

REFERENCES

- [1] M. Cannon and B. Kouvaritakis, "Infinite horizon predictive control of constrained continuous-time linear systems," *Automatica*, vol. 36, pp. 943–955, 2000.
- [2] W. H. Chen, D. J. Ballance, and J. O'Reilly, "Model predictive control of nonlinear systems: Computational burden and stability," *Proc. Inst. Elect. Eng. D*, vol. 147, no. 4, pp. 387–394, 2000.
- [3] F. A. Cuzzola, J. C. Geromel, and M. Morari, "An improved approach for constrained robust model predictive control," *Automatica*, vol. 38, no. 7, pp. 1183–1189, 2002.
- [4] L. Dai, *Singular Control Systems*. New York: Springer-Verlag, 1989, vol. 118, Lecture Notes in Control and Information Sciences.
- [5] C. E. Garcia, D. M. Prett, and M. Morari, "Model predictive control: Theory and practice—A survey," *Automatica*, vol. 25, no. 3, pp. 335–348, 1989.
- [6] M. Johansson and A. Rantzer, "Computation of piecewise quadratic Lyapunov functions for hybrid systems," *IEEE Trans. Automat. Contr.*, vol. 43, pp. 555–559, Apr. 1998.
- [7] M. V. Kothare, V. Balakrishnan, and M. Morari, "Robust constrained model predictive control using linear matrix inequalities," *Automatica*, vol. 32, no. 10, pp. 1361–1379, 1996.
- [8] B. Kouvaritakis, J. A. Rossiter, and J. Schuurmans, "Efficient robust predictive control," *IEEE Trans. Automat. Contr.*, vol. 45, pp. 1545–1549, Sept. 2000.
- [9] F. L. Lewis, "A survey of linear singular systems," *Circuits, Syst., Signal Processing*, vol. 5, no. 1, pp. 3–36, 1986.
- [10] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, pp. 789–814, 2000.
- [11] J. B. Rawlings and K. R. Muske, "The stability of constrained receding horizon control," *IEEE Trans. Automat. Contr.*, vol. 38, pp. 1512–1516, Nov. 1993.
- [12] A. Weinmann, *Uncertain Models and Robust Control*. New York: Springer-Verlag, 1991.
- [13] F. Wu, "LMI-based robust model predictive control and its application to an industrial CSTR problem," *J. Process Control*, vol. 11, pp. 649–659, 2001.
- [14] L. Xie, "Output feedback H_∞ control of systems with parameter uncertainty," *Int. J. Control*, vol. 63, no. 4, pp. 741–750, 1996.
- [15] S. Xu and C. Yang, "Robust stabilization of generalized state-space systems with uncertainty," *Int. J. Control*, vol. 72, no. 18, pp. 1659–1664, 1999.
- [16] ———, "Stabilization of discrete-time singular systems: A matrix inequalities approach," *Automatica*, vol. 35, pp. 1613–1617, 1999.
- [17] L. Zhang, B. Huang, and J. Lam, "LMI synthesis of H_2 and mixed H_2/H_∞ controllers for singular systems," *IEEE Trans. Circuits Syst. II*, vol. 50, pp. 615–626, Sept. 2003.

Instability of a Tandem Network and Its Propagation Under RED

Richard J. La

Abstract—The interaction between a random early detection (RED) gateway and transmission control protocol (TCP) connections has been shown to lead to a rich set of nonlinear phenomena in single bottleneck cases. We extend this work and study the interaction of TCP connections with RED gateways in a simple tandem network, using a nonlinear first-order discrete-time model. We demonstrate that the nonlinear behavior of TCP can result in both smooth and nonsmooth bifurcations, leading to chaos. We show that the instabilities can be induced at both bottlenecks by changing the system parameters only at one of the bottlenecks while fixing the parameters at the other, thus demonstrating the propagation of instability. Moreover, we show that locally sufficient conditions for stability based on single node analysis are not sufficient for global network stability.

Index Terms—Bifurcation, congestion control, network stability.

I. INTRODUCTION

The random early detection (RED) mechanism, proposed by Floyd and Jacobson [3], attempts to control the congestion level at a bottleneck by monitoring the average queue size. Although the RED mechanism is conceptually simple and easy to understand, its interaction with transmission control protocol (TCP) connections has been shown to be rather complex and is not well understood [7], [9], [12]. Ranjan *et al.* have used a simple nonlinear model to investigate the behavior of a simple single link network with a RED gateway and TCP connections [12]. They have demonstrated that such a system leads to nonlinear phenomena, such as oscillations and chaos, if the control parameters are not selected carefully [12].

Although the RED was proposed almost a decade ago, it has not been widely deployed in practice, mainly due to a lack of agreement on parameter setting and/or evidence that it significantly improves the performance [1]. This problem is further complicated by a lack of understanding of how congestion in one part of network affects another part of the network [5]. In this note we take the first step toward remedying this situation. Our starting point is [12], which discusses the parametric sensitivity of RED in the context of only *single* bottleneck. Hence, [12] does not address the issues in multiple bottleneck cases we are interested in. We extend the model for single bottleneck cases in [12] to a *tandem* network and study the interaction of the RED mechanism with TCP connections. This tandem network can be viewed as a network with two dominant bottleneck links.

It turns out even this simple tandem network has enough structure to illustrate some of important issues in multiple bottleneck cases that are of interest to us. We first show the existence of both smooth and nonsmooth bifurcations, i.e., classical period doubling bifurcation and border collision bifurcation, as parameters are varied, which lead to queue oscillations at the bottlenecks. We then demonstrate that these instabilities can be induced by varying the parameters only at one of the bottlenecks, while keeping the parameters at the other bottleneck fixed. This suggests that, in some cases of general networks, a queue oscillation induced by one node can spread to other nodes, making it difficult

Manuscript received September 1, 2002; revised February 7, 2004. Recommended by Associate Editor Y. Wardi.

The author is with the Department of Electrical and Computer Engineering, the University of Maryland, College Park, MD 20742 USA (e-mail: hyongla@isr.umd.edu).

Digital Object Identifier 10.1109/TAC.2004.829635

to isolate the source of the instability. Furthermore, the locally sufficient conditions for stability based on single node analysis are shown to be *not* sufficient to guarantee the global network stability. This result suggests that the stability conditions obtained from the single node analysis may not be strong enough to provide a guideline for *independently* setting RED parameters at the gateways in a general network, and some form of coordination may be necessary to ensure global stability.

The rest of the note is organized as follows. Sections II and III present the nonlinear first-order discrete-time model that is used for our analysis. Section IV discusses the local stability of equilibrium points. Section V presents a numerical example based on our analytical model.

II. NONLINEAR FIRST-ORDER DISCRETE-TIME MODEL

We consider a simple network of two links that are shared by many connections, as shown in Fig. 1. As mentioned before, this network can be viewed as a network with two dominant bottleneck links but with other links that are not bottlenecks. We denote the set of connections that traverse both links l_1 and l_2 by $\mathcal{I}_1, \mathcal{I}_1 = \{1, \dots, N_1\}$, and the connections that traverse only the second link l_2 by $\mathcal{I}_2, \mathcal{I}_2 = \{1, \dots, N_2\}$. All connections are assumed to be TCP Reno connections that are explicit congestion notification (ECN) capable [4]. The capacity of links l_1 and l_2 are denoted by C_1 and C_2 , respectively, and the buffer sizes at nodes R_1 and R_2 by B_1 and B_2 , respectively. The packet size of the connections is denoted by M , which can be interpreted as the average packet size of the connections. We assume that the RED queue management mechanism with ECN capability is implemented at nodes R_1 and R_2 in order to control the average queue size at the routers. A RED gateway marks/drops a packet with a probability p , which is a function of the average queue size q^{ave} as follows [3]¹:

$$p(q^{\text{ave}}) = \begin{cases} 0, & \text{if } q^{\text{ave}} < q_{\min} \\ 1, & \text{if } q^{\text{ave}} > q_{\max} \\ \frac{q^{\text{ave}} - q_{\min}}{q_{\max} - q_{\min}} p_{\max}, & \text{otherwise} \end{cases} \quad (1)$$

where q_{\min} and q_{\max} are the lower and higher threshold values, and p_{\max} is the selected marking/drop probability when $q^{\text{ave}} = q_{\max}$. The average queue size is updated at the time of packet arrival according to the exponential averaging

$$q_{\text{new}}^{\text{ave}} = (1 - w)q_{\text{old}}^{\text{ave}} + w \cdot q_{\text{cur}} \quad (2)$$

where q_{cur} is the queue size at the time of arrival, and w is the exponential averaging weight, which determines the time constant of the averaging mechanism and how fast the RED can react to time-varying load.²

On one hand, the averaging weight should be selected small enough so that transient, temporary congestion does not result in an oscillation of the packet marking/drop probability. On the other hand, the weight should be set large enough so that the RED can react to changes in load in a timely manner. These are conflicting goals, and the selection of the parameters affects the interaction of the RED mechanism with adaptive sources, such as TCP.

The connections are assumed to be long-lived, and the set of connections remains fixed for the time period of interest. Although here we consider only long-lived connections, similar behavior has been observed when short-lived connections coexist with long-lived connections [9]. In order to have a tractable model we assume that all connec-

¹In practice, an RED gateway marks/drops a packet with a modified probability in order to lead to a more uniform marking/dropping pattern [3].

²This is a simplification of the formula that does not take into consideration idle periods of the queue.




Fig. 1. Network model.

tions in \mathcal{I}_1 have the same roundtrip propagation delay (RTPD) d_1 and all connections in \mathcal{I}_2 have the same RTPD d_2 . Rather than interpreting this assumption as a requirement that the connections must have the same propagation delay, one should consider the delay d_j as the effective delay that represents the average propagation delay of the connections in \mathcal{I}_j , or this could describe a case where the bottleneck links have large propagation delays that dominate the roundtrip delays of connections, e.g., intercontinental Internet links. This allows us to reduce the problem with $N_1 + N_2$ connections to a two-connection system that represents the set of connections and study its behavior. See [8] for *ns-2* simulation results with heterogeneous roundtrip delays.

Given the roundtrip time (RTT) R and packet marking/drop probability p , the stationary throughput of a TCP Reno connection can be approximated by

$$T(p, R) = \frac{MK}{\sqrt{p}R} + o\left(\frac{1}{\sqrt{p}}\right) \quad (3)$$

where K is some constant in $[1, \sqrt{8/3}]$ [10], [11], when packet losses are detected through triple-duplicate acknowledgment. In this note, we assume $K = \sqrt{3}/2$ [11]. We approximate TCP throughput with $MK/\sqrt{p}R$ to facilitate our analysis. Although we use this simple approximation for TCP throughput to facilitate our analysis, our qualitative results do not depend on this particular form of approximation, and are consequences of rather benign nonlinear behavior of TCP and hold with more sophisticated models for TCP throughput.

TCP adjusts its transmission rate depending on whether it has detected a packet mark/drop or not. Therefore, a network with an AQM mechanism can be viewed as a feedback system, where TCP sources adjust their transmission rates based on the feedback from the AQM mechanism. This feedback is in the form of marked/dropped packets and is delayed by an RTT of the connection. This can be modeled as a stroboscopic map where the instant of observation is approximately one RTT. Since the AQM mechanism should allow enough time for the connections to react to control actions before the control action changes significantly, it is natural to model the system as a discrete-time system. Moreover, our results in [13] illustrate that there exists a natural discrete-time model corresponding to delay-differential equations which arises in studying system stability of a rate control mechanism with a communication delay, and the stability of one can be inferred from that of the other when the delay is large. In other words, the stability of the system given by delay-differential equations can be determined by studying the stability of the underlying discrete-time system. Therefore, our model here can be interpreted as the underlying discrete-time model corresponding to the system given by delay-differential equations (e.g., [7] and [9]).

We use a nonlinear dynamic first-order discrete-time model to analyze the interaction of the RED gateways with TCP connections, which was first proposed by Firoiu and Borden [2]. We define the control system as follows. The packet marking probabilities $\underline{p}_k = (p_k^1, p_k^2)$ at period $k, k \geq 1$, where p_k^j is the marking probability at node $R_j, j = 1, 2$, determine the throughput of the connections and the queue

sizes $\underline{q}_{k+1} = (q_{k+1}^1, q_{k+1}^2)$ at period $k+1$, based on the system constraints. The queue sizes at period $k+1$ are used to compute the average queue sizes $q_{k+1}^{j, \text{ave}}$, $j = 1, 2$, at period $k+1$ according to the exponential averaging rule in (2). Then, the average queue sizes $\underline{q}_{k+1}^{\text{ave}} = (q_{k+1}^{1, \text{ave}}, q_{k+1}^{2, \text{ave}})$ are used to calculate the packet marking probabilities \underline{p}_{k+1} at period $k+1$, which are the control variable of the AQM mechanism. This can be mathematically written as follows:

$$\text{plant function : } \underline{q}_{k+1} = G(\underline{p}_k) \quad (4)$$

$$\text{averaging function : } \underline{q}_{k+1}^{\text{ave}} = A(\underline{q}_k^{\text{ave}}, \underline{q}_{k+1}) \quad (5)$$

$$\text{control function : } \underline{p}_{k+1} = H(\underline{q}_{k+1}^{\text{ave}}) \quad (6)$$

where $A(\underline{q}_k^{\text{ave}}, \underline{q}_{k+1}) = (1 - \underline{w}) \cdot \underline{q}_{k+1}^{\text{ave}} + \underline{w} \cdot \underline{q}_k^{\text{ave}}$ as given in (2), $\underline{w} = (w_1, w_2)$, and the control function $H(\underline{q}_{k+1}^{\text{ave}}) = \underline{p}(\underline{q}_{k+1}^{\text{ave}})$ in (1). These multiplications are component-wise, and the underline is used to denote a vector. The exact form of the plant function $G(\cdot)$ depends on the system parameters such as the number of connections, nature of connections, roundtrip delays, etc., and is described in the following section.

III. INTERACTION OF RED GATEWAYS WITH TCP CONNECTIONS

We now describe the plant function $G(\cdot)$ that is used for analyzing the system. In order to compute the plant function we assume the following: Given the packet drop probability at period k the aggregate throughput of the connections is given by the stationary throughput formula in (3). The use of stationary throughput formula for computing the plant function can be justified as follows. As mentioned earlier the exponential averaging weight should be chosen sufficiently small so that the average queue size q^{ave} does not fluctuate much due to transient, temporary fluctuations in the arrival rate. This implies that the exponential averaging weight should be small enough so that the time constant determined by the averaging weight would be larger than the effective RTT of the connections in order to avoid a fast oscillation in the packet marking/drop probability. When the exponential averaging weight is reasonably small, as it should be, and the number of connections is large with a nonnegligible RTT, the TCP connections' dynamics operate at a faster time-scale than the evolution of the average queue size and the aggregate throughput (or workload) presented by the connections sees a quasi-stationary behavior before the average queue size changes significantly.

We now proceed to define the plant function. Suppose $\underline{p}_k = (p_k^1, p_k^2)$ denote the packet marking probabilities at period k , $k \geq 1$. Let $p_k^3 = 1 - (1 - p_k^1)(1 - p_k^2) = p_k^1 + p_k^2 - p_k^1 \cdot p_k^2$. We define p^u to be

$$p^u = \inf \left\{ p \in [0, 1] \mid \frac{MK}{\sqrt{p}d_1} \leq \frac{C_1}{N_1} \right\} = \left(\frac{N_1 MK}{C_1 d_1} \right)^2 \quad (7)$$

which is the smallest probability such that all $p_k^3 \geq p^u$ lead to $q_{k+1}^1 = 0$ and is assumed to be smaller than one. Based on packet marking probabilities \underline{p}_k we compute the queue sizes $\underline{q}_{k+1} = (q_{k+1}^1, q_{k+1}^2)$ at period $k+1$ by considering different cases as follows.

Case 1)— $p_k^3 \geq p^u$: We define

$$p^{2,u}(\underline{p}_k) = \inf \left\{ p \in \mathfrak{R}_+ \mid \frac{N_2 MK}{\sqrt{p}d_2} \leq C_2 - \frac{N_1 MK}{\sqrt{p_k^3}d_1} \right\}. \quad (8)$$

This is the smallest probability that results in $q_{k+1}^2 = 0$ given $p_k^3 \geq p^u$.

Case 1a)— $p_k^2 \geq p^{2,u}(\underline{p}_k)$: In this case from (7) and (8) one can see that $q_{k+1}^2 = 0$.

Case 1b)— $p_k^2 < p^{2,u}(\underline{p}_k)$: In this case, $q_{k+1}^1 = 0$ and $q_{k+1}^2 = \min(B_2, \tilde{q}^2(\underline{p}_k))$, where $\tilde{q}^2(\underline{p}_k)$ is the solution to

$$\frac{N_1 MK}{\sqrt{p_k^3} \left(d_1 + \frac{q^2 M}{C_2} \right)} + \frac{N_2 MK}{\sqrt{p_k^2} \left(d_2 + \frac{q^2 M}{C_2} \right)} = C_2. \quad (9)$$

If we assume that $\bar{d}_1 = d_2$, then $\tilde{q}^2(\underline{p}_k) = (N_1 K / \sqrt{p_k^3}) + (N_2 K / \sqrt{p_k^2}) - (C_2 d_2 / M)$.

Case 2)— $p_k^3 < p^u$: There exist a set of $(q^1, q^2) \geq \underline{0}$ such that

$$\frac{N_1 MK}{\sqrt{p_k^3} \left(d_1 + \frac{q^1 M}{C_1} + \frac{q^2 M}{C_2} \right)} = C_1. \quad (10)$$

Let $q_k^{2, \text{max}} = (N_1 K C_2 / \sqrt{p_k^3} C_1) - (d_1 C_2 / M)$, i.e., q^2 that satisfies (10) with $q^1 = 0$, and $q_k^{2u} = \max(0, (N_2 K C_2 / (\sqrt{p_k^2} (C_2 - C_1))) - (d_2 C_2 / M))$, where the second term is the solution to $(N_2 MK / (\sqrt{p_k^2} (d_2 + (q^2 M / C_2)))) = C_2 - C_1$.

Case 2a)— $q_k^{2u} \leq \min(B_2, q_k^{2, \text{max}})$: In this case, the queue size $q_{k+1}^2 = q_k^{2u}$ and $q_{k+1}^1 = \min(B_1, \tilde{q}^1(q_{k+1}^2, \underline{p}_k))$, where $\tilde{q}^1(q_{k+1}^2, \underline{p}_k)$ is q^1 that satisfies (10) with $q^2 = q_{k+1}^2 = q_k^{2u}$, i.e., $\tilde{q}^1(q_{k+1}^2, \underline{p}_k) = (N_1 K / \sqrt{p_k^3}) - (d_1 C_1 / M) - (q_{k+1}^2 C_1 / C_2) = (N_1 K / \sqrt{p_k^3}) - (d_1 C_1 / M) - (q_k^{2u} C_1 / C_2)$.

Case 2b)— $q_k^{2u} > \min(B_2, q_k^{2, \text{max}})$: In this case the queue size $q_{k+1}^2 = \min(B_2, \tilde{q}^2(\underline{p}_k))$ and $q_{k+1}^1 = \max(0, \min(B_1, \tilde{q}^1(q_{k+1}^2, \underline{p}_k)))$, where $\tilde{q}^2(\underline{p}_k)$ is the solution to (9), and $\tilde{q}^1(q_{k+1}^2, \underline{p}_k) = (N_1 K / \sqrt{p_k^3}) - (C_1 d_1 / M) - (C_1 q_{k+1}^2 / C_2)$ as in the previous case.

From the aforementioned information, we can compute \underline{q}_{k+1} as a function of \underline{p}_k , and hence we have the plant function shown in (11) and (12) at the bottom of page to complete our discrete-time model described in Section II. From (4) to (6) and the plant function $G(\cdot)$ given in (11) and (12), we define a mapping

$$\underline{q}_{k+1}^{\text{ave}} = (1 - \underline{w}) \cdot \underline{q}_k^{\text{ave}} + \underline{w} \cdot A G \left(H \left(\underline{q}_k^{\text{ave}} \right) \right) := g \left(\underline{q}_k^{\text{ave}}, \underline{\rho} \right) \quad (13)$$

where $\underline{\rho}$ summarizes the system parameters, including the exponential averaging weights \underline{w} , and multiplications are component-wise. Equation (13) maps the average queue sizes at period k to the average queue sizes at period $k+1$.

IV. STABILITY OF THE SYSTEM

A fixed point of the mapping $g(\cdot, \cdot)$ is a vector of average queue sizes \underline{q}^* such that $\underline{q}^* = g(\underline{q}^*, \underline{\rho})$. Given the system parameters one can find such a fixed point \underline{q}^* , if there is any, by solving (13) with the given $A(\cdot)$, $G(\cdot)$, and $H(\cdot)$. The local stability of the fixed point \underline{q}^* can be studied by looking at the eigenvalues of the Jacobian matrix $J = [J_{ij}, i, j = 1, 2]$ of the mapping $g(\cdot, \cdot)$ in (13), where $J_{ij} = (\partial g_i / \partial q^j)$. Depending on where the fixed point $\underline{q}^* = (q^{1*}, q^{2*})$ of the system lies, there could be several different cases. In this section we only consider the case where \underline{q}^* satisfies case (2a) in Section III,

which is the most interesting case from the operational point of view. In this case, $G_1(\underline{p})$ and $G_2(\underline{p})$, where $\underline{p} = H(\underline{q}^*)$, are given by

$$\begin{aligned} G_1(\underline{p}) &= \min \left(B_1, \frac{N_1 K}{\sqrt{p_k^3}} - \frac{N_2 K C_1}{\sqrt{p_k^2} (C_2 - C_1)} + \frac{(d_2 - d_1) C_1}{M} \right) \\ G_2(\underline{p}) &= \max \left(0, \frac{N_2 K C_2}{\sqrt{p_k^2} (C_2 - C_1)} - \frac{d_2 C_2}{M} \right) \end{aligned} \quad (14)$$

where $p^1 = ((q^{1*} - q_{\min}^1)/(q_{\max}^1 - q_{\min}^1))p_{\max}^1$, $p^2 = ((q^{2*} - q_{\min}^2)/(q_{\max}^2 - q_{\min}^2))p_{\max}^2$, and $p^3 = p^1 + p^2 - p^1 \cdot p^2$. We assume that the fixed point is an interior point, i.e., $0 < \underline{q}^* < \underline{B}$.

From (14), one can see that $g_2(\cdot, \rho)$ does not depend on q^{1*} and, hence, $J_{21} = 0$. Therefore, the eigenvalues of the Jacobian matrix are given by J_{11} and J_{22} , where

$$\begin{aligned} J_{11} &= 1 - w_1 - w_1 \cdot \frac{N_1 K (1 - p^2) \alpha^1}{2(p^1 + p^2 - p^1 \cdot p^2)^{1.5}} \\ J_{22} &= 1 - w_2 - w_2 \cdot \frac{N_2 K C_2}{2\sqrt{p^2} (q^{2*} - q_{\min}^2) (C_2 - C_1)} \end{aligned}$$

and $\alpha^j = p_{\max}^j / (q_{\max}^j - q_{\min}^j)$, $j = 1, 2$. We investigate how these eigenvalues behave as the exponential averaging weights vary in the following section.

V. NUMERICAL RESULTS

A bifurcation diagram shows the qualitative changes in the nature or the number of fixed points of a dynamical system with varying parameters. In this section we only vary the exponential averaging weights and study the stability of the system. However, similar results can be obtained with any of the system or control parameters (see [12]). The x -axis is the parameter that is being varied, and the y -axis plots the set of fixed solutions (with a period of one or higher) corresponding to the value of the parameter. For generating the bifurcation diagrams, in each run we randomly select four random initial average queue sizes, $\underline{q}_1^{\text{ave}}(0)$, $\underline{q}_2^{\text{ave}}(0)$, $\underline{q}_3^{\text{ave}}(0)$, and $\underline{q}_4^{\text{ave}}(0)$, and these average queue sizes evolve according to the map $g(\cdot, \cdot)$ in (13), i.e.,

$$\begin{aligned} \underline{q}_i^{\text{ave}}(k) &= g(\underline{q}_i^{\text{ave}}(k-1), \underline{\rho}), \\ \text{for } k &= 1, \dots, 1,000 \text{ and } i = 1, 2, 3, \text{ and } 4. \end{aligned}$$

We plot $\underline{q}_i^{\text{ave}}(k)$, $k = 991, \dots, 1,000$ and $i = 1, 2, 3, 4$. Hence, if there is a single stable fixed point or attractor \underline{q}^* of the system at some value of the parameter, all $\underline{q}_i^{\text{ave}}(k)$ will converge to \underline{q}^* and there will be only one point along the vertical line at the value of the parameter. However, if there are two stable fixed points, $\underline{q}_1^{\text{ave}}$ and $\underline{q}_2^{\text{ave}}$, with a period of two, i.e., $g(\underline{q}_i^{\text{ave}}, \underline{\rho}) \neq \underline{q}_i^{\text{ave}}$ and $g(g(\underline{q}_i^{\text{ave}}, \underline{\rho})) = \underline{q}_i^{\text{ave}}$, $i = 1, 2$, then there will be two points along the vertical line and the average queue sizes will alternate between $\underline{q}_1^{\text{ave}}$ and $\underline{q}_2^{\text{ave}}$.

The parameters used in the numerical examples are as follows:

$$\begin{aligned} \underline{q}_{\max} &= (600, 1200), \quad \underline{q}_{\min} = (200, 400), \quad \underline{p}_{\max} = \left(\frac{1}{8}, \frac{1}{8}\right) \\ \underline{C} &= (12, 30) \text{ Mbps}, \quad K = \sqrt{\frac{3}{2}}, \quad \underline{B} = (1000, 2000) \\ d_1 = d_2 &= 0.1 \text{ s}, \quad M = 4,000 \text{ bits}, \quad N_1 = N_2 = 100. \end{aligned}$$

A. Instability of the Tandem Network

Fig. 2(a) and (b) plot the set of stable fixed points as a function of exponential averaging weights $\underline{w} = w \cdot (1, 1)$, where w is the value along the x -axis. The y -axis plots the average queue sizes per flow, i.e., $q^{1,\text{ave}}/N_1$ and $q^{2,\text{ave}}/(N_1 + N_2)$. Fig. 2(c) and (d) show the actual queue sizes per flow. One can see that there is a unique fixed point of the system $\underline{q}^* = (2.12 \cdot N_1, 2.78 \cdot (N_1 + N_2))$ for $w < 0.3862$. At $w = 0.3862$, the initial period doubling bifurcation occurs, and the queue size begins to oscillate. This can be also verified by computing the eigenvalue J_{22} and showing that it hits the unit circle at -1 [6].

A careful investigation of the bifurcation diagrams reveals that the second period doubling bifurcation at $w = 0.464$ is not a classical period doubling bifurcation, but rather a period doubling bifurcation due to a border collision. Here, a border is a region at which the behavior of queue sizes changes from one case to another in Section III. If a fixed point collides with a border, there is a discontinuous change in the Jacobian matrix of the mapping. This border collision bifurcation leads to a cascade of bifurcations, resulting in chaos as shown in Fig. 2.

This border collision bifurcation illustrates the impact of an instability on the system throughput. Note that the distance between the initial period doubling bifurcation point and the border collision bifurcation point is relatively short. Hence, once the system becomes unstable, the queue sizes quickly start oscillating widely, often leading to empty queues and a waste of resources.

$$\begin{aligned} G_1(\underline{p}_k) &= q_{k+1}^1 \\ &= \begin{cases} 0, & \text{if } p_k^3 \geq p^u \\ \min \left(B_1, \frac{N_1 K}{\sqrt{p_k^3}} - \frac{d_1 C_1}{M} - \frac{q_k^{2u} C_1}{C_2} \right), & \text{if } p_k^3 < p^u \text{ and } q_k^{2u} \leq \min(B_2, q_k^{2,\text{max}}) \\ \max \left(0, \min \left(B_1, \tilde{q}^1 \left(\min \left(B_2, \tilde{q}^2(\underline{p}_k) \right), \underline{p}_k \right) \right) \right), & \text{otherwise} \end{cases} \end{aligned} \quad (11)$$

and

$$\begin{aligned} G_2(\underline{p}_k) &= q_{k+1}^2 \\ &= \begin{cases} 0, & \text{if } p_k^3 \geq p^u \text{ and } p_k^2 \geq p^{2,u}(\underline{p}_k) \\ \min \left(B_2, \tilde{q}^2(\underline{p}_k) \right), & \text{if } p_k^3 \geq p^u \text{ and } p_k^2 < p^{2,u}(\underline{p}_k) \\ q_k^{2u}, & \text{if } p_k^3 < p^u \text{ and } q_k^{2u} \leq \min(B_2, q_k^{2,\text{max}}) \\ \min \left(B_2, \tilde{q}^2(\underline{p}_k) \right), & \text{otherwise} \end{cases} \end{aligned} \quad (12)$$




Fig. 2. Bifurcation diagrams. (a) $q^{1,ave}$, (b) $q^{2,ave}$, (c) q^1 , and (d) q^2 .

B. Propagation of Instability

In this section, we only vary the exponential averaging weight of the second bottleneck link l_2 . Fig. 3 plots the average queue sizes per flow, i.e., $q^{1,ave}/N_1$ and $q^{2,ave}/(N_1 + N_2)$, as a function of $w = (0.1, w)$. One can see that, as in Fig. 2, there is a unique fixed point of the system $q^* = (2.12 \cdot N_1, 2.78 \cdot (N_1 + N_2))$ for $w < 0.3862$. At $w = 0.3862$ the initial period doubling bifurcation takes place as in the previous subsection. Hence, although w_1 is much smaller in this case (w_1 is fixed at 0.1), the initial bifurcation occurs at the same value of w as in Fig. 2. In fact the border collision bifurcation, which leads to chaos, occurs at the same value as well, as shown in Figs. 2 and 3. This can be explained as follows. In this example, the fixed point q^* satisfies case (2a), and the eigenvalue of the Jacobian matrix with larger magnitude for the given range of w in this example is J_{22} , which does not depend on q^{1*} because $g_2(\cdot, \rho)$ is a function only of $q^{2,ave}$ as shown in Section IV.

This example demonstrates that in a general network an instability caused by one bottleneck may propagate to other parts of the network. Hence, even if the majority of the routers are properly configured, if a handful of routers are misconfigured, then a large portion of the network may experience an oscillation in queue sizes, which may make it difficult to isolate the source of instability and correct it.

This numerical example also illustrates the fact that locally sufficient conditions for stability are not enough to guarantee the global network stability as will be shown here. Note from [12] that the initial period doubling bifurcation point in a single bottleneck case is given by

$$w^* = \frac{2}{1 + \frac{NK}{2\sqrt{\alpha}(q^* - q_{\min})^{1.5}}} \quad (15)$$

where $\alpha = p_{\max}/(q_{\max} - q_{\min})$, and N is the number of connections. The fixed point q^* can be computed as the positive real solution to the following third-order polynomial:

$$(q^* - q_{\min}) \left(q^* + \frac{dC}{M} \right)^2 - \frac{(NK)^2}{\alpha} = 0.$$

Suppose that we isolate each bottleneck link and assume that the other link is not a bottleneck. In other words, we select one bottleneck link in the tandem network at a time and remove the other bottleneck and the connections not traversing the selected bottleneck link. This allows us to study how congestion at one bottleneck affects the queue behavior




Fig. 3. Bifurcation diagrams. (a) $q^{1,ave}$. (b) $q^{2,ave}$.

at another and the stability of the overall network. Equation (15) yields the bifurcation points of 0.5674 and 0.4692 for the first and second bottleneck links, respectively, when they are considered separately. Therefore, from Figs. 2 and 3 one can see that the locally sufficient conditions that $w_1 < 0.5674$ and $w_2 < 0.4692$ based on the single node analysis, are not sufficient to guarantee the global network stability. The $ns-2$ simulation results for these numerical examples are given in [8].

This has rather serious practical implications. First, if the instability propagates from one bottleneck to another, it is difficult to track down its source since every bottleneck that exhibits an instability needs to be checked. Moreover, even if every bottleneck is checked, it may still be difficult to locate the source. Second, even when a network manager can somehow estimate all system parameters and satisfy the locally sufficient conditions for stability based on the single node analysis [9], [12], these locally sufficient conditions at individual nodes do not guarantee the global stability of the network, and some of the nodes may still experience an oscillation in queue size. These issues on parameter selection pose a serious challenge to network engineers since different sets of bottlenecks connections traverse (determined by the routing algorithms) may belong to different domains, and ensuring the global network stability may require communication and cooperation between these domains.

REFERENCES

- [1] M. Christiansen, K. Jeffay, D. Ott, and F. D. Smith, "Tuning RED for web traffic," *IEEE/ACM Trans. Networking*, vol. 9, pp. 249–264, June 2001.
- [2] V. Firoiu and M. Borden, "A study of active queue management for congestion control," presented at the INFOCOM 2000, Tel Aviv, Israel, Mar. 2000.
- [3] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE Trans. Networking*, vol. 1, pp. 397–413, July 1993.
- [4] S. Floyd, "TCP and explicit congestion notification," *ACM Comput. Commun. Rev.*, vol. 24, pp. 10–23, Oct. 1994.
- [5] S. Floyd and E. Kohler, "Internet research needs better models." [Online] <http://www.icir.org/floyd>
- [6] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, 1st ed. New York: Springer-Verlag, 1983.
- [7] C. V. Hollot, V. Misra, D. Towsley, and W. Gong, "A control theoretic analysis of RED," in *Proc. IEEE INFOCOM 2001*.

- [8] R. J. La, "Instability of a tandem network and its propagation under RED." [Online] <http://www.ece.umd.edu/~hyongla>
- [9] S. H. Low, F. Paganini, J. Wang, S. Adlakha, and J. C. Doyle, "Dynamics of TCP/RED and a scalable control," in *Proc. IEEE INFOCOM 2002*.
- [10] M. Mathis, J. Semke, J. Mahdavi, and T. Ott, "The macroscopic behavior of the TCP congestion avoidance algorithm," *Comput. Commun. Rev.*, vol. 27, no. 3, 1997.
- [11] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP Reno performance: a simple model and its empirical validation," *IEEE/ACM Trans. Networking*, vol. 8, pp. 133–145, Apr. 2000.
- [12] P. Ranjan, E. H. Abed, and R. J. La, "Nonlinear instabilities in TCP-RED," in *Proc. IEEE INFOCOM 2002*.
- [13] P. Ranjan, R. J. La, and E. H. Abed, "Tradeoffs in rate control with communication delay," *IEEE/ACM Trans. Networking*, Oct. 2004, to be published.