systems with complex behavior," *Physica D*, vol. 197, pp. 101–133, 2004.
- [41] M. Dellnitz and O. Junge, *Set oriented numerical methods for dynamical systems*. World Scientific, 2002, pp. 221–264.
- [42] G. Froyland, *Nonlinear Dynamics and Statistics: Proceedings, Newton Institute, Cambridge, 1998*. Birkhauser, 2001, ch. Extracting dynamical behaviour via Markov models, pp. 283–324.
- [43] K. Salamatian and S. Vaton, "Hidden Markov modeling for network communication channels," *SIGMETRICS Performance Evaluation Review*, vol. 29, no. 1, pp. 92–101, 2001.
- [44] R. Laalaoua, T. Czachorski, and T. Atmaca, "Markovian model of RED mechanism," in *Proc. of the First IEEE/ACM Internat. Symp. on Cluster Computing and the Grid*, May 2001, pp. 610–617.
- [45] P. G. Mehta, M. Hessel, and M. Dellnitz, "Symmetry of attractors and the Perron-Frobenius operator," *Journal of Difference Equations & Applications*, September 2006.
- [46] UCB, LBNL, and VINT, "Network simulator ns (version 2)," <http://www.isi.edu/nsnam/ns/>.
- [47] P. G. Mehta and U. Vaidya, "On stochastic analysis approaches for comparing dynamical systems," in *Proc. of the 44th IEEE Conference on Decision and Control*, Seville, Spain, December 2005, pp. 8082–87.
- [48] M. Golubitsky and I. Stewart, *The Symmetry Perspective. From Equilibrium to Chaos in Phase Space and Physical Space*. Basel: Birkhäuser, 2002.
- [49] H. K. Khalil, *Nonlinear Systems*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, 1996.
- [50] M. Dellnitz, G. Froyland, and O. Junge, "The algorithms behind GAIO – set oriented numerical methods for dynamical systems," in *Ergodic Theory, Analysis, and Efficient Simulation of Dynamical Systems*, B. Fiedler, Ed. Springer, 2001, pp. 145–174.
- [51] S. M. Ross, *Applied Probability with Optimization Applications*. San Francisco, CA: holden-Day, Inc., 1970.
- [52] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," in *Proc. of the IEEE*, vol. 77, no. 2, February 1989, pp. 257–286.
- [53] J. Navratil and R. L. Cottrell. Abing. [Online]. Available: <http://www-iepm.slac.stanford.edu/tools/abing/>