

Control of Non-Equilibrium Dynamics in Communication Networks

Tansu Alpcan, Prashant G. Mehta, Tamer Başar, and Umesh Vaidya

Abstract—We consider the analysis and control problem of Active Queue Management (AQM) in communication networks. Specifically, we are interested in the complex and non-equilibrium dynamic queue behavior arising in nonlinear AQM models. We propose a dynamical systems based approach for stochastic modeling of the non-equilibrium behavior in these deterministic models. The asymptotic dynamics are interpreted using invariant measures of certain stochastic operators. For computational tractability, we consider set-oriented numerical methods to construct finite-dimensional Markov models including control Markov chains and hidden Markov models. We pose and solve the AQM control problem using a Markov Decision Process (MDP) based framework and present results of numerous simulations that show persistent non-equilibrium queue behavior with an optimal control strategy.

I. INTRODUCTION

Communication networks such as the Internet exhibit a wide variety of non-equilibrium and complex 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. Kelly's framework for network capacity optimization [4], [5] together with control-oriented stability analysis of these solutions in the context of congestion control problems has had an important impact [6]. The game theoretic approach to network control and optimization has been another avenue of research [7]–[9]. However, much of this research has focussed on a single-point equilibrium solution and analysis of its stability properties.

This paper is concerned with the analysis and control of non-equilibrium behavior in communication networks. Specifically, we consider the Active Queue Management (AQM) problem. AQM provides a mechanism by which a link (router) sends 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) and its variations are other well-known examples of AQM algorithms [6].

There exists significant evidence in literature for complex and chaotic queue behavior in Internet-type networks and their models [3], [10]–[12]. From a practical viewpoint, the control of such behavior could play an important role in the performance of the overall congestion control scheme. As a result, a number of AQM schemes with different characteristics has been proposed and studied by the research community [13]–[16].

In this paper, 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. 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 operator [17]–[19]. While, the dynamical model propagates the initial condition, the Perron-Frobenius (P-F) operator propagates uncertainty in initial condition. The upshot is that it is frequently 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.

Next, we use the recent set-oriented numerical methods for discretization of the dynamical systems; cf. [20], [21]. The idea is to use these simulation based methods to construct finite-dimensional Markov chains from the dynamic model. These Markov chains are used for both the computational analysis as well as for control design. A somewhat different Markov modeling of communication networks has additionally been considered in [22], [23]. However, application of the stochastic modeling approach for control of non-equilibrium behavior in AQM appears to be novel.

For control, we propose a Markov Decision Process (MDP) based framework for optimization of the asymptotic dynamics. Even though, the control framework is quite general, we use 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 (HMM) where only the queue size is observed. Both the analysis and control are verified using simulations in MATLAB.

The outline of this paper is as follows. In Section II, the network model of user and queue behavior is presented. In the following three Sections III–V, stochastic modeling of the network model together with its discrete approximation as

Research of T. Alpcan and T. Başar was supported in part by the NSF ITR Grant CCR 00-85917.

T. Alpcan is with Deutsche Telekom AG Laboratories, Technische Universität Berlin, Ernst-Reuter-Platz 7, 10587 Berlin, Germany. alpcan@control.csl.uiuc.edu

P. G. Mehta is with the Department of Mechanical & Industrial Engineering, University of Illinois at Urbana-Champaign, 1206 W. Green Street, Urbana, IL 61801 mehtapg@uiuc.edu

T. Başar is with the Coordinated Science Laboratory, University of Illinois, 1308 West Main Street, Urbana, IL 61801 USA. tbasar@control.csl.uiuc.edu

U. Vaidya is with the Department of Electrical & Computer Engineering, Iowa State University, Ames, IA 50014 ugvaidya@iastate.edu

finite-dimensional Markov models – control Markov chains and hidden Markov model – is presented. In Sec. VI, an MDP based framework for estimation and control of these models is summarized. The simulation results demonstrating control of non-equilibrium behavior in AQM are described in Sec. VII, and conclusions appear in Sec. VIII.

II. THE NETWORK MODEL

We consider a network consisting of a single bottleneck link of fixed capacity C shared by M users. Instead of conducting a packet level analysis of the network we use a discrete-time network model based on fluid approximations [6], [7]. Each user is associated with a unique connection for simplicity and transmits with a nonnegative flow rate x_i over this bottleneck link. The i^{th} user is assumed to follow a TCP-like additive-increase multiplicative-decrease flow control scheme responding to the observed rate $0 \leq p \leq 1$ of marking (or depending on the implementation, dropping) of its packets,

$$x_i(n+1) = x_i(n) + \kappa \left[\frac{1}{d_i} - \beta_i x_i(n)^2 p(n) \right]^+, \quad (1)$$

where κ denotes the step-size, and d_i and β_i denote the user-specific rate increase and decrease parameters, respectively. The function $[x]^+ = x$ for nonnegative values of x and is zero otherwise. Here, n denotes the discrete time instance normalized with respect to the round trip time (RTT) of users. However, the effect of information (feedback) delay is not taken into account.

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

$$q(n+1) = q(n) + \left[\sum_{i=1}^M x_i(n) - C \right]^+. \quad (2)$$

For simulations, we sometimes consider an ideal buffer of infinite size in order to compare various queue management schemes. Finally, $p(\cdot)$ in Eq. (1) is set by the AQM control and will be discussed in Sec. IV.

III. STOCHASTIC MODELING OF NETWORK DYNAMICS

We study the deterministic network model's dynamics using a stochastic framework. For stochastic modeling, it is first assumed that all the users are symmetric and have the same parameters, $d_i = d_j$ and $\beta_i = \beta_j \forall i, j = 1, \dots, M$, as well as same initial conditions. Consequently, the user flow rates are also symmetric $x = x_i \ i = 1, \dots, M$. The system dynamics in the symmetric fixed-point space are

$$\begin{aligned} x(n+1) &= x(n) + \kappa \left[\frac{1}{d} - \beta x(n)^2 p(n) \right]^+ \\ q(n+1) &= q(n) + [Mx(n) - C]^+ \end{aligned} \quad (3)$$

We denote these equations as a dynamical system $T : X \times Q \rightarrow \mathbb{R}^2$, where $X \subset \mathbb{R}^1$ is the compact state-space for $x(\cdot)$ and $Q \subset \mathbb{R}^1$ is the compact state-space for $q(\cdot)$. We define $S \doteq X \times Q$.

Remark III.1. It is possible to approximately treat the asymmetric multi-user case just like the symmetric one by *formally* replacing the symmetric flow rate x in (3) with the average flow rate of asymmetric users $(1/M) \sum_{i=1}^M x_i$. Admittedly, such a study of the asymmetric case is only an approximation. Sec. VII describes some simulation results with the asymmetric multi-user model.

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

$$\mathbb{P}[\mu](A) = \mu(T^{-1}(A)), \quad (4)$$

where $A \subset \mathcal{B}(S)$, the Borel σ -algebra of S and $\mu \in \mathcal{M}(S)$, the measure space on S . 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.,

$$\mathbb{P}\mu_1 = 1 \cdot \mu_1. \quad (5)$$

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., [17].

The set-oriented numerical methods have recently been employed for constructing efficient finite-dimensional approximations of the P-F operator; cf. [20], [21]. 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 μ_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

$$P_{ij} = \frac{m(T^{-1}(D_j) \cap D_i)}{m(D_i)} \quad (6)$$

on the “measure space” \mathbb{R}^L ; m is the Lebesgue measure [21], [24]. The resulting matrix is non-negative and if $T : D_i \rightarrow S$,

$$\sum_{j=1}^L P_{ij} = 1, \quad (7)$$

i.e., P is a Markov or a row-stochastic matrix. P is interpreted as a randomly perturbed approximation of \mathbb{P} and P converges to \mathbb{P} in L^2 as the partition gets finer and finer [18].

The partition \mathcal{S}_L for the stochastic approximation of the network model is constructed by taking a uniform 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

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

X_1	0-15	Q_1	0-10
X_2	15-30	Q_2	10-20
\vdots	\vdots	\vdots	\vdots
X_9	120-135	Q_9	80-90
X_{10}	150-	Q_{10}	90-

of computational constraints, we chose 10 quantization levels for the user flow rate and the bottleneck link queue size. The two quantized partitions are denoted as $\mathcal{X} = [X_1, \dots, X_{10}]$ and $\mathcal{Q} = [Q_1, \dots, Q_{10}]$ and the partition size $L = 100$. 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 I 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.

For the network dynamical system (3), a Markov Model (MM) consists of a Markov chain with states in \mathcal{S} and transition probabilities (entries of P in Eq. (6)) 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. (3)) is replaced its stochastic approximation,

$$\mu(n+1) = \mu(n)P, \quad (8)$$

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

One approach for numerically approximating P in Eq. (6) is to use a Monte Carlo algorithm. 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. (3). 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 $\sum_{[k:n_k^{out} \in j]} 1$ denotes the number of points k such that $n_k^{out} \in j$.

As one takes finer partitions, the invariant measure of Markov matrix P converges to a weak limit μ^* that approximates the invariant measure of the P-F operator \mathbb{P} . The invariant measure of \mathbb{P} is a stochastic counterpart of the asymptotic dynamics (attractor set) of the original dynamical system (3); see Theorem 3.1 in [21]. 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 1 compares the invariant measure μ^* for the MM corresponding to (3) with the time-averaged (asymptotic) dynamics of its simulation; cf., [24]. As shown in the figure, the MM is fairly accurate in describing the asymptotic behavior of the system.

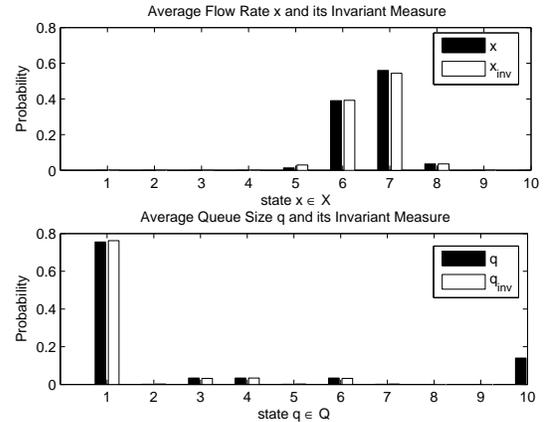


Fig. 1. The time-averaged dynamics and invariant measures from the MM for (top) the user flow rate x and (bottom) the queue size q ; 1-10 are the 10 quantization bins \mathcal{X} and \mathcal{Q} (see Table I)

Remark III.2. We use the simple network model in Eq. (3) as an example to illustrate the proposed approach. Real networks exhibit much more chaotic and complex dynamics than these simplified models. Thus, aspects of the stochastic approach are perhaps even more suitable to the real networks.

IV. CONTROL MARKOV CHAINS

Given the user and queue dynamical system in (3), we define the AQM control structure using three separate packet marking (or dropping) schemes p_1 , p_2 , and p_3 . 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 scheme p_2 and p_3 can be interpreted as a simple variant of RED and a very aggressive marking algorithm, respectively. In all of the schemes, 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 an upper threshold q_{max} (corresponds to $p_u = 1$). If the current queue size is between the two thresholds ($q_{min} < q < q_{max}$), the packets are marked randomly in scheme 2 ($p_2 = 0.5$), and all the packets are marked in scheme 3 ($p_3 = 1$).

The AQM control modifies the dynamical system (3) and leads to control Markov chains P^u , where $u \in \{1, 2, 3\}$ corresponds to the choice of control scheme p_u . The P_{ij}^u denotes the probability of the next state being in D_j conditioned on the current state being in D_i and control being 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 III.

Finally, the number and properties of these active queue management (AQM) schemes are chosen for simplicity and illustrative purposes. Our analysis can be extended to a more complex control structure in a straightforward manner.

V. HIDDEN MARKOV MODEL

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

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

Remark V.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} . This behavior is needed in order to ensure the observability of the system and states in the HMM. Otherwise, given the observed queue states \mathcal{Q} (through emission matrix E) it would not be possible to estimate the full state \mathcal{S} .

VI. OPTIMIZATION, CONTROL, AND ESTIMATION

Now that we have a Markov model on a finite state space \mathcal{S} with finitely many control actions $u = \{1, 2, 3\}$, 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 $s \rightarrow \hat{u}(s)$ such that a certain expected reward

$$\max_{\hat{u}(s(0))} E \left[\sum_k \alpha^k R(s(k)) \right] \quad (9)$$

is maximized over a time horizon. The expectation is taken with respect to the Markov model and α denotes the discount factor. It is chosen as 0.3 for the results in this paper. The reward function $R(s)$ is defined over the states in \mathcal{S} and is depicted in Figure 2. 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 state 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.

The solution of the optimization problem in (9) is obtained by solving a linear program

$$\begin{aligned} \min \quad & \sum_s z_s \\ \text{subject to} \quad & z_s \geq R'(s) + \alpha \sum_s P_{sj}^u z_s \quad \forall s, u \\ & z_s \geq 0 \quad \forall s. \end{aligned} \quad (10)$$

Here z_s for $s = 1, \dots, L$ are auxillary variables, and $R' = R + c$, where c is chosen such that R' is non-negative; cf.,

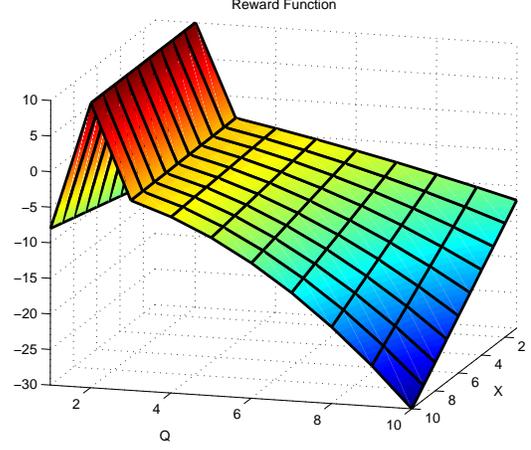


Fig. 2. The reward function on $\mathcal{X} \times \mathcal{Q}$.

Chapter 6 in [25]. Figure 3 depicts the resulting stationary optimal policy solution \hat{u} . We note that the stationary policy,

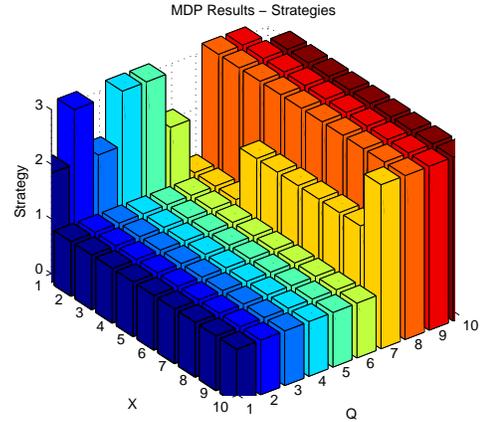


Fig. 3. The optimal strategy (actions) obtained as a result of solving the MDP.

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.

Using the optimal policy \hat{u} , we also computed the invariant probability measures of the controlled Markov matrix $P^{\hat{u}}$ for user flow rates (in \mathcal{X}) and queue size (in \mathcal{Q}). These measures were found to compare favorably with the asymptotic distributions obtained from the time-domain simulations of the closed-loop dynamical system (3) with the optimal policy. The latter are shown in Figure 4. We note that either of these results point to non-equilibrium queue behavior with the optimal policy. The details of the simulation results will be discussed in the following section.

To test the optimal policy for the practical situation where the user flow rate is hidden, we applied the Viterbi algorithm [26] to estimate the state (in \mathcal{S}) from the time-series

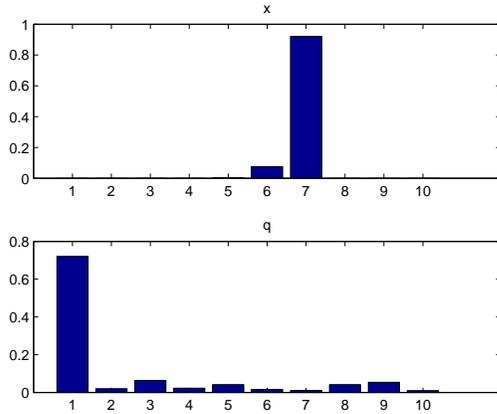


Fig. 4. The asymptotic stationary probabilities of the states in (top) \mathcal{X} (flow rate) and (bottom) \mathcal{Q} (queue size) under optimal AQM control.

of the simulation. In particular, given the HMM consisting of E, P^u ($u = 1, 2, 3$) 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 5 depicts the typical estimation results obtained with a window based implementation of the Viterbi algorithm.

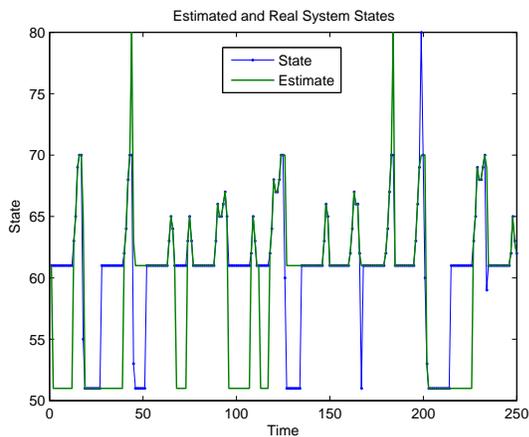


Fig. 5. Estimated \hat{s} and actual system states $s \in \mathcal{S} = \mathcal{X} \times \mathcal{Q}$ under AQM control.

VII. SIMULATIONS

In simulations carried out with Matlab, we chose the bottleneck link capacity as $C = 1,000$, and the number of users as $M = 10$. Table II summarizes these together with other simulation parameters for the dynamical system model in (3). A Monte Carlo algorithm with 100,000 sample points was used for constructing the Markov models (matrices P^u) with each of the three AQM schemes.

TABLE II
SIMULATION MODEL PARAMETERS FOR SYMMETRIC USERS

Link capacity	$C = 1000$
# of users	$M = 10$
Queue bounds	$q_{min} = 10, q_{max} = 100$
User parameters	$\kappa = 0.01, d = 0.01, \beta = 0.1$

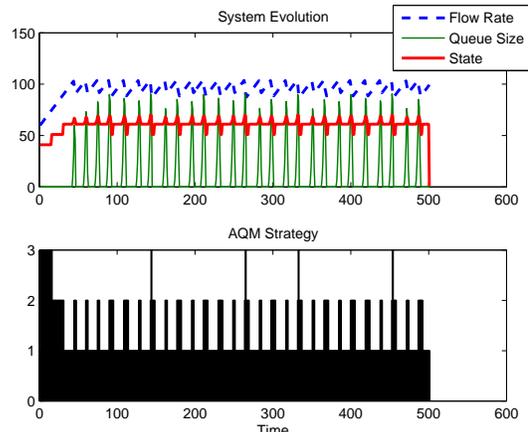


Fig. 6. Results of the simulation with a state-dependent optimal policy: (top) The evolutions of user flow rate x , link queue size q , the state s of the MM, and (bottom) the AQM scheme deployed versus time are shown.

We first simulated the system with the assumption that the full state is known. The AQM control actions correspond to the stationary optimal policy computed off-line as a result of the MDP described in Section VI. Figure 6 depicts the evolution of the controlled system – x (averaged over 10 users), q , and state s – as a function of time. Also shown is the sequence of control actions, i.e., specific AQM scheme at any given time instance. Note that the average queue size is approximately 15 with near capacity utilization. Averages of the non-equilibrium solution are consistent with the choice of reward function used for MDP and shown in Figure 2. In order to compare this with the standard AQM, we next simulated the system with AQM scheme 1 (droptail) used for all times. These results are shown in Figure 7. 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

- 1) larger user capacity utilization,
- 2) smaller queue fluctuations, and
- 3) 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 6) is notable.

In the subsequent simulation, we considered the more practical case where only the queue size was assumed to be known. For the HMM, we used the Viterbi algorithm (Sec-

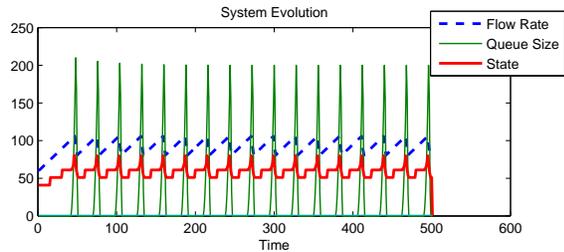


Fig. 7. Results of the simulation with the drop-tail AQM scheme: the evolutions of user flow rate x , link queue size q , the state s of the MM versus time are shown.

TABLE III
SIMULATION MODEL PARAMETERS FOR ASYMMETRIC USERS

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

tion VI) for state estimation. We formally assumed certainty equivalence to hold and treated the estimate as the state. Using stationary policy with respect to estimates still yielded promising results as shown in Figure 8. The estimation errors had little detrimental effect on the performance of the algorithm. Analysis of this will be the subject of future investigations.

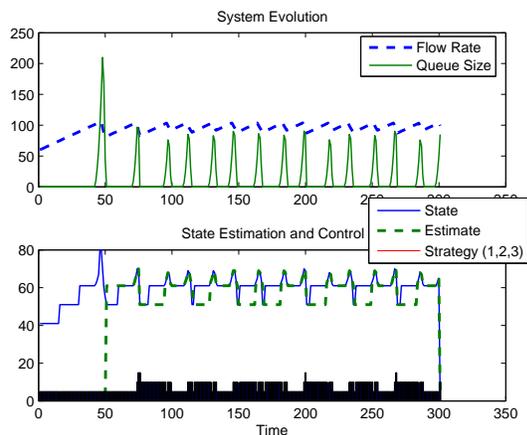


Fig. 8. Results of the simulation with a state estimation based feedback control: (top) the evolutions of user flow rate x and link queue size q ; (bottom) the real s vs. estimated \hat{s} state of the HMM and the AQM scheme deployed vs. time are shown.

Finally, we also carried out simulation based studies for an asymmetric multi-user case with $M = 10$ number of users. Table III summarizes the simulation parameters. The users are asymmetric because the parameters d_i and β_i for individual user i is picked from a uniform distribution whose range is indicated in the Table. In order to test the robustness of our method, we used the stationary policy and

estimation based on the reduced order symmetric model. The symmetric user flow rate variable x was *formally* replaced by the average flow rate for the M asymmetric users. As a result, there are modeling errors introduced because of the reduction in dimension (M states to a single state) and averaging. Note that the simulations were carried out with the $M + 1$ -dimensional asymmetric multi-user dynamical system (1) and (2). Figure 9 depicts the results of the simulation with estimation and control using the symmetric stationary policy. Remarkably, even though the dynamic behavior now is more complex (see for example, the queue trajectory), the performance shows good average capacity utilization (for x) and queue size q within limits as dictated by the reward function. Analysis of this will be considered in the journal version of this paper.

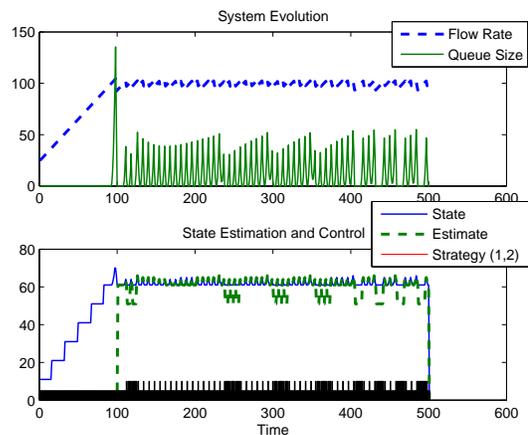


Fig. 9. Results of the simulation with a state estimation based feedback control for the asymmetric multi-user case: (top) the evolutions of average flow rate x and link queue size q ; (bottom) the real s and estimated state \hat{s} of the HMM and the AQM scheme deployed versus time are shown.

VIII. CONCLUSION

In this paper, we outlined a dynamical systems based framework for stochastic modeling, model computations, analysis, estimation, and MDP based control of nonlinear problems in communication networks. The framework was demonstrated for the problem of Active Queue Management (AQM). Even though the results were presented using a simple AQM control structure and a particular choice of optimality criterion, the framework is applicable more generally. The two ideas of set-oriented computations of Markov models and symmetry-based averaging over user flow rates were effectively used to manage complexity in estimation and control design. Extensive simulations with asymmetric multi-user models demonstrate both the applicability and robustness of our approach. A unified and rigorous theoretical framework for control and optimization of non-equilibrium dynamic behavior in communication networks together with its applications is a subject of continuing research investigation.

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] R. Srikant, *The Mathematics of Internet Congestion Control*, ser. Systems & Control: Foundations & Applications. Boston, MA: Birkhauser, 2004.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] V. Firoiu and M. Borden, "A study of active queue management for congestion control," in *Proc. IEEE Infocom*, Tel Aviv, Israel, March 2000.
- [13] 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
- [14] —, "Analysis and design of an adaptive virtual queue (AVQ) algorithm for active queue management," in *Proc. of ACM SIGCOMM*, August 2001, pp. 123–134.
- [15] 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.
- [16] 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.
- [17] A. Lasota and M. C. Mackey, *Chaos, Fractals, and Noise: Stochastic Aspects of Dynamics*. New York: Springer-Verlag, 1994.
- [18] M. Dellnitz and O. Junge, "On the approximation of complicated dynamical behavior," *SIAM Journal on Numerical Analysis*, vol. 36, pp. 491–515, 1999.
- [19] I. Mezic and A. Banaszuk, "Comparison of systems with complex behavior," *Physica D*, vol. 197, pp. 101–133, 2004.
- [20] M. Dellnitz and O. Junge, *Set oriented numerical methods for dynamical systems*. World Scientific, 2000, pp. 221–264.
- [21] G. Froyland, *Nonlinear Dynamics and Statistics: Proceedings, Newton Institute, Cambridge, 1998*. Birkhauser, 2001, ch. Extracting dynamical behaviour via Markov models, pp. 283–324.
- [22] K. Salamatian and S. Vaton, "Hidden Markov modeling for network communication channels," *SIGMETRICS Performance Evaluation Review*, vol. 29, no. 1, pp. 92–101, 2001.
- [23] 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.
- [24] 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.
- [25] S. M. Ross, *Applied Probability with Optimization Applications*. San Fransisco, CA: holden-Day, Inc., 1970.
- [26] 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.