Equilibrium Properties of Rate Control Protocols

Abhijit Kiran Valluri

University of Maryland, College Park, MD, USA

Email: abhijitv@umd.edu

Abstract—We analyze the stability of the Rate Control Protocol (RCP) using two different models that have been proposed in literature. Our objective is to better understand the impact of the protocol parameters and the effect different forms of feedback have on the stability of the network. We also highlight that different time scales, depending on the propagation delay relative to the queuing delay, have an impact on the nonlinear and the stochastic properties of the protocol fluid models. To better understand some of the nonlinear properties, we resort to local bifurcation analysis where we exhibit the existence of a Hopf type bifurcation that then leads to stable limit cycles. Our work serves as a step towards a more comprehensive understanding of the nonlinear fluid models that have been used as representative models for RCP.

I. Introduction

Congestion in the Internet is just as critical now as ever before, with the proliferation of high speed mobile and broadband internet and the large increase in demand for several high bandwidth applications such as multimedia streaming and real-time communications. Understandably, the Internet community is actively working to develop congestion control for such applications, through the "RTP Media Congestion Avoidance Techniques" (RM-CAT) working group[1].

Real-time applications perform poorly in the presence of TCP cross-traffic, due to the inherent fairness issues that arise from TCP's Additive Increase Multiplicative Decrease (AIMD) congestion control algorithms. Numerous studies have exhibited that standard TCP AIMD protocols [1], [2] are unsuitable for next generation networks, and repeated attempts have been made to improve it. Explicit Congestion control Protocols (ECPs) such as [3] aim to provide better fairness. The need for more explicit feedback is well recognized which in turn motivates the requirement for a comprehensive theoretical framework within which to design transport protocols. Therefore, in this paper, we reconsider the Rate Control Protocol (RCP) [4] that has been proposed earlier and study its stability properties.

RCP has received significant attention from the research community: [5] develops some stability properties of a max-min RCP in small buffer regime, [6] computationally develops some sufficient conditions for local stability, [7] considers a dynamic environment with RCP flows arriving and departing over a single link, [8] develops an α-fair variant of RCP and investigates some of the associated local stability properties, [9] develops a NetFPGA hardware implementation of RCP, [10] develops a new congestion controller called Proportional Integral Queue Independent RCP, [11] performs local bifurcation analysis for some congestion control algorithms, and [12] develops an experimental framework used to evaluate several ECPs. The range of these studies exhibit the difficulty of developing a new transport protocol.

RCP routers obtain rate estimates from two forms of feedback: one is based on rate mismatch and the other is from the queue size. Thus far the role of both forms of feedback has not been well understood. [6] provides some understanding of the role of queuing dynamics on the local stability of explicit congestion control algorithms. Nevertheless, there are still no clear guidelines on choosing vital protocol parameters that influence stability and link utilization.

In this paper, we study two nonlinear fluid models of RCP and analyze the role played by queue feedback in RCP performance. We claim, using packet level simulations from one of our collaborators in [8], that non-switched system is more appropriate and in this regime we develop necessary and sufficient conditions for local stability. We also computationally study the inherent nonlinear aspects of the protocol. We then study another model for RCP, where the queue is not modeled as a separate fluid quantity, but is a deterministic representation for the underlying stochastics. In this model, we further develop the understanding of RCP by considering local stability and a local bifurcation analysis. This is done by varying the network parameters, which impact the stability and the link utilization.

The rest of the paper is structured as follows. In Section II, we outline and analyze two models for RCP. In Section III, we summarize our contributions and conclude the paper.

II. Stability Analysis of two RCP models

In this section, we outline two nonlinear dynamical systems models for the RCP [6], [8]. These models have previously been motivated with the objective to help design and better understand the performance of RCP.

In the operation of RCP, the feedback from the routers to the end-systems is time-delayed which makes it important to understand the stability properties of the nonlinear models. Here, we develop both necessary and sufficient conditions for stability of RCP that are not previously shown in literature, and explore the consequences of such local stability conditions being violated, both analytically and numerically, where bifurcation phenomena may readily occur.

A. Model A

The protocol strives to estimate the fair rate through a single bottleneck link from the rate mismatch and the queue size.

---

[1]URL: https://tools.ietf.org/wg/rmcat/

In order to understand the performance of the protocol, the following nonlinear dynamical system for the rate and the queue has been proposed [6], [13]:

$$\frac{d}{dt} R(t) = \frac{R(t)}{C\overline{T}} \left(a(C - y(t)) - \beta \frac{q(t)}{\overline{T}} \right) \qquad (1)$$

where

$$y(t) = \sum_s R(t - T_s) \qquad (2)$$

and

$$\frac{d}{dt} q(t) = [y(t) - C] \qquad q(t) > 0$$
$$= [y(t) - C]^+ \qquad q(t) = 0, \qquad (3)$$

using the notation $x^+ = \max(0, x)$. Here $R(t)$ is the rate being updated by the router, C is the link capacity, $y(t)$ is aggregate load at the link, $q(t)$ is the queue size, T_s is the round trip time (RTT) of flow s, and $\overline{T}$ is the average round trip time, over the flows present. In the formulation of the RCP equation (1), a and β are non-negative dimensionless parameters. It is important to understand the impact that these parameters would have on the performance of the protocol.

The nonlinear rate equation (1) utilizes two forms of feedback: one for the rate mismatch which is characterized by $C - y(t)$, and another for the instantaneous queue size, $q(t)$. The rate mismatch term causes the rate to increase if the utilization is lower than the link capacity C and the queue feedback term serves to decrease the feedback rate as the queue size in the router starts to build up.

Some sufficient conditions for the local stability of the system (1), (2), (3), about its equilibrium point, were derived in [6], using techniques developed for a "switched" linear control system with a time delay. The analysis takes into account the discontinuity in the system dynamics which would occur as the queue size approaches zero. However, we note that this analysis applies to the fluid model rather than a packet level description of the protocol.

The sufficient conditions, on the non-negative and dimensionless parameters a and β, take the functional form

$$a < \frac{\pi}{2} \qquad (4)$$

and $\beta < f(a)$ where $f(\cdot)$ is a positive function that depends on $\overline{T}$.

We refer to some packet level simulations performed for RCP using a discrete event simulator, reported in [8, Fig. 5]. These simulations demonstrate that the queue, in equilibrium, may not sit at zero but rather the mean queue size would be close to it. The authors consider a network consisting of a single bottleneck link with a capacity C of one packet per unit time and 100 Poisson sources and RTT 100 time units. The parameters with queue feedback are: $a = 0.5$, $\beta = 1$ for a utilization of 90%; the parameters without queue feedback are: $a = 1$, $\beta = 0$, with capacity γC, where $\gamma = 0.9$ to target a utilization of 90% as before. It was observed that with queue feedback the queue size is unstable and has oscillations. On the other hand, without queue feedback, the queue size term is stable and has a non zero equilibrium value.

This motivates our subsequent stability analysis for necessary and sufficient conditions. Clearly, the two forms of feedback are playing a non-trivial role which would not be apparent from a linear system. Due to this, we also need to understand the protocol behavior when conditions for stability are violated.

1) Stability analysis: We shall now understand the local stability of Model A. Since we wish to focus on the nonlinearity of the model, we consider the following modified equation for the queue dynamics, instead of (3)

$$\dot{q}(t) = [y(t) - C] \qquad \forall q(t). \qquad (5)$$

Assume that all flows through the bottleneck link have a common RTT, T. Our system is now represented by (1), (2), (5). The fixed points for these equations are:

$$R^\star = C/n, \qquad q^\star = 0, \qquad (6)$$

where n is the number of flows. Upon linearizing (1), (2), (5) about the fixed point and on further simplification, we get

$$\dot{r}(t) = -\frac{a}{T}(r(t - T)) - \frac{\beta}{nT^2} q(t)$$
$$\dot{q}(t) = nr(t - T), \qquad (7)$$

with $r(t) = R(t) - C/n$. Using a Laplace transform, we get

$$S^2 T^2 e^{TS} + aTS + \beta = 0, \qquad (8)$$

where S is the complex argument in the frequency domain. We consider two cases: when $\beta = 0$ and $\beta > 0$ corresponding to queue feedback being absent and present, respectively.

a) Without queue feedback: Here, (8) becomes

$$STe^{TS} + a = 0. \qquad (9)$$

We let $TS = \lambda$, and introduce a new parameter, η, in (9):

$$\lambda e^{\lambda \eta} + a = 0. \qquad (10)$$

Substituting $\lambda = j\omega$, and comparing real and imaginary parts above, we get $\omega \eta = (2m + 1)\pi/2$ and $\omega = a$. Hence, $a = \pi/(2\eta)$ when $m = 1$. We also note that when $\eta = 0$, the only root of (10) is $\lambda = -a$ that is stable. Now, using Rouché's theorem [14], we find that the system represented by (10) is stable if $\eta < \pi/(2a)$. With $\eta = 1$, we get back our original characteristic equation, and the stability condition becomes

$$a < \frac{\pi}{2}. \qquad (11)$$

This is a necessary and sufficient condition for local stability for the system (1), (2), (5) when $\beta = 0$, which implies that there is no queue feedback.
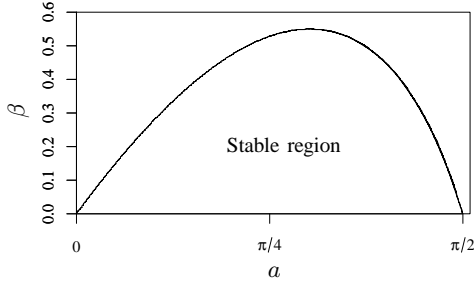
Fig. 1. Stability chart for Model A

b) With queue feedback: Here, (8) becomes ($TS = \lambda$):

$$\lambda^2 e^{\lambda} + a\lambda + \beta = 0. \tag{12}$$

Once again, we introduce a new parameter η as follows:

$$\lambda^2 e^{\eta\lambda} + a\lambda + \beta = 0. \tag{13}$$

By substituting $\lambda = j\omega$ in the above equation we get:

$$-\omega^2 \cos(\eta\omega) + \beta = 0, \text{ and } -\omega^2 \sin(\eta\omega) + a\omega = 0. \tag{14}$$

Upon simplifying the above equations, we get:

$$\omega = \pm\sqrt{\frac{a^2 + \sqrt{a^4 + 4\beta^2}}{2}}, \tag{15}$$

as ω^2 is non-negative. $\omega \sin(\eta\omega) = a$ as $\omega \neq 0$ if $a > 0$, so

$$\eta = \frac{1}{\omega} \sin^{-1}\left(\frac{a}{\omega}\right). \tag{16}$$

Now, if a and β are fixed and $\eta = 0$, the roots of (13) have negative real parts, as a is positive, and are stable. Once again, using Rouché's theorem, and with $\eta = 1$, we get

$$1 < \frac{1}{\omega} \sin^{-1}\left(\frac{a}{\omega}\right) \tag{17}$$

as the stability condition for the linearized system corresponding to equation (12). This can be simplified as

$$\tan\frac{\sqrt{a^2 + \sqrt{a^4 + 4\beta^2}}}{\sqrt{2}} < \frac{a}{\beta}\frac{\sqrt{a^2 + \sqrt{a^4 + 4\beta^2}}}{\sqrt{2}}, \tag{18}$$

which is the necessary and sufficient condition for the local stability for the system (1), (2), (5) when $\beta > 0$; see Fig. 1.

2) Numerical results: We performed numerical computations for Model A and presented the results in Fig. 2. We chose $\beta = 0.3$, RTT $T = 0.1$ time units, capacity $C = 100,000$ packets per unit time, number of flows $s = 100$, and varied the parameter a. We plotted the bifurcation diagrams and the phase portraits, both with and without the switch in the queue dynamics. The results confirm our analytically derived stability region, with the system becoming unstable when a leaves the stability region. Additionally, the results show the nonlinear nature of the instability in the form of limit cycles, which are visually apparent in the phase portraits. The cycle for $a = 1.1$,

in the non-switched case, and $a = 1$, in the switched case, are so small that they had to be represented using a square box.

Note that the dynamics of Model A undergo a bifurcation at the boundary of the stability region. The analysis in Section-IIA1 is done for the system (1), (2), (5). We have neglected the switching in the dynamics of the queue size, using (5), to highlight the nonlinearity in the dynamics of the rate term. Nevertheless, the actual behavior of RCP is better represented by the system (1), (2), (5). [8, Fig. 5] shows that the mean queue size is not zero. Hence, the switch in the queue size dynamics does not play a vital role, since the equilibrium queue size is non-zero. Furthermore, we observe from Fig. 2 that the dynamics of the rate term undergoes a bifurcation with respect to the parameter a, both with and without switching in the queue dynamics. The phase portraits show that both cases are topologically equivalent as we can obtain one from the other by distorting the phase portrait appropriately. It is, therefore, reasonable to analyze Model A without the switch in the queue size dynamics.

We now present another model used to represent RCP.

B. Model B

Model A accounts for the queue term explicitly via the differential equations (3), (5). We now outline a small buffer model of RCP. In this regime, the queue size fluctuates so rapidly that it becomes impossible to respond to and control its actual size. Instead, RCP behaves as if it is responding to a *distribution* of the queue size. Therefore, at the time scale pertinent for the convergence of the system, the *mean* queue size is more important. It is also assumed that the queuing delay is negligible compared to the propagation delay, which conforms with the small buffer assumption.

A small queue variant of RCP that is proportionally fair is described by the following nonlinear differential equations [8]

$$\frac{d}{dt}R_j(t) = \frac{aR_j(t)}{C_j\overline{T}_j(t)}\left(C_j - y_j(t) - b_jC_jp_j(y_j(t))\right) \tag{19}$$

where

$$y_j(t) = \sum_{r:j\in r} x_r(t - T_{rj}) \tag{20}$$

is the aggregate load at resource j summed over all the routes, r, containing resource j; $x_r(t)$ is the flow rate leaving the source of route r; $p_j(y_j)$ is the mean queue size at link j when the load there is y_j; capacity of resource j is C_j; and

$$\overline{T}_j(t) = \frac{\sum_{r:j\in r} x_r(t)T_r}{\sum_{r:j\in r} x_r(t)} \tag{21}$$

is the average RTT of packets passing through resource j. Here, T_r is the sum of the propagation delay from the source of the flow on route r to resource j (T_{rj}) and the propagation delay from the resource j to the source of the flow on route r (T_{jr}). In (19), a and b_j are non-negative dimensionless parameters. Let the flow rate $x_r(t)$ leaving the source of route r at time t be given by
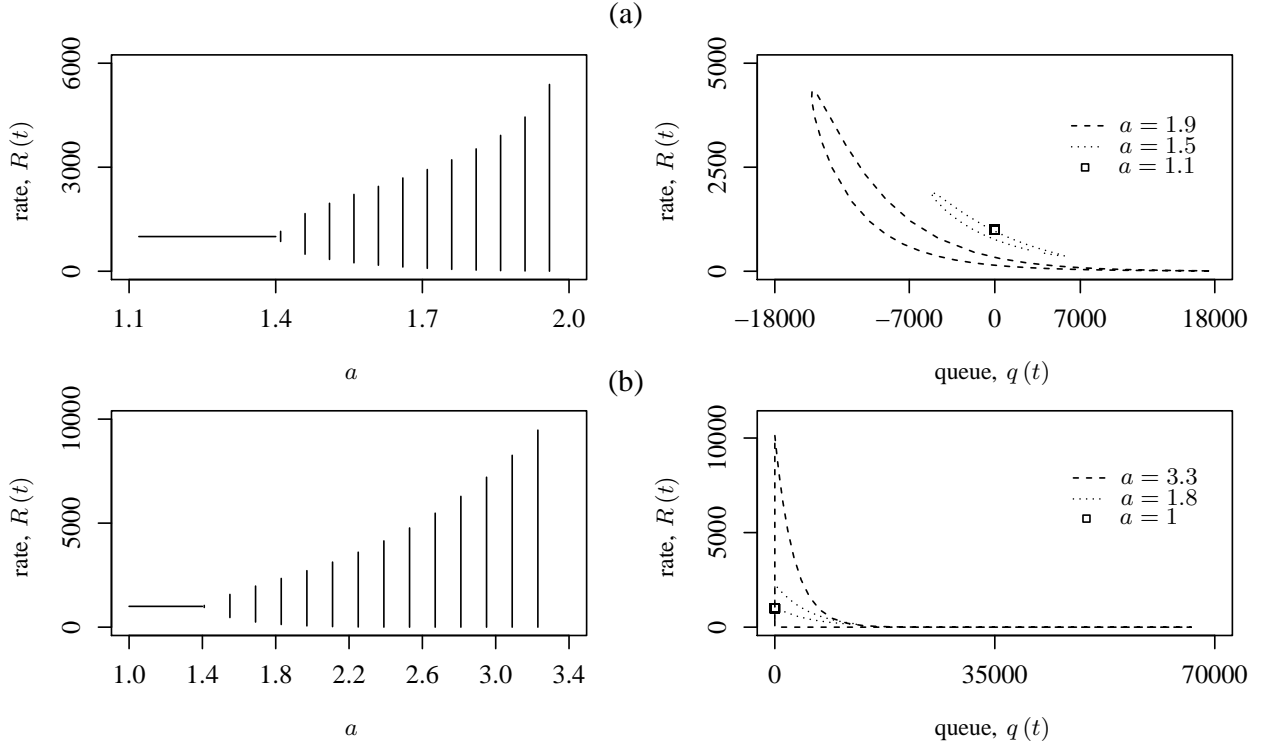
Fig. 2. Numerical computations for Model A: (a) Bifurcation diagram (left) and phase portrait (right) for the "non-switched" case using (5), with $\beta = 0.3$ and a is varied, (b) Bifurcation diagram (left) and phase portrait (right) for the "switched" case using (3), with $\beta = 0.3$ and a is varied. Other parameter values are: RTT, $T = 0.1$ time units; Capacity, $C = 100,000$ packets per unit time; Number of flows, $s = 100$.

$$x_r(t) = w_r \left(\sum_{j \in r} R_j (t - T_{jr})^{-1} \right)^{-1}, \qquad (22)$$

where w_r is the weight given to route r. We can obtain an expression for the mean queue size in the following way: consider the arriving workload at resource j is Gaussian over a time period τ, with mean $y_j \tau$ and variance $y_j \tau \sigma_j^2$. Then the workload present at the queue is a reflected Brownian motion, with mean under its stationary distribution of

$$p_j(y_j) = \frac{y_j \sigma_j^2}{2(C_j - y_j)}. \qquad (23)$$

In essence, the queue size term is being modeled by $p_j(y_j)$ as described in (23). The parameter σ_j determines how variable the traffic at resource j is. For instance, if $\sigma_j = 1$, then the traffic is Poisson.

We should note that the parameter a in both the models is the same. The parameter b_j of Model B can be related to the parameters a and β of Model A by the equation

$$b_j = \frac{\beta}{aC_j \overline{\overline{T}}_j}. \qquad (24)$$

The parameter b_j affects the utilization of resource j at equilibrium. From (23) and considering the condition for equilibrium of the system in (19), (20), we can evaluate the utilization of resource j, ρ_j, after simplification, as

$$\rho_j = 1 - \sigma_j \sqrt{\left(\frac{b_j}{2} \right)} + O\left(\sigma_j^2 b_j \right). \qquad (25)$$

1) Stability analysis: In [8], a sufficient condition for the local stability of Model B for heterogeneous propagation delays was derived. It was shown that with queue feedback, a sufficient condition for stability is $a < \pi/2$, while without queue feedback, it is $a < \pi/4$.

Here, we present the necessary and sufficient conditions for the local stability of this model, for a homogeneous propagation delay, both with and without queue feedback. We note that we get similar sufficient conditions as mentioned in [8]. We also prove, analytically, the existence of a Hopf bifurcation at the edge of the stability region in both cases.

a) With queue feedback: Let the network consist of a single link with capacity C, a single route and a common RTT, τ, for all the flows. We drop the subscripts, j and r, for clarity. All the flows send Poisson traffic, hence, $\sigma = 1$. We take $w = 1$ as this only affects the equilibrium point. For this scenario, the general rate equation in (19) becomes

$$\frac{d}{dt} R(t) = \frac{aR(t)}{C\tau} (C - y(t) - bCp(y(t))) \qquad (26)$$

where

$$y(t) = R(t - \tau), \text{ and } p(y) = \frac{y}{2(C - y)}. \quad (27)$$

We now linearize (26) using the Taylor expansion about the equilibrium. Let $R(t) = r(t) + R^\star$, where $R^\star$ is the equilibrium value of $R(t)$, and $r(t)$ is a small perturbation about the equilibrium. Hence, (26) becomes

$$\frac{d}{dt}R(t) = \dot{r}(t) = -a\left(\frac{R^\star + C}{C\tau}\right)r(t - \tau), \quad (28)$$

where $R^\star = C\left(b + 4 - \sqrt{b^2 + 8b}\right)/4$. More simply,

$$\dot{r}(t) = -\kappa r(t - \tau), \quad (29)$$

where $\kappa = a\left(2 + b/4 - \sqrt{b^2/16 + b/2}\right)/\tau$. By taking the Laplace transform of (29), we get the characteristic equation

$$\lambda + \kappa e^{-\lambda\tau} = 0, \quad (30)$$

where λ is the complex argument in the frequency domain. Once again, we introduce a new parameter η as,

$$\lambda + \kappa e^{-\lambda\tau\eta} = 0. \quad (31)$$

By substituting $\lambda = j\omega$ and comparing real and imaginary parts, we get $\kappa\tau\eta = \pi/2$. We note that when $\eta = 0$, the only solution of (31) is $\lambda = -\kappa$, which is stable. Hence, using Rouché's theorem, and substituting $\eta = 1$ to get the characteristic equation (30), we note that $\kappa\tau < \pi/2$ is the necessary and sufficient condition for the local stability of Model B, which is expanded as

$$a\left(2 + \frac{b}{4} - \sqrt{\frac{b^2}{16} + \frac{b}{2}}\right) < \frac{\pi}{2}. \quad (32)$$

If we substitute $b \to 0$ in (32), we get

$$a < \pi/4. \quad (33)$$

This is the sufficient condition for local stability of Model B, with queue feedback as $b > 0$. This can be seen from Fig. 3, where $\forall b > 0$, the region $a < \pi/4$ falls in the stable region.

In (32), we note that as $b \to \infty$, $a \to \pi/2$. This can be shown by the fact that

$$\lim_{b \to \infty}\left(1 + \frac{b}{4} - \sqrt{\frac{b^2}{16} + \frac{b}{2}}\right) = 0. \quad (34)$$

b) Without queue feedback: Consider the same scenario as in the previous case and choose $b = 0$ in (26). We linearize (26), with $b = 0$, as

$$\dot{r}(t) = -\left(\frac{aR^\star}{C\tau}\right)r(t - \tau). \quad (35)$$

where $R^\star = C$, in this case. Substituting $\zeta = a/\tau$ in (35), we get the characteristic equation as

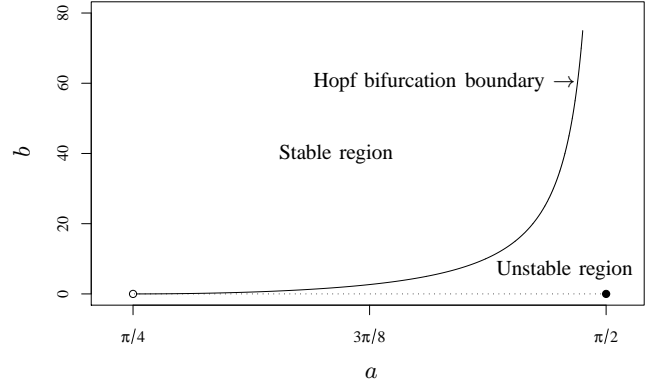$$\lambda + \zeta e^{-\lambda\tau} = 0. \quad (36)$$

Fig. 3. Stability chart for Model B

Following the same approach as in Section-IIB1a, we get

$$a < \frac{\pi}{2} \quad (37)$$

as the necessary and sufficient condition for the local stability of the Model B when $b = 0$, that is without queue feedback.

We plot the stability chart for this model in Fig. 3. Note that the region above the curve represents the stability region for $b > 0$, whereas when $b = 0$, the system is stable iff $a < \pi/2$; hence there is a discontinuity at $b = 0$, as confirmed by (32) and (37).

c) Hopf bifurcation: We show the system in (30), i.e. with queue feedback, undergoes a Hopf bifurcation as a is varied. Consider (30) by rewriting κ as $\kappa = \xi a/\tau$ with $\xi = 2 + b/4 - \sqrt{b^2/16 + b/2}$. Then, we get

$$\lambda + \frac{a\xi}{\tau}e^{-\lambda\tau} = 0. \quad (38)$$

Differentiating the above equation with respect to a, we get

$$\frac{d\lambda}{da} = \frac{1}{a\tau - \frac{\tau}{\xi}e^{\lambda\tau}}. \quad (39)$$

Define $\text{Re}(z)$ as the real part of z, and $\text{sgn}(\cdot)$ as the sign function. Then,

$$\text{sgn}\left(\text{Re}\left(\frac{d\lambda}{da}\right)\right) = \text{sgn}\left(\text{Re}\left(\frac{1}{a\tau - \frac{\tau}{\xi}e^{\lambda\tau}}\right)\right)$$
$$= \text{sgn}\left(\text{Re}\left(\xi a - e^{\lambda\tau}\right)\right) = \text{sgn}\left(\text{Re}\left(a + \frac{a}{\lambda\tau}\right)\right) = \text{sgn}(a) > 0, \quad (40)$$

where the second to last equality is obtained by substituting for $e^{\lambda\tau}$ from (38), and the last equality is because $\text{Re}(\lambda) = 0$ at the boundary. This shows that at the boundary of the stability region $\text{Re}\left(\frac{d\lambda}{da}\right) \neq 0$, while $\text{Re}(\lambda) < 0$ inside the region. Hence, λ changes signs across the region and the fixed point becomes unstable, i.e. we say that it undergoes

a Hopf bifurcation. The condition $\operatorname{Re}\left(\frac{d\lambda}{da}\right) \neq 0$ is called the transversality condition for a Hopf bifurcation.

Similarly, the above analysis also applies for the case of no queue feedback with (36), along the same lines as (30).

C. Impact of queue feedback

In context to the parameters chosen for packet level simulation results in [8, Fig. 5], the stability chart in Fig. 1 shows that $a = 0.5$ and $\beta = 1$ is outside the provably stable region, corroborating the results in [8, Fig. 5]. In the absence of queue feedback, to target a utilization of 90%, $a = 1$, $\beta = 0$ and $\gamma = 0.9$ where C is replaced with γC. In this case, the RCP model only reacts to rate mismatch. We note that the simulations in [8] did not produce any deterministic instabilities, and our analysis suggests that at these parameters the protocol is locally stable. Therefore, in this regime, the presence of queue feedback causes the queue to be *less* accurately controlled, suggesting a fundamental difference between these two forms of feedback and also showing evidence in favor of no queue feedback.

For model B, we have analytically shown that a Hopf bifurcation arises at the boundary of the stability region, which signifies the emergence of limit cycles. Of course, it is important to determine the stability of the bifurcating periodic orbit. An analytical characterisation of the stability or instability of the bifurcating limit cycle is beyond the scope of this paper. However, the computations performed for Model B, shown in Fig. 4, suggest that the limit cycles could indeed be stable. In Fig. 4, with queue and without queue feedback, we observed stable limit cycles that were plotted in the corresponding bifurcation diagrams. We observe that the corresponding phase portraits are topologically equivalent. We would like to highlight that the computations done for Model A, represented in Fig. 2, also provide evidence for the emergence of stable limit cycles as protocol parameters are varied. To that end, the packet level simulations shown in [8], which exhibit nonlinear oscillations appear to be induced via a Hopf type bifurcation. Our computational and analytical results for the Hopf bifurcation together provide rather strong evidence for a thorough investigation on the nonlinear dynamical characteristics for the various RCP models outlined in this paper.

D. Impact of utilization

We shall now look at how varying utilization impacts stability. From Fig. 4(b) and 4(d), we see that as utilization is decreased, the system enters a limit cycle at a later value of parameter a. The amplitude of the limit cycle also grows slower. This effect is more pronounced with queue feedback. This suggests to us that as utilization is decreased, the stability of the system increases.

E. Impact of parameter a

We observe in Models A and B that the parameter a appears along with the rate mismatch term, $C - y(t)$. Due to this, parameter a affects the speed at which the equilibrium rate

is attained as the magnitude of the rate mismatch feedback changes with a. We now show this using theoretical analysis.

Note that if $\operatorname{Re}\left(\frac{d\lambda}{da}\right) < 0$ and $\operatorname{Re}\left(\lambda\right) < 0$, then we observe that the solution of the linearized system in (29) decays faster as a increases. This motivates our analysis. Using the second to last expression in (40), the condition for $\operatorname{Re}\left(\frac{d\lambda}{da}\right) < 0$ is

$$\operatorname{Re}\left(\lambda\right) > -\frac{1}{\tau}. \tag{41}$$

We solve (36) for $\lambda = -1/\tau$ and obtain $a = 1/e$. We also note that when $a = 0$, $\lambda = 0 > -1/\tau$, hence $\operatorname{Re}\left(\frac{d\lambda}{da}\right) < 0$. Obviously, there exists a particular value of a after which $\lambda < -1/\tau$ and this value of a is $1/e$. Therefore, we note that the rate of convergence increases upto $a = \frac{1}{e}$, and then decreases again beyond that. Simulations confirming this analysis have been presented in [4].

Furthermore, we note again that parameter a also affects the stability of the system, as shown by our analysis in this section and in Fig. 2 and 4, and causes limit cycles for sufficiently large a. From the above analysis, we conclude that the optimal value of a for the fastest rate of convergence is

$$a = \frac{1}{e}. \tag{42}$$

III. CONCLUSIONS

Our focus, in this paper was on the equilibrium properties of two specific models of RCP. These two nonlinear models have only been studied for sufficient conditions to ensure local stability. To that end, for both these models, we developed necessary and sufficient conditions for local stability, under certain conditions. As conditions for stability get violated, bifurcations may occur. For both the models, we explored the consequences of parameters violating the stability conditions, and we plotted the respective bifurcation plots for the emerging stable limit cycles. For the small buffer variant of RCP, we also analytically showed that the bifurcation would be a Hopf bifurcation, which does signify the emergence of limit cycles. We used this insight, to help explain the potential destabilizing effect of having two forms of feedback in the protocol definition of RCP.

This, we believe, sheds light on a key architectural question concerning the design of RCP, i.e. whether the protocol needs to estimate the fair rate from both rate mismatch and from the queue size. Hopf type bifurcations, occurring due to the presence of queue size in the feedback, open additional questions regarding the nonlinear properties of the fluid models and their relationship with protocol design.
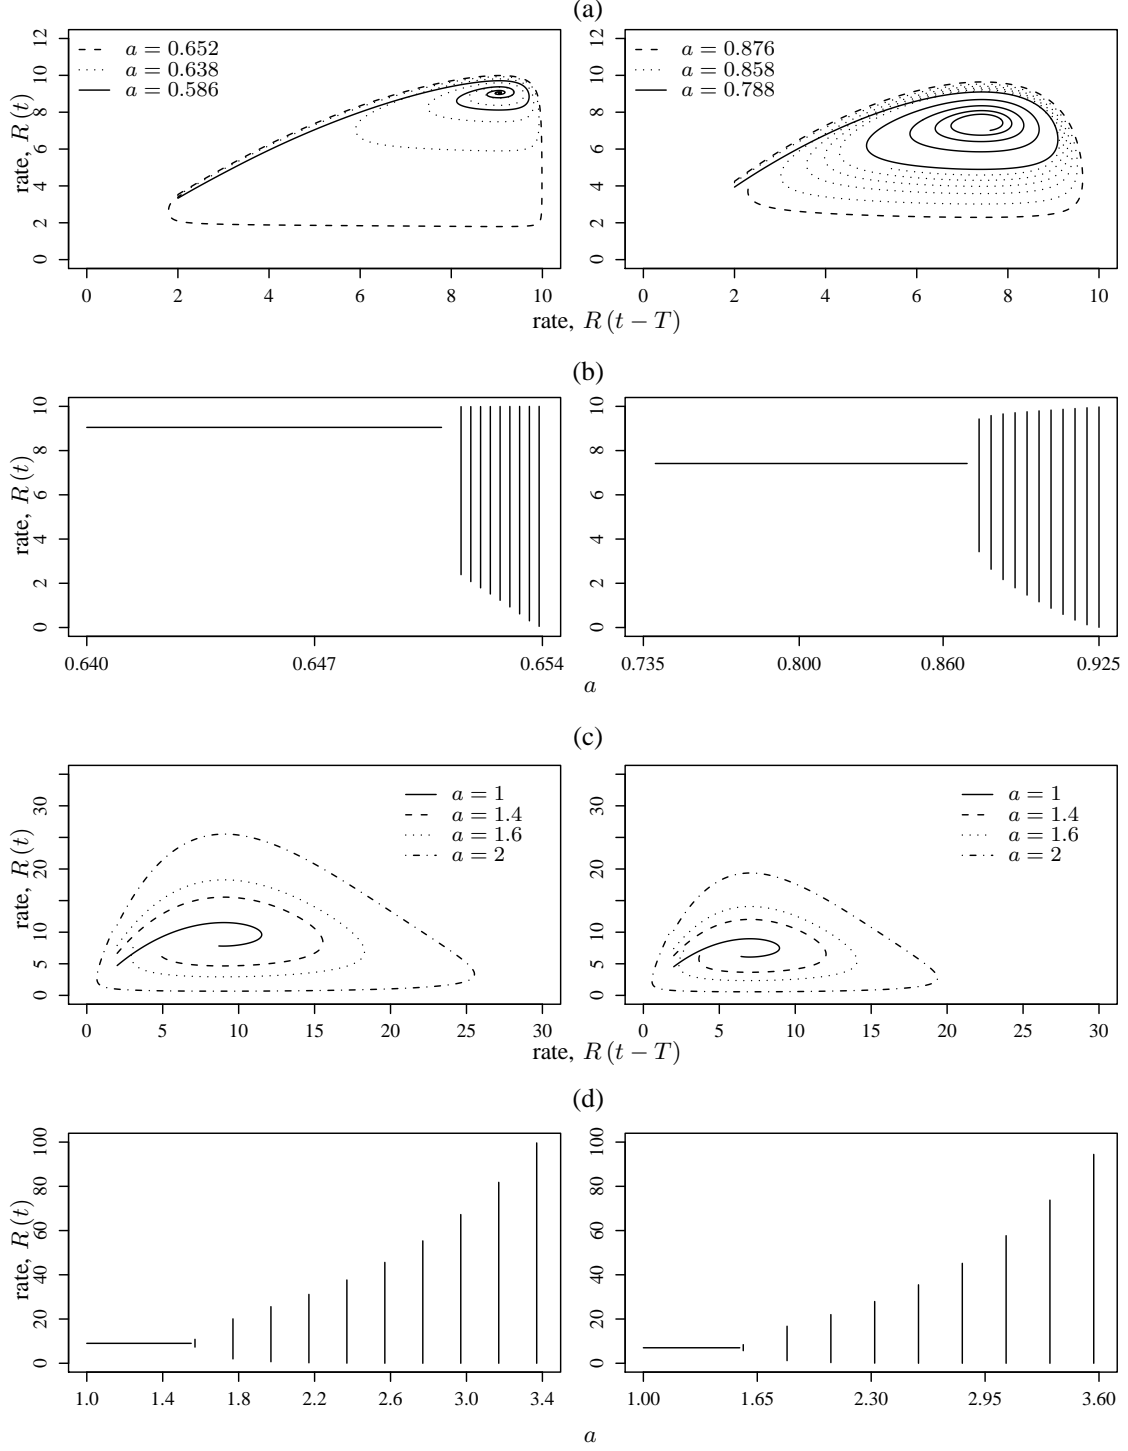
Fig. 4. Numerical computations for Model B – (a) Phase portrait, with queue: Utilization = 90% (left), 70% (right), (b) Bifurcation diagram, with queue: Utilization = 90% (left), 70% (right), (c) Phase portrait, without queue: Utilization = 90% (left), 70% (right), (d) Bifurcation diagram, without queue: Utilization = 90% (left), 70% (right). The values of the parameters are: RTT, $T = 1$ time unit; Capacity, $C = 10$ packets per unit time; with queue feedback – $b = 0.02$ for 90% utilization , $b = 0.18$ for 70% utilization; without queue feedback – $b = 0$, $\gamma = 0.9$ for 90% utilization, $\gamma = 0.7$ for 70% utilization. The values of parameter a in (a) are chosen to be proportionally spaced from the bifurcation boundary.

REFERENCES

[1] S. Floyd, "High-speed TCP for Large Congestion Windows," Internet Engineering Task Force, RFC 3649, 2003.

[2] T. Kelly, "Scalable TCP: Improving Performance in Highspeed Wide Area Networks," *SIGCOMM Comput. Commun. Rev.*, vol. 33, no. 2, pp. 83–91, Apr. 2003.

[3] D. Katabi, M. Handley, and C. Rohrs, "Congestion Control for High Bandwidth-delay Product Networks," *SIGCOMM Comput. Commun. Rev.*, vol. 32, no. 4, pp. 89–102, Aug. 2002.

[4] N. Dukkipati, "Rate control protocol (rcp): Congestion control to make flows complete quickly," Ph.D. dissertation, Dept. Elec. Engg., Stanford University, 2007.

[5] T. Voice and G. Raina, "Stability Analysis of a Max-Min Fair Rate Control Protocol (RCP) in a Small Buffer Regime," *Automatic Control, IEEE Transactions on*, vol. 54, no. 8, pp. 1908–1913, Aug 2009.

[6] H. Balakrishnan, N. Dukkipati, N. McKeown, and C. Tomlin, "Stability Analysis of Explicit Congestion Control Protocols," *Communications Letters, IEEE*, vol. 11, no. 10, pp. 823–825, October 2007.

[7] A. Lakshmikantha, R. Srikant, N. Dukkipati, N. McKeown, and C. Beck, "Buffer Sizing Results for RCP Congestion Control Under Connection Arrivals and Departures," *SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 1, pp. 5–15, Dec. 2008.

[8] F. Kelly, G. Raina, and T. Voice, "Stability and Fairness of Explicit Congestion Control with Small Buffers," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 3, pp. 51–62, Jul. 2008.

[9] N. Dukkipati, G. Gibb, N. McKeown, and J. Zhu, "Building a RCP (Rate Control Protocol) Test Network," in *High-Performance Interconnects, 2007. HOTI 2007. 15th Annual IEEE Symposium on*, Aug 2007, pp. 91–98.

[10] S. Jain and D. Loguinov, "PIQI-RCP: Design and Analysis of Rate-Based Explicit Congestion Control," in *Quality of Service, 2007 Fifteenth IEEE International Workshop on*, June 2007, pp. 10–20.

[11] G. Raina, "Local Bifurcation Analysis of Some Dual Congestion Control Algorithms," *Automatic Control, IEEE Transactions on*, vol. 50, no. 8, pp. 1135–1146, Aug 2005.

[12] Y. Zhang, S. Jain, and D. Loguinov, "Towards experimental evaluation of explicit congestion control," *Computer Networks*, vol. 53, no. 7, 2009.

[13] N. Dukkipati and N. McKeown, "Why Flow-completion Time is the Right Metric for Congestion Control," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 59–62, Jan. 2006.

[14] S. Lang, *Complex Analysis*, 4th ed.   Springer, 2003.