

A NON-EQUILIBRIUM ANALYSIS & CONTROL FRAMEWORK FOR COMMUNICATION NETWORKS

Tansu Alpcan * **Paul Wang** and **Prashant G. Mehta** **
Tamer Başar ***

* *Deutsche Telekom Laboratories Ernst-Reuter-Platz 7, Berlin
10587 Germany. E-mail: tansu.alpcan@telekom.de*

** *Dept. of Mechanical Sci. & Eng., University of Illinois
1206 W. Green Street, Urbana IL 61801 USA
E-mails: (paulwang, mehtapg)@uiuc.edu*

*** *Coordinated Science Lab., University of Illinois
1308 West Main Street, Urbana IL 61801 USA
E-mail: 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. 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 initial results obtained confirm the theoretical analysis and exhibit promising performance when compared with well-known alternative schemes under persistent non-equilibrium queue behavior.

Keywords: Communication networks, active queue management, Markov models, optimal control.

1. INTRODUCTION

Communication networks such as the Internet exhibit a wide variety of complex *non-equilibrium* dynamic behavior. Examples of such behavior include self-excited user flow rate oscillations in the presence of delays, dynamic synchronization of the flows passing through the same bottleneck link, and chaotic behavior of user flows and queues at the routers. However, the optimization and control methods for such networks are primarily based upon an equilibrium framework.

We note the Kelly's framework (Kelly et al., 1998) or game theoretic framework (Alpcan and Başar, 2005) and refer to (Srikant, 2004) for an overview of the literature on the subject.

This paper is concerned with methods for 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 router sends advanced congestion notification to the users in order to regulate

the flow rates. This can be accomplished either by dropping packets or marking them. Since the earliest Droptail scheme, several AQM schemes have been proposed and studied in literature including RED, E-RED, REM, AVQ, and BLUE (Srikant, 2004).

We are primarily motivated by the analysis and control of non-equilibrium queue behavior which arises primarily as a result of nonlinear dynamics. 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 this paper, we represent for modeling as well as 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 (Lasota and Mackey, 1994) While, the dynamical model propagates the initial condition, the Perron-Frobenius (P-F) operator propagates uncertainty in initial condition. The main advantage of this approach is the often easier representation of complex asymptotic dynamic behavior as invariant probability measures of the P-F operator.

We use the recent set-oriented numerical methods for discretization of the dynamical systems (Dellnitz and Junge, 2002; Froyland, 2001). 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 for both the computational bifurcation analysis as well as for control design. The stochastic modeling approach we take enables us to carry out a *bifurcation analysis* to understand qualitative changes in queue behavior, *identification* of the fluid model parameters from packet level ns-2 simulations, and *control synthesis* for shaping non-equilibrium behavior.

The outline for the rest of the paper is as follows: In Section 2, a well-known network model of user and queue behavior is presented. Next, a stochastic modeling of the network model 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 4, an MDP-based framework for optimization and control of these models is summarized. The ns-2 simulation results demonstrating control of non-equilibrium behavior under the AQM algorithm developed are described in Section 5. The paper ends with the concluding remarks of Section 6.

2. DETERMINISTIC FLUID MODEL

2.1 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 (Srikant, 2004; Alpcan and Başar, 2005). Each user is associated with a unique connection for simplicity and transmits data with a nonnegative flow rate $x_i \in \mathbb{R}^+ \doteq [0, \infty)$ over this bottleneck link. 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).

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. (Hollot et al., 2002). The function $p(\cdot)$ 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

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

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 all $i, j \in \{1, \dots, M\}$. As a result of this symmetry, the linear subspace $S = \{\underline{x} \in \mathbb{R}^{+M} : x_i = x_1\}$ is a fixed-point space, which we will also refer to as the synchrony subspace. In particular, the subspace S is positively invariant with respect to dynamics of (1).

2.2 Synchronization to symmetric fixed-point space

Theorem 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 (Khalil, 1996). 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 $S \cap \Omega$.

The Lyapunov function is taken to be

$$V(\underline{x}) = \frac{1}{2} \sum_{i=2}^M (x_i - x_1)^2. \quad (4)$$

Using (1), the time derivative of (4) is

$$\dot{V}(\underline{x}) = -\kappa\beta \sum_{i=2}^M (x_i - x_1)^2 \cdot (x_i + x_1) \cdot p. \quad (5)$$

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

$$\Omega = \left\{ \underline{x} \in \mathbb{R}^{+M} : x_i \leq \max \left(C + B, \frac{1}{\sqrt{\beta d}} \right) \right\}. \quad (6)$$

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.

Finally, we show that $M \subset S \cap \Omega$. Using (5), 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),

$$x_i(t) = x_i(0) + \frac{\kappa}{d}t, \quad (7)$$

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

Let us denote the symmetric user flow rate as x . Then, in the fixed-point space the system dynamics are

$$\begin{aligned} \dot{x}(t) &= \kappa \left(\frac{1}{d} - \beta x(t)^2 p(t) \right), \\ \dot{q}(t) &= \begin{cases} Mx(t) - C & q \in (0, B), \\ \min(0, Mx(t) - C) & q = B, \\ \max(0, Mx(t) - C) & q = 0. \end{cases} \end{aligned} \quad (8)$$

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.

2.3 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 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 conclusions regarding the equilibrium solution:

(a) 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.

(b) For values of $\beta > \beta^*$, no equilibrium solution can exist. The asymptotic dynamics are non-equilibrium but lie in the fixed-point space S .

We omit the proof of the theorem due to space limitations.

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 stable, is clearly not desirable. As the feedback gain β is increased, one reaches the critical value

$$\beta^* = \frac{1}{d} \left(\frac{C}{M} \right)^{-2} \quad (9)$$

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 (8) 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




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

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.

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 (9), 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 $M < M^*$.

For the fluid-approximation (8), 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 (9) the parameter

$$\kappa\beta = \frac{\kappa}{d} \left(\frac{C}{M^*} \right)^{-2}. \quad (10)$$

Table 1 depicts the identified nominal model parameters for the fluid approximation (8). 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.

Table 1. Model parameters

Link capacity	$C = 1000$
# of users	$M = 10, M^* = 210$
Queue bounds	$Q_{min} = 0, Q_{max} = 100$
User parameters	$\kappa = 0.05, d = 0.01, \beta = 4.4$

3. STOCHASTIC MODEL AND BIFURCATION ANALYSIS

We characterize stochastic descriptions such as invariant measures (and their numerical approximations) for the dynamics of the deterministic model in Section 2. The stochastic model presented here is included for completeness and for a detailed description we refer to (Alpcan et al., 2006).

Let us consider a dynamical system $T : X \times Q \rightarrow \mathbb{R}^{+2}$ obtained by sampling the solutions of the ODE (8). In particular, let $\phi(t; x_0, q_0)$ be the solution operator for the fluid approximation ODE (8) and set $T(x_0, q_0) = \phi(\Delta t; x_0, q_0)$. The sampling time Δt is taken to be sufficiently small and $(x_0, q_0) \in X \times Q$ denote the initial condition. $X \subset \mathbb{R}^+$ 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 $\mathbb{P}[\mu](A) = \mu(T^{-1}(A))$, 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., $\mathbb{P}\mu_1 = 1 \cdot \mu_1$.

We utilize set-oriented numerical methods for constructing efficient finite-dimensional approximations of the P-F operator (Dellnitz and Junge, 2002; Froyland, 2001). The approximation arises as a Markov (transition probability) matrix defined with respect to a finite (L) partition \mathcal{S} of the phase space S . The resulting matrix is non-negative and if $T : D_i \rightarrow S$, then $\sum_{j=1}^L P_{ij} = 1$, i.e. P is a Markov or a row-stochastic matrix. Notice that, P is interpreted as a randomly perturbed approximation of \mathbb{P} and P converges to \mathbb{P} in the L^2 norm sense as the partition gets finer and finer.

We construct here the partition \mathcal{S} for the stochastic approximation of the network model by uniformly quantizing the user flow-rates (in X) and queue size (in Q)

between a lower and upper bound in way similar to the one in (Alpcan et al., 2006). The two quantized partitions are denoted as $\mathcal{X} = [X_1, \dots, X_{12}]$ and $\mathcal{Q} = [Q_1, \dots, Q_{11}]$ and the partition size $L = 132$. We define $\mathcal{S} \doteq \mathcal{X} \times \mathcal{Q}$.

One approach for numerically approximating the Markov matrix P is to use a Monte Carlo algorithm with several short term simulations using the dynamical system T . As one takes finer partitions, the invariant measure of P converges to a weak limit that approximates the invariant measure of the P-F operator \mathbb{P} .

We next carry 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 3 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. The L^1 -error between the invariant measure computed using the MM and the ergodic average of the time-series data is at most 5% indicating that the MM is fairly accurate in capturing the asymptotic behavior with the time-domain simulation over a range of values of M .

4. CONTROL MARKOV CHAIN AND OPTIMIZATION

There are infinitely many possibilities for probabilistic AQM control schemes (such as RED) where packet drop probability is defined as a function of the queue size, which in turn corresponds to selecting a specific $p(\cdot)$ function in the deterministic fluid model (8). Each such marking (or dropping) scheme p can be analyzed using the stochastic models of the previous section. In this paper, we define a non-equilibrium queue management (NEQM) algorithm as an illustrative example using four separate packet marking (or dropping) schemes p_u where $u \in U$ and $U := 0, \dots, 3$. The scheme p_0 corresponds to the droptail whereas the schemes p_1 through p_3 can be interpreted as progressively aggressive variants of RED. Table 2 summarizes the packet marking for the intermediate queue-size, where $q_{min} = 10$ and $q_{max} = 100$.

Table 2. Control Policies

Policy	$q \leq q_{min}$	$q_{min} < q < q_{max}$	$q \geq q_{max}$
p_0	0	0	1
p_1	0	0.3	1
p_2	0	0.6	1
p_3	0	1	1

Each AQM scheme, p_u , modifies the dynamical system (8) and leads to a different *control Markov chain* P^u , where $u \in U$ corresponds to the one in 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 . We next pose the control problem

as a Markov Decision Process (MDP) by defining a specific real-valued reward function $R(s)$ (shown in Figure 4) over the set of states \mathcal{S} and set of the control action U . The reward function reflects the expectation that a positive but moderate queue size will ensure maximum capacity utilization while preventing large queue sizes. We




Fig. 4. The reward function on $\mathcal{X} \times \mathcal{Q}$.

solve the MDP obtained using standard methods in a way similar to the one in (Alpcan et al., 2006). Figure 5 depicts the resulting stationary optimal policy solution μ^* on $\mathcal{X} \times \mathcal{Q}$.




Fig. 5. The optimal policy μ^* on $\mathcal{X} \times \mathcal{Q}$ obtained by solving the MDP.

Using the optimal policy μ^* , we also compute the invariant probability measures of the controlled Markov matrix P^{μ^*} for user flow rates (in \mathcal{X}) and queue size (in \mathcal{Q}). The invariant measures exhibit low queue sizes along with full link capacity usage. In addition, they compare favorably with the asymptotic distributions obtained from the time-domain simulations of the dynamical system (8) under the optimal policy.

5. SIMULATIONS

We 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.

Figure 6(a) and (b) compares the queue and the user behavior obtained with RED, and the NEQM. We observe more efficient queue utilization under the




Fig. 3. Bifurcation diagram in terms of the invariant measure: (left) ns-2 dynamics, (middle) the time-averaged dynamics, and (right) the invariant measure of the Markov model as the parameter M increases.

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 have observed in other simulations that there is a strong contrast in the mean user flow rate between NEQM and the droptail scheme, which exhibits heavy oscillations.




Fig. 6. Results of the ns-2 simulations with (a) RED and (b) the NEQM schemes.

6. CONCLUSION

In this paper, we have outlined an ergodic theoretic framework for stochastic modeling, model computations, analysis, identification, and MDP based optimal control of nonlinear dynamics in communication networks, specifically the AQM problem. Through equilibrium and local bifurcation analysis as well as ns-2 simulations, we have illustrated the fact that methods for analysis and optimization of non-equilibrium

queue behavior are *essential* for a better understanding of the problem. In addition, based on the stochastic approximations of a simplified fluid model, an optimization problem has been defined and solved via standard MDP methods to obtain a specific optimal stationary AQM policy, called NEQM, as an example. The NEQM scheme developed has been observed to perform better than the droptail and RED schemes in packet-level ns-2 simulations.

Acknowledgment: The authors thank Uday Shanbhag and Umesh Vaidya for helpful discussions.

REFERENCES

- T. Alpcan and T. Başar. A utility-based congestion control scheme for Internet-style networks with delay. *IEEE Transactions on Networking*, 13(6): 1261–1274, December 2005.
- 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.
- M. Dellnitz and O. Junge. *Set oriented numerical methods for dynamical systems*, pages 221–264. World Scientific, 2002.
- G. Froyland. *Nonlinear Dynamics and Statistics: Proceedings, Newton Institute, Cambridge, 1998*, chapter Extracting dynamical behaviour via Markov models, pages 283–324. Birkhauser, 2001.
- C. V. Hollot, V. Misra, D. Towsley, and W.B. Gong. Analysis and design of controllers for AQM routers supporting TCP flows. *IEEE Transactions on Automatic Control*, pages 945–959, June 2002.
- F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- H. K. Khalil. *Nonlinear Systems*. Prentice Hall, Upper Saddle River, NJ, 2nd edition, 1996.
- A. Lasota and M. C. Mackey. *Chaos, Fractals, and Noise: Stochastic Aspects of Dynamics*. Springer-Verlag, New York, 1994.
- R. Srikant. *The Mathematics of Internet Congestion Control*. Systems & Control: Foundations & Applications. Birkhauser, Boston, MA, 2004.