Nonlinear Stability Analysis for a Class of TCP/AQM Networks¹

C.V. Hollot ECE Department University of Massachusetts Amherst, MA 01003

Abstract

Recent work has shown the benefit of using proportional feedback in TCP/AQM networks. By proportional feedback we mean the marking probability is proportional to the instantaneous queue length. Our earlier work relied on linearization of nonlinear fluidflow models of TCP. In this work we address these nonlinearities directly and establish some stability results when the marking is proportional. In the case of delayfree marking, we show the system's equilibrium point to be asymptotically stable for all proportional gains. In the more realistic case of delayed feedback, we establish local asymptotic stability and quantify a region of attraction.

1 Introduction

The reliable and efficient exchange of information across the Internet has been a key ingredient to its explosive growth and utilization. At the heart of this information exchange is the Transmission Control Protocol (TCP). Under TCP, a sender has authority to set its transmission rate using a window flow-control mechanism. The sender continuously probes the network's available bandwidth and increases its window size to garner maximum share of network resource. For each successful end-to-end packet transmission TCP increases the sender's window size. Conversely, TCP cuts the window in half whenever a sender's packet does not reach the receiver. Such packet losses can affect network performance by decreasing the sender's effective transmission rate and increasing delay due to packet retransmission. By itself, TCP has no information of network mechanisms contributing to packet loss - such as the congested router shown in the simple senderreceiver connection of Figure 1. Thus, routers must assume a role in network management by sensing congestion and preemptively signaling TCP rather than have it react to unreceived packets.

Y. Chait MIE Department University of Massachusetts Amherst, MA 01003



Figure 1: Router congestion results in lost packets. The receiver then signals sender to decrease window size.

The simplest form of such active queue management (AQM), termed *drop tail*, drops arriving packets when the router's buffer is full. Drawbacks of this scheme include flow-synchronization (see [1]) and performance degradation due to the excessive time-outs and restarts arising when the trailing end of a sequence of data packets is dropped. Motivated by drop-tail's inefficiencies, the random early detection (RED) scheme was introduced in [1]. Rather than waiting for buffer overflow to occur, RED anticipates congestion by measuring the router's average queue length and throttling the sender's rate accordingly. Since TCP is an end-toend protocol. RED achieves this signaling indirectly by randomly marking packets and routing them to the receiver.¹ The receiver, in turn, completes the feedback by acknowledging the receipt of marked packets to the sender; this is depicted in Figure 2 where we emphasize the implicit, delayed, feeding-back of acknowledgment packets.² Upon receipt of such acknowledgments, the sender adjusts its rate according to the TCP algorithm. The randomness in RED's packet-marking scheme was meant to eliminate flow-synchronization and introduce fair-marking while queue-averaging was introduced to attenuate the effects of bursty traffic on the feedback signal. A crucial drawback in deploying RED stems from tuning difficulties where the performance of RED can approach that of a drop-tail router; e.g., see [2],

 $^{^1\}mathrm{This}$ work is supported in part by the National Science Foundation under Grant CMS-9800612 and by DARPA under Contract DOD F30602-00-0554.

 $^{^{1}}$ By marking packets we refer to the action of either dropping packets or setting their explicit congestion notification (ECN) bit.

 $^{^{2}}$ This time delay is equivalent to one round-trip time which is comprised of propagation and queuing delays.



Figure 2: RED randomly marks packets to anticipate congestion.

[3] and [4]. Motivated by these deficiencies in the basic RED mechanism, researchers have proposed modifications; e.g., see [5] - [7]. Our recent research is motivated by a desire to design advanced AQM schemes which explicitly rely on dynamic modeling and feedback control principles. Central to our approach is the recognition that AQM schemes are essentially feedback control systems and that the principles of control theory can provide critical insight and guidance into the analysis and design of such schemes. While such principles can be found in the study of ATM networks (see for example [8] and [9] and the references cited therein) they have not been applied to TCP-controlled flows. Their absence from the design scene so far is apparently due to a lack of an analytical model of TCP. Fortunately, this roadblock has been recently removed in [10] with the introduction of a fluid-flow model that expresses TCP in a language that allows control engineers to analyze and design AQM schemes.³ Indeed, in and [13], [14] and [15] we have done just that by:

- 1. Relating key network parameters such as TCP load, router capacity and round-trip time to the stability and performance of AQM systems.
- 2. Analyzing RED and showing that queue *averaging* is not advisable from a feedback control viewpoint.
- 3. Proposing improved AQM schemes which amount to classical proportional (P) and proportional-integral (PI) feedback controllers.⁴

The objective of this paper is to continue this line of research and more formally address the nonlinearities appearing in the TCP fluid-flow model [10].⁵ Our work

in [14] and [15] relied on small-signal linearization of these nonlinearities. Specifically, assuming a proportional packet-marking strategy we will show in this paper: 1) the resulting closed-loop dynamics are globally asymptotically stable if the marking is delay-free. 2) the equilibrium point is local asymptotic stability under delayed-marking. In the latter case we will identify a region of convergence whose radius is inversely proportional to the gain-delay product.

In the next section we introduce the nonlinear TCP dynamics and a proportional AQM scheme which together forms the basis of our analysis. In Sections 3 and 4 we conduct stability analyzes for the delay-free and delayed marking cases respectively. We conclude in Section 5.

2 Dynamics of TCP/AQM

2.1 A fluid-flow model of TCP

In [10], a dynamic model of TCP behavior was developed using fluid-flow and stochastic differential equation analysis. Simulation results demonstrated that this model accurately captured the dynamics of TCP. In this paper we use a simplified version of that model which ignores the TCP timeout mechanism. This model relates the average value of key network variables and is described by the following coupled, nonlinear differential equations:

$$\begin{aligned} \dot{W}(t) &= \frac{1}{R(t)} - \frac{W(t)}{2} \frac{W(t - R(t))}{R(t - R(t))} p(t - R(t)) \\ \dot{q}(t) &= N(t) \frac{W(t)}{R(t)} - C \end{aligned} \tag{1}$$

where \dot{x} denotes the time-derivative of x and

- $W \doteq$ average TCP window size (packets);
- $q \doteq$ average queue length (packets);

$$R \doteq \text{round-trip time} = \frac{q}{C} + T_p \text{ (secs)};$$

- $C \doteq$ queue capacity (packets/sec);
- $T_p \doteq$ propagation delay (secs);
- $N \doteq$ number of TCP sessions;
- $p \doteq$ probability of packet mark.

The queue length q and window-size W are positive and bounded quantities; i.e., $q \in [0, q_{max}]$ and $W \in [0, W_{max}]$ Also, the marking probability p takes value only in [0, 1]. We illustrate these differential equations in the block diagram of Figure 3 which highlights TCP window-control and queue dynamics. We now consider an AQM scheme in which the probability of packet

 $^{^{3}}$ We'd also like to point out other recent TCP models in [11] and [12]. In [13] we discuss them in the context of our fluid-flow model.

⁴Both P and PI controllers have been suggested by others; see [16] and [17] respectively. One of our contributions was to provide feedback-control rationale for these schemes and to give systematic design rules.

⁵Related, nonlinear stability analysis for optimization-based flow control has been conducted in [18]. While [18] considers rate-based congestion control, our paper addresses window-based

schemes that necessarily rely on the nonlinear dynamics of TCP. We also address the effect of time-delay in marking which is missing in [18].

mark is proportional to the instantaneous queue length.



TCP window control

Figure 3: Block-diagram of TCP's congestion-avoidance mode.

2.2 Proportional AQM scheme

Active queue management is a core process wherein packets are marked as a function of queue length; see Figure 4. Marking packets in this way amounts to reducing TCP source rates as queue length grows. The objective is to manage buffer size as a means for regulating buffer utilization and queuing delay. Recent AQM schemes include RED, random exponential marking (REM) in [17] and PI control. In RED, the loss is made roughly proportional to average queue length. In both REM and PI, the loss is a weighted sum of instantaneous queue length and its integral. In this paper we will analyze the stability of a proportional scheme wherein loss probability is proportional to the instantaneous queue length and described by $p = K_p q$ with $K_p > 0$ and $p \in [0, 1]$. Combining with (1) gives the closed-loop dynamics

$$\dot{W}(t) = \frac{1}{R(t)} - \frac{W(t)}{2} \frac{W(t - R(t))}{R(t - R(t))} K_p q(t - R(t))$$

$$\dot{q}(t) = N(t) \frac{W(t)}{R(t)} - C.$$
 (2)

We now assume that the TCP load N and roundtrip time R are time-invariant; i.e., $N(t) \equiv N$ and $R(t) \equiv R$. The former assumes that the number of TCP flows is constant, while the latter may be a good approximation when round-trip time is dominated by the propagation delay. This occurs when the capacity C is large. Additionally, in [13] we have shown that the right-hand side of the \dot{W} equation in (2) can be approximated by

$$\frac{1}{R} - \frac{W(t)W(t)}{2R}K_pq(t-R)$$

when the window size satisfies $W \gg 1$. Under these conditions the equilibrium points (W_0, q_0) for (2) are given by

$$W_0 = \frac{RC}{N}; \quad q_0 = \frac{2N^2}{R^2 C^2 K_p}.$$
 (3)

In the transformed variables $\tilde{W} \doteq W - W_0$, $\tilde{q} \doteq q - q_0$, (2) becomes

$$\dot{\tilde{W}}(t) = -\frac{(\tilde{W}(t) + W_0)^2}{2R} K_p \tilde{q}(t - R) \qquad (4) - \frac{\tilde{W}^2(t) + 2\tilde{W}(t)W_0}{2R} K_p q_0 \dot{\tilde{q}}(t) = \frac{N}{R} \tilde{W}(t). \qquad (5)$$

The equilibrium point for (5) is then $(\tilde{W}, \tilde{q}) = (0, 0)$. Note that $\tilde{q} \ge -q_0$ since q > 0.



Figure 4: AQM signals TCP sources to reduce their window sizes as a function of instantaneous queue length.

3 Delay-free marking

In this section we conduct a first analysis under the assumption of delay-free marking; i.e., $p(t) = K_p q(t)$. We will show that the AQM system's equilibrium point (W_0, q_0) is asymptotically stable for all positive gains K_p . In the next section we will study the case of delayed marking. For simplicity we drop the explicit dependence on time and write (5) as

$$\dot{\tilde{W}} = -\frac{(\tilde{W}(t) + W_0)^2}{2R} K_p \tilde{q} - \frac{\tilde{W}^2 + 2\tilde{W}W_0}{2R} K_p q_0$$
$$\dot{\tilde{q}} = \frac{N}{R} \tilde{W}.$$
(6)

Eliminating \tilde{W} from (6) gives the second-order differential equation

$$\ddot{\tilde{q}} + \beta b(\dot{\tilde{q}}) K_p(\tilde{q} + q_0) + \alpha K_p \tilde{q} = 0$$
(7)

where

$$b(\dot{\tilde{q}}) \doteq 0.5\dot{\tilde{q}}^2 + C\dot{\tilde{q}};$$

$$\alpha \doteq \frac{C^2}{2N};$$

$$\beta \doteq \frac{1}{N}.$$
(8)

Consider the positive-definite Lyapunov candidate

$$V(\tilde{q}, \dot{\tilde{q}}) = \frac{1}{2}\dot{\tilde{q}}^2 + \frac{1}{2}\alpha K_p \tilde{q}^2.$$
 (9)

Along solutions to (7) compute the time-derivative

Since $\dot{\tilde{q}} = \dot{q} = \frac{NW}{R} - C$ and $\frac{NW}{R} > 0$ then $\dot{\tilde{q}} > -C$. Hence, $\dot{\tilde{q}}b(\dot{\tilde{q}}) = \dot{\tilde{q}}^2(0.5\dot{\tilde{q}} + C) \ge 0$. Consequently,

$$\dot{V} = -\beta \dot{\tilde{q}} b(\dot{\tilde{q}}) K_p(\tilde{q} + q_0) \le 0$$

for all $\dot{\tilde{q}} \geq 0$ and all $\tilde{q} \geq -q_0$. This proves the equilibrium point is stable. To conclude asymptotic stability we invoke a standard invariant set result, see [19], which states that trajectories of (6) converge to the largest invariant set contained in

$$\mathcal{M} \stackrel{\doteq}{=} \{ (\tilde{q}, \dot{\tilde{q}}) : \dot{V} = 0 \}$$
$$= \{ (\tilde{q}, \dot{\tilde{q}}) : \dot{\tilde{q}} = 0 \text{ or } \tilde{q} = -q_0 \}$$

However, the only invariant set contained in \mathcal{M} is the equilibrium point (0,0). Indeed, if $(\tilde{q}(t), \dot{\tilde{q}}(t))$ is equal to $(\tilde{q}(t), 0)$ or $(-q_0, \dot{\tilde{q}}(t))$ then, using (7), $(\tilde{q}(t^+), \dot{\tilde{q}}(t^+)) \notin \mathcal{M}$. This proves asymptotic stability which we now summarize:

Theorem 1: The equilibrium point of (6) is asymptotically stable for all $K_p > 0$.

4 Delayed marking

We now consider the case when packet-marking is delayed by the round-trip time R. In this case, the proportional AQM algorithm is $p(t-R) = K_p q(t-R)$ and we write (7) as

$$\ddot{\tilde{q}}(t) + \beta b(\dot{\tilde{q}}(t)) K_p(\tilde{q}(t-R) + q_0) + \alpha K_p \tilde{q}(t-R) = 0 \quad (10)$$

where q_0 , α and β are defined in (3) and (8). In the subsequent analysis we will rely on the fact that $\dot{\tilde{q}}$ is bounded. This follows since $\dot{\tilde{q}} = \frac{NW}{R} - C$ and since the window size W is bounded. Using Lyapunov-Razumikhin formalism for delay-differential equations (see [20]) we again consider the positive-definite candidate (9). Using the identity

$$\tilde{q}(t-R) = \tilde{q} - \int_{-R}^{0} \dot{\tilde{q}}(t+\theta)d\theta,$$

we compute

$$\dot{V} = -\beta \dot{\tilde{q}} b(\dot{\tilde{q}}) K_p(\tilde{q}+q_0) + \dot{\tilde{q}}(\beta b(\dot{\tilde{q}})+\alpha) K_p \int_{-R}^{0} \dot{\tilde{q}}(t+\theta) d\theta$$

Following the formalism, we first assume a $\lambda>1$ such that

$$V(\xi) \le \lambda V(t) \tag{11}$$

for all $t - R \le \xi \le t$. Asymptotic stability is demonstrated if condition (11) implies $\dot{V} < 0.^6$ To this end, (11) indeed implies

$$\dot{\tilde{q}}(\xi) \le \lambda \|z(t)\|$$

for all $t - R \leq \xi \leq t$ where $z(t) \doteq [\alpha \tilde{q}(t) \ \dot{\tilde{q}}(t)]$. Hence,

$$\int_{-R}^{0} \dot{\tilde{q}}(t+\theta)d\theta \le \lambda R \|z(t)\|$$

for all $t - R \le \xi \le t$. As a result,

$$\dot{V} \le -\beta \dot{\tilde{q}} b(\dot{\tilde{q}}) K_p(\tilde{q} + q_0 - \lambda R \| z(t) \|) + \alpha K_p \lambda R \| z(t) \| \dot{\tilde{q}}.$$
(12)

Since both \tilde{q} and $\dot{\tilde{q}}$ are bounded it is possible to take $\alpha K_p = \frac{K_p}{N}$ small enough to force $K_p ||z(t)|| \approx K_P |\dot{\tilde{q}}|$ and make negligible the second term on the right-hand side of (12). In this case,

$$\dot{V} \le -\beta \dot{\tilde{q}} b(\dot{\tilde{q}}) K_p(\tilde{q} + q_0 - \lambda R |\dot{\tilde{q}}|).$$

Since $q_0 > 0$, there exists a $(\tilde{q}, \dot{\tilde{q}})$ -neighborhood over which $K_p(\tilde{q} + q_0 - \lambda R |\dot{\tilde{q}}|) > 0.^7$ Using an invariant-set argument as in the previous section we conclude (10) to be locally asymptotically stable. A region of attraction is the level-set interior:

$$V_L \doteq \{ (\tilde{q}, \dot{\tilde{q}}) : V(\tilde{q}, \dot{\tilde{q}}) < L \}$$

$$(13)$$

where L satisfies

$$(\tilde{q}, \dot{\tilde{q}}) \in V_L \quad \Rightarrow \quad K_p(\tilde{q} + q_0 - \lambda R |\dot{\tilde{q}}|) > 0.$$

A suitable L (see Appendix) is

$$L = \min\left\{\frac{1}{4K_P^2 R^2}, \frac{N^4}{16R^6 C^4 K_p^2}\right\}.$$
 (14)

We summarize.

Theorem 2: For sufficiently small $\frac{K_p}{N} > 0$, the equilibrium point of (10) is locally asymptotically stable. A region of attraction is given by (13) and (14).

5 Conclusion

In this paper we have conducted a stability analysis of TCP/AQM based on the fluid-flow model developed in [10]. We considered the simplest of packet-marking rules in which marking is proportional to buffer queue

 $^{^{6}}$ See Theorem 4 in [20].

⁷We must include $K_p > 0$ in this inequality to account for $p \in [0, 1]$ which, in turn, implies $K_p(\tilde{q} + q_0) \in [0, 1]$.

length. We show that the resulting closed loop dynamics are globally asymptotically stable when the marking is delay-free. When the marking's action is delayed by the round-trip time, the stability results become local in nature. We show that the region of attraction is inversely proportional to the gain-delay product. There are at least three areas in which we would like to expand this analysis. First, to include proportional-integral marking as introduced in [15] and [17]. Secondly, to include a queuing delay component $\left(\frac{q}{C}\right)$ in the round-trip time. The time-delay then becomes state-dependent. Thirdly, to allow for time-varying load flow as modeled by time-variation in N.

6 Acknowledgment

We would like to thank Professor Vishal Misra of the Computer Science Department, Columbia University for many fruitful discussions.

A Region of Attraction

We seek an L > 0 such that

$$(\tilde{q}, \dot{\tilde{q}}) \in V_L \quad \Rightarrow \quad K_p(\tilde{q} + q_0 - \lambda R |\dot{\tilde{q}}|) > 0$$

where

$$V_L \doteq \{ (\tilde{q}, \dot{\tilde{q}}) : V(\tilde{q}, \dot{\tilde{q}}) < L \}$$

and

$$V(\tilde{q},\dot{\tilde{q}}) = \frac{1}{2}\dot{\tilde{q}}^2 + \frac{1}{2}\alpha K_p \tilde{q}^2$$

Since $K_p(\tilde{q} + q_0) < 1$, then $L = \min\{L_1, L_2\}$ where $L_2 = V(0, \frac{1}{K_p R}) = \frac{1}{4K_p^2 R^2}$ and L_1 renders

$$V(\frac{1}{R}\tilde{q} + \frac{1}{R}q_0, \dot{\tilde{q}}) = L_1$$

to have solution $\tilde{q} \geq q_0$. Computation gives $L_1 = \frac{N^4}{16R^6C^4K_p^2}$.

References

[1] S. Floyd and V. Jacobson, "Random Early Detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, Vol. 1, August 1997.

[2] M. Christiansen, K. Jeffay, D. Ott, and F. Smith, "Tuning Red for web traffic," in *Proceedings of ACM/SIGCOMM*, 2000.

[3] M. May, T. Bonald, and J.-C. Bolot, "Analytic Evaluation of RED Performance," in *Proceedings of IN-FOCOM 2000*.

[4] W. Chang Feng, D. D. Kandlur, D. Saha, and K. G. Shin, "A self configuring red gateway." IEEE INFOCOM 1999. [5] T. J. Ott, T. V. Lakshman, and L. H. Wong, "SRED: Stabilized RED," in *Proceedings of INFOCOM* 1999.

[6] W. Feng, D. D. Kandlur, D. Saha, and K. G. Shin, "A Self-Configuring RED Gateway," in *Proceedings of INFOCOM 1999.*

[7] D. Lin and R. Morris, "Dynamics of Random Early Eetection," in *Proceedings of ACM/SIGCOMM*, 1997.

[8] Y. Zhao, S.Q. Li and S. Sigarto, "A linear dynamic model for design of stable explicit-rate ABR control schemes," in *Proceedings of IEEE Infocom'97*, 1997.

[9] E. Altman, T. Basar and R. Srikant, "Robust rate control for ABR sources," in *Proceedings of IEEE INFOCOM'98*, 1998.

[10] V. Misra, W. B. Gong, and D. Towsley, "Fluidbased Analysis of a Network of AQM Routers Supporting TCP Flows with an Application to RED," in *Proceedings of ACM/SIGCOMM*, 2000.

[11] S. Mascolo, "Congestion control in highspeed communication networks," *Automatica*, vol. 35, pp. 1921–1935, March 1999.

[12] F. Kelly, "Mathematical modelling of the Internet," in *Mathematics Unlimited - 2001 and Beyond*, 2000.

[13] C. V. Hollot, V. Misra, D. Towsley, and W. B. Gong, "Analysis and design of controllers for aqm routers supporting tcp flows." submitted to special issue of *IEEE Transactions on Automatic Control* on "Systems and Control Methods for Communication Networks."

[14] C. V. Hollot, V. Misra, D. Towsley, and W. B. Gong, "A control theoretic analysis of RED." in *Proceedings of INFOCOM 2001*.

[15] C. V. Hollot, V. Misra, D. Towsley, and W. B. Gong, "On designing improved controllers for aqm routers supporting tcp flows." in *Proceedings of INFO-COM 2001*.

[16] M. May, C. Diot, B. Lyles, and J. Bolot, "Influence of Active Queue Management Parameters on Aggregate Traffic Performance." Work in progress. ftp://ftp.sprintlabs.com/diot/aqm.zip.

[17] S. Athuraliya, V. H. Li, S.H. Low and Q. Yin, "REM: Active Queue Management," in *IEEE Network*, May/June, 2001.

[18] F. Paganini, "On the Stability of Optimization-Based Flow Control," in *Proceedings of the American Control Conference*, Arlington, VA, 2001, pg. 4869.

[19] J.E. Slotine and W. Li, *Applied Nonlinear Control*, Prentice Hall, Englewood Cliffs, NJ, 1991.

[20] S.-I. Niculescu, E.I. Verriest, L. Dugard and J.-M. Dion, "Stability and Robust Stability of Time-Delay Systems: A Guided Tour," Stability and Control of Time-delay Systems, Eds L. Dugard and E.I. Verriest, Lecture Noes in Control and Information Sciences, Springer, 1998.