Some Equilibrium Properties of a Rate Control Protocol

A Project Report

submitted by

ABHIJIT KIRAN VALLURI

in partial fulfilment of the requirements
for the award of the degree of

BACHELOR OF TECHNOLOGY

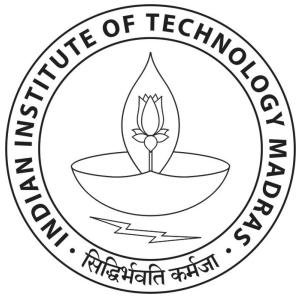

DEPARTMENT OF ELECTRICAL ENGINEERING
INDIAN INSTITUTE OF TECHNOLOGY MADRAS.

May 2011

THESIS CERTIFICATE

This is to certify that the thesis titled **Some Equilibrium Properties of a Rate Control Protocol**, submitted by **Abhijit Kiran Valluri**, to the Indian Institute of Technology, Madras, for the award of the degree of **Bachelor of Technology**, is a bona fide record of the research work done by him under our supervision. The contents of this thesis, in full or in parts, have not been submitted to any other Institute or University for the award of any degree or diploma.

Dr. Gaurav Raina
Research Guide
Visiting Faculty
Dept. of Electrical Engineering
IIT Madras, 600 036

Place: Chennai

Date: May 6, 2011

ACKNOWLEDGEMENTS

I am grateful to my guide, Dr. Gaurav Raina, for being a constant source of inspiration and ideas. His encouragement and motivation showed me the way through the project to its completion. I also express my gratitude to Prof. K.M.M. Prabhu, my faculty advisor, who has provided moral support, especially during the first year of my stay at IIT Madras. I am highly indebted to Ashish Krishna for his help and support throughout the project. I also express my gratitude to Shashank Jain, Ganesh Patil, D. Praveen Raja, Nizar M. and Shankar Raman M. J. for their helpful discussions and for their support during the various stages of my project. My heartfelt thanks to my parents who have been a source of support and affection, without which I could not have achieved any progress. Further, I take this opportunity to thank my classmates and friends who have provided a great atmosphere for both my academic and non-academic pursuits during the course of my study at IIT Madras.

Abhijit Kiran Valluri

ABSTRACT

KEYWORDS: RCP; Stability; Bifurcations; Feedback.

There is growing interest in explicit congestion control protocols as they promise a fair and a stable network.

In this thesis, we consider the Rate Control Protocol (RCP) and some different models that have been proposed to represent it. 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.

TABLE OF CONTENTS

LIST OF FIGURES

ABBREVIATIONS

RCP	Rate Control Protocol
RTT	Round Trip Time
AIMD	Additive Increase Multiplicative Decrease
TCP	Transport Control Protocol

NOTATION

a	Parameter a in the RCP rate evolution equation
β	Parameter β in the RCP rate evolution equation
b	Parameter b in a proportionally fair RCP rate evolution equation
$\dot{x}$	Derivative of x with respect to time
j	When used as a number in an equation, it represents $\sqrt{-1}$ (not to be confused with the subscript j, denoting resource j, used in some equations)

CHAPTER 1

Introduction

With ever increasing demands on the Internet, current architectures and protocols are under strain as users have high expectations on the quality of service they expect from communication networks. It is well known that congestion control is one such area where there has been a need to update the currently implemented protocols. Explicit congestion control is a promising avenue in the quest for fair and stable flow control algorithms. In the class of explicit flow control algorithms, the Rate Control Protocol (RCP) is a potential protocol for future high bandwidth-delay product environments.

A flow control algorithm should be fair, stable, and offer a low delay, low loss and high utilisation network. TCP uses loss as its feedback mechanism to manage its flow and congestion control algorithms. Thus it could be expected that such implicit feedback could make it difficult to enhance flow completion times. Additionally, numerous studies have exhibited that the standard Additive Increase Multiplicative Decrease (AIMD) TCP protocols (5; 6) are unsuitable for next generation networks. The need for more explicit feedback is well recognised, which in turn motivates the requirement for a comprehensive theoretical framework within which to design transport protocols.

1.1 Prior work on RCP

RCP has received a lot of attention from the research community in recent years. In (10), some stability properties of a max-min RCP in small buffer regime were developed. In (1), some sufficient conditions for local stability of explicit congestion protocols, such as RCP and eXplicit Congestion control Protocol (XCP), were developed computationally. In (8), a dynamic environment with RCP flows arriving and departing over a single link was considered. In (7), an α-fair variant of RCP was developed and some of the associated local stability properties were investigated. Also, in (3), a NetFPGA hardware

implementation of RCP was developed. The range of these studies exhibit the rather difficult nature of the problem of developing a new transport protocol.

Fluid models have arisen as a powerful tool within which to address some aspects of protocol design and for performance evaluation. To that end, in this thesis, we too will focus on the analysis of some nonlinear fluid models to better understand some of the equilibrium properties of RCP. In particular, 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. In order to understand the role played by modelling the queue as a deterministic fluid quantity, the authors in (1) study the stability of explicit congestion controllers by considering a linear switched control system with time delay using discretised Lyapunov functionals. They develop sufficient conditions to guide local stability based on computational analysis. Additionally, there are protocol parameters which influence stability and link utilisation; guidelines on how to choose these parameters are not fully developed.

1.2 Problem statement

1.2.1 Questions addressed

In this thesis, we deal with the following issues and answer the questions they pose.

1. We consider the role played by queue feedback in RCP performance. We provide evidence to suggest that at the packet level operation of the protocol, the non-switched queue dynamics, in RCP, is more appropriate and in this regime we develop necessary and sufficient conditions for local stability of RCP.

- We observe, from the results of some packet level simulations, the inherent nonlinear aspects of the protocol. We then study this computationally and explore the phenomena that occur when the conditions for local stability of the system are violated.

2. We then study another model for RCP, where the queue is not modelled 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 utilisation.

- Subsequently, we understand further aspects of this model, such as the impact of the various time scales, the effects of utilisation and the impact of the parameter a on the stability of the system.
- We also address the issue of queue feedback on the stability of the system, which is a key concern.

1.2.2 Why is this work significant?

Rate controlled protocols have a claim to be able to provide a network with high and controllable utilisation, very low loss, high throughput and short flow completion times. We, therefore, would like to explain the various phenomena under varying network conditions in which the system can run the protocols.

Moreover, since this is an engineered system, we have control over the microscopic rules that lead to desirable macroscopic properties. These desirable properties are stability, high utilisation, low loss and high throughput. Of this, stability is of top priority, as we need engineered systems to be well behaved so as to be able to provide guarantees on quality of service, such as high utilisation and throughput. Therefore, studies focused on issues of stability, and the impact of various parameters on the stability of the system are important, from an engineering perspective, as they aid us in choosing the correct design parameters that lead to an ideal functioning of the system.

1.3 Organisation of the thesis

The rest of this thesis is organised as follows. In Chapter 2, we outline and analyse the stability of some models for RCP. In Chapter 3, we study some more properties of RCP with respect to the small buffer variant of RCP. We summarise our contributions and outline avenues for future work in Chapter 4. Details regarding the design of the packet level simulator used for some of the results presented is described in Appendix A.

CHAPTER 2

Models of RCP

The first step towards understanding any protocol is modelling. To that end, in this chapter, we outline two nonlinear dynamical systems models for the Rate Control Protocol (RCP) (1; 7). 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. Most of the previous analysis has focused on conditions that may help to ensure that the nonlinear system is locally stable. To better understand the impact of stability a key focus will be on developing some necessary and sufficient conditions for stability, and exploring computationally and analytically the consequences of such local stability conditions being violated where bifurcation phenomena may readily occur.

2.1 Model A

The protocol strives to estimate the fair rate through a single bottleneck link from the rate mismatch and the queue size. In order to understand the performance of the protocol, the following nonlinear dynamical system for the rate and the queue has been proposed (1; 4):

$$\frac{d}{dt}R(t) = \frac{R(t)}{C\overline{T}}\left(a(C - y(t)) - \beta\frac{q(t)}{\overline{T}}\right) \tag{2.1}$$

where

$$y(t) = \sum_{s} R(t - T_s) \tag{2.2}$$

and

$$\begin{aligned} \frac{d}{dt}q\left(t\right) &= \left[y\left(t\right) - C\right] && q\left(t\right) > 0 \\ &= \left[y\left(t\right) - C\right]^{+} && q\left(t\right) = 0, \end{aligned} \tag{2.3}$$

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 (2.1), 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 (2.1) utilises two forms of feedback: one for the rate mismatch which is characterised by $C - y\,(t)$, and another for the instantaneous queue size, $q\,(t)$. The rate mismatch term causes the rate to increase if the utilisation 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 of equations (2.1), (2.2), (2.3), about its equilibrium point, were derived in (1), using analytics which were developed for a "switched" linear control system with a time delay. The analysis took due consideration of the discontinuity in the system dynamics, which would occur as the queue size approaches zero. This aspect of the model formulation was explicitly taken into consideration. We do, however, observe that the analysis was applicable to the fluid model, as opposed to a packet level description of the protocol.

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

$$a < \frac{\pi}{2} \tag{2.4}$$

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

We now present some packet level simulations, taken from (7) with permission, that have been performed for RCP using a discrete event simulator, details regarding which are presented in the Appendix A. A key objective of these simulations is to exhibit some nonlinear properties of RCP. They also serve to exhibit that the queue, in equilibrium, may not sit at zero but rather the mean queue size would be close to it.

The network being simulated consists of a single bottleneck link with a capacity of one packet per unit time and 100 Poisson sources. The RTT has been chosen as 100 time units. The parameters with queue feedback are: $a = 0.5$, $\beta = 1$ for a utilisation of

90%; the parameters without queue feedback are: $a = 1$, $\beta = 0$. When queue feedback is removed, we substitute the capacity, C, with γC to target a utilisation of $\gamma \times 100\%$. Hence, we choose $\gamma = 0.9$ to get a target utilisation of 90% when queue feedback is removed. In Fig. 2.1, we show the traces of queue size term, with and without queue feedback.

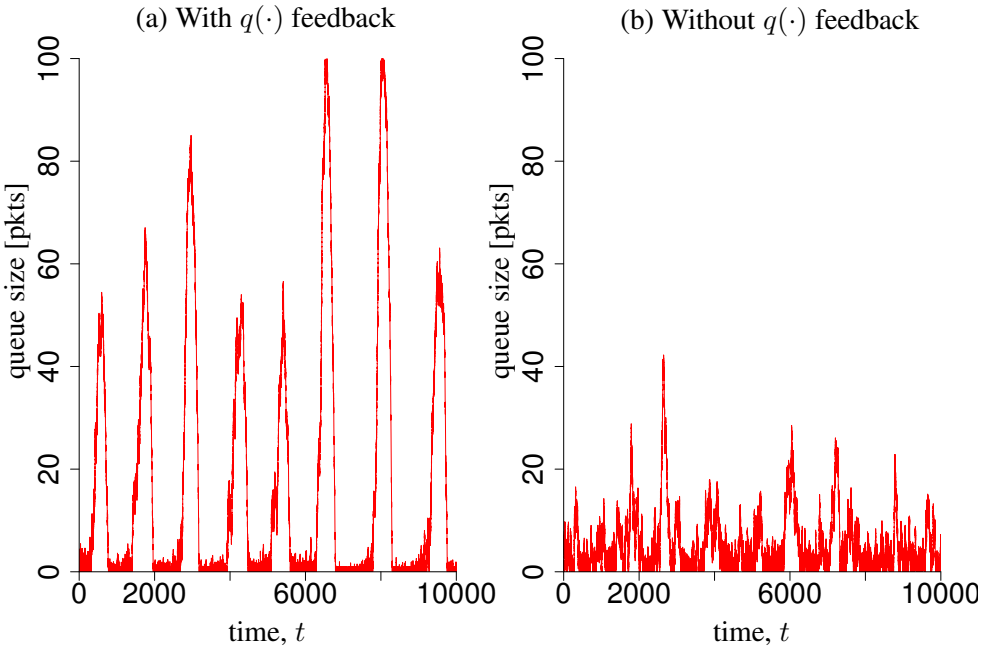

Figure 2.1: Traces from a packet-level simulation of a single bottleneck link with capacity of 1 packet per unit time, 100 RCP sources, round trip time of 100 time units, and a target link utilisation of 90%. The parameter values used are (a) $a = 0.5$; $\beta = 1$, and (b) $a = 1$; $\beta = 0$; $\gamma = 0.9$.

With queue feedback, and for the parameter values chosen, we observe deterministic instabilities in the queue size term in the form of nonlinear oscillations. On the other hand, without queue feedback, for the chosen parameter values, we observe that the queue size term is stable. However, the queue size has a non zero equilibrium value, which helps us to motivate 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 behaviour when conditions for stability are violated.

2.1.1 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 equation (2.3)

$$\dot{q}(t) = [y(t) - C] \qquad \forall q(t). \tag{2.5}$$

Let us consider that all the flows through the bottleneck link have a common RTT, T. Our system is now represented by the equations (2.1), (2.2), (2.5). The fixed points for these equations are:

$$\begin{aligned} R^{\star} &= C/n \\ q^{\star} &= 0, \end{aligned} \tag{2.6}$$

where n is the number of flows. Upon linearising equations (2.1), (2.2), (2.5) about the fixed point and on further simplification, we get

$$\begin{aligned} \dot{r}(t) &= -\frac{a}{T}(r(t-T)) - \frac{\beta}{nT^2} q(t) \\ \dot{q}(t) &= nr(t-T), \end{aligned} \tag{2.7}$$

where $R(t) - C/n = r(t)$. Substituting $\dot{q}(t+T) = nr(t)$ in equation (2.7) and taking a Laplace transform, we get:

$$S^2 T^2 e^{TS} + aTS + \beta = 0, \tag{2.8}$$

where S is the complex argument in the frequency domain that is obtained by taking the Laplace transform of the time domain equation (2.7). We consider two cases: when $\beta = 0$, which translates to the queue feedback being absent, and when $\beta > 0$ that implies that the queue feedback is present.

Without queue feedback

In this case, we substitute $\beta = 0$ in equation (2.7) and we get

$$STe^{TS} + a = 0. \tag{2.9}$$

If we let $TS = \lambda$, then equation (2.9) becomes

$$\lambda e^{\lambda} + a = 0. \tag{2.10}$$

To analyse this case, we introduce a new parameter, η, in equation (2.10) as follows:

$$\lambda e^{\lambda\eta} + a = 0. \tag{2.11}$$

Then, we substitute $\lambda = j\omega$, where $j = \sqrt{-1}$, and equate the real and imaginary parts to obtain

$$\begin{aligned} \omega \cos(\omega\eta) &= 0 \\ \omega \sin(\omega\eta) &= a, \end{aligned} \tag{2.12}$$

which simplifies to $\omega\eta = (2m+1)\,\pi/2$ and $\omega = a$. However, we consider the smallest value of $\omega\eta$ that satisfies equation (2.12) for which $a = \pi/(2\eta)$. We also note that when $\eta = 0$, the only root of equation (2.11) is $\lambda = -a$ that is stable. Now, using Rouché's theorem (9), we find that the system represented by equation (2.11) is stable if $\eta < \pi/(2a)$. With $\eta = 1$, we get back our original characteristic equation, and the stability condition becomes

$$a < \frac{\pi}{2}. \tag{2.13}$$

This is a necessary and sufficient condition for local stability for the system of equations (2.1), (2.2), (2.5) when $\beta = 0$, which implies that there is no queue feedback.

With queue feedback

With $\beta > 0$ we get the following characteristic equation from equation (2.8) by letting $TS = \lambda$:

$$\lambda^2 e^{\lambda} + a\lambda + \beta = 0. \tag{2.14}$$

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

$$\lambda^2 e^{\eta\lambda} + a\lambda + \beta = 0. \tag{2.15}$$

By substituting $\lambda = j\omega$ ($j = \sqrt{-1}$) in the above equation and equating the real and imaginary parts to zero, we get:

$$\begin{aligned} -\omega^2 \cos(\eta\omega) + \beta &= 0 \\ -\omega^2 \sin(\eta\omega) + a\omega &= 0. \end{aligned} \tag{2.16}$$

Upon simplifying the above equations, we get:

$$\omega^4 - a^2\omega^2 - \beta^2 = 0 \tag{2.17}$$

$$\Rightarrow \omega = \pm\sqrt{\frac{a^2 + \sqrt{a^4 + 4\beta^2}}{2}}, \tag{2.18}$$

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

$$\Rightarrow \eta = \frac{1}{\omega} \sin^{-1}\left(\frac{a}{\omega}\right). \tag{2.19}$$

Now, if a and β are fixed and $\eta = 0$, our characteristic equation becomes $\lambda^2 + a\lambda + \beta = 0$ whose roots have negative real parts as a is positive. Hence, when $\eta = 0$, the roots of equation (2.15) are stable. Once again, using Rouché's theorem, we get

$$\eta < \frac{1}{\omega} \sin^{-1}\left(\frac{a}{\omega}\right) \tag{2.20}$$

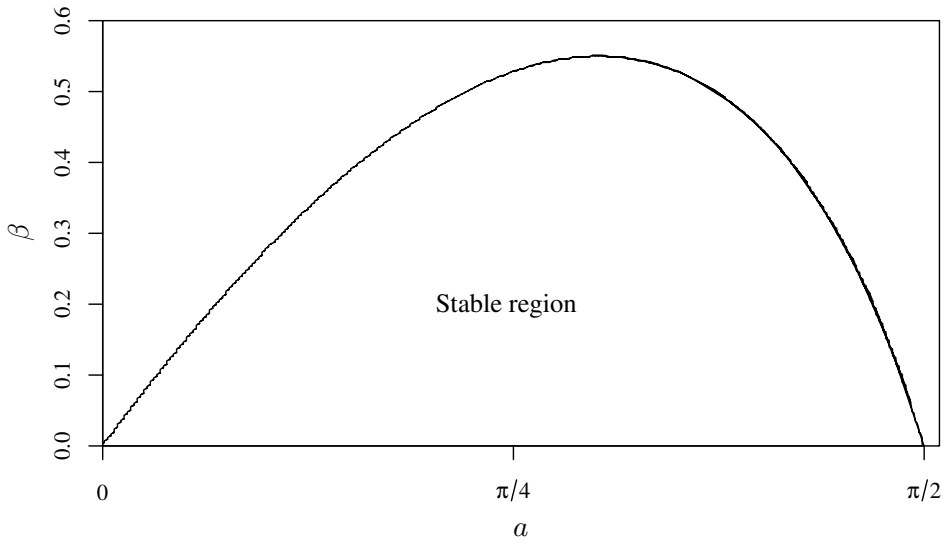

Figure 2.2: Stability chart for Model A

as the stability condition for the system corresponding to the characteristic equation (2.15). With $\eta = 1$, we obtain our original characteristic equation (2.14) from the characteristic equation (2.15). Therefore, the stability condition for the linearised system corresponding to the characteristic equation (2.14) is:

$$1 < \frac{1}{\omega} \sin^{-1}\left(\frac{a}{\omega}\right), \tag{2.21}$$

which can be simplified as,

$$\tan \frac{\sqrt{a^2 + \sqrt{a^4 + 4\beta^2}}}{\sqrt{2}} < \frac{a}{\beta} \frac{\sqrt{a^2 + \sqrt{a^4 + 4\beta^2}}}{\sqrt{2}}, \tag{2.22}$$

which is the necessary and sufficient condition for the local stability for the system of equations (2.1), (2.2), (2.5) when $\beta > 0$. In Fig. 2.2, we show the region in the $(a,\, \beta)$ parameter space where the linearised model of RCP is stable.

We performed numerical computations for Model A and presented the results in Fig. 2.3. We chose $\beta = 0.3$ and varied the parameter a. We chose RTT, $T = 0.1$ time units; capacity, $C = 100,000$ packets per unit time; number of flows, $s = 100$. We plotted the bifurcation diagrams and the phase portraits, both with and without the switch in the queue dynamics. Our first observation is that as the parameter a is varied across the boundary of the stability region, which we derived, the system becomes

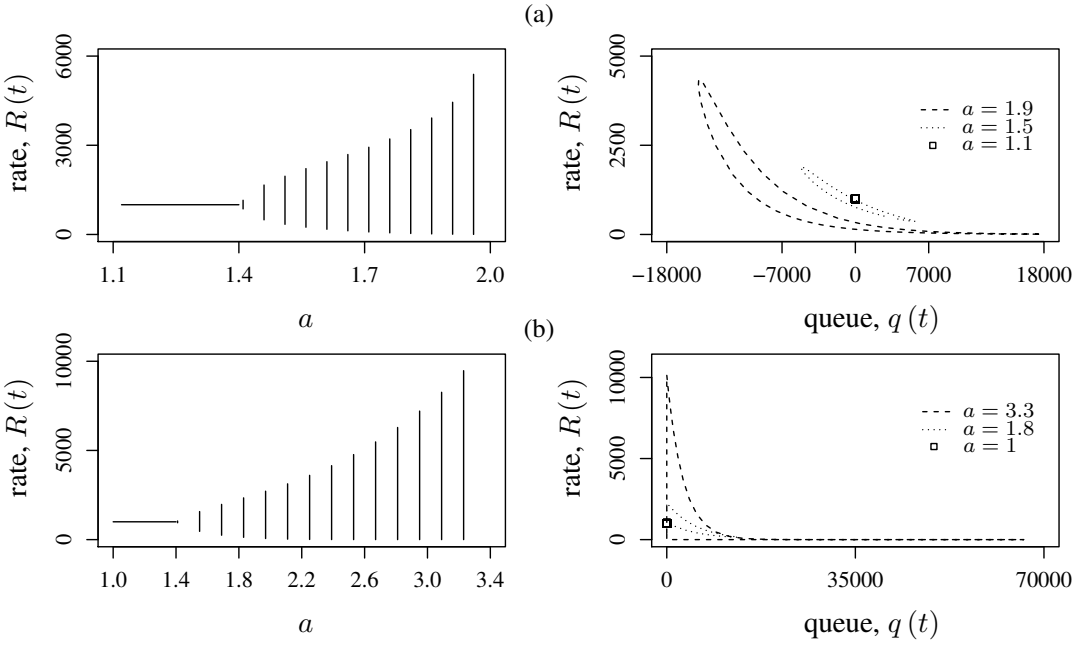

Figure 2.3: Numerical computations for Model A: (a) Bifurcation diagram (left) and phase portrait (right) for the "non-switched" case, with $\beta = 0.3$ and a is varied, (b) Bifurcation diagram (left) and phase portrait (right) for the "switched" case, with $\beta = 0.3$ and a is varied. Other parameter values are: RTT, $T = 0.1$ time units; Capacity, $C = 100,000$ packets per unit time; Number of flows, $s = 100$.

unstable. However, unlike linear systems, which blow up exponentially when unstable, the dynamics of the rate term enters a nonlinear cycle.

These cycles 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.) We also note that the fixed point becomes unstable. When a fixed point switches its stability as a parameter is varied, it is said to undergo a bifurcation. Hence, we note that the dynamics of Model A undergo a bifurcation at the boundary of the stability region. We note that the analysis in section 2.1.1 is done for the system of equations (2.1), (2.2), (2.5). We have neglected the switching in the dynamics of the queue size to highlight the nonlinearity in the dynamics of the rate term. Nevertheless, we note that the actual behaviour of RCP is better represented by the system of equations (2.1), (2.2), (2.5). For instance, from Fig. 2.1 we note that the mean queue size in the packet level simulations was not zero. Hence, the switch in the queue size dynamics does not play a vital role, since the equilibrium queue size

is non zero. To further understand the difference between the two scenarios, we again refer to the results of the numerical computations performed for Model A, presented in Fig. 2.3. We note that in both the scenarios, the dynamics of the rate term undergoes a bifurcation with respect to the parameter a. The phase portraits show that both cases are *topologically equivalent* as we can obtain one phase portrait from the other by distorting it appropriately. Rigorously speaking, *topologically equivalence* implies that there exists a *homeomorphism* (a continuous deformation with a continuous inverse) that maps one local phase portrait onto the other, such that a trajectory is mapped to a corresponding trajectory and the sense of time (direction of the arrows on the phase portrait) is maintained. It is, therefore, reasonable to analyse Model A without the switch in the queue size dynamics.

Elaborating on topological equivalence, we note that the fixed points and closed orbits of two topologically equivalent phase portraits share the same stability as we can obtain one phase portrait from the other by merely distorting it (without allowing for ripping of the phase portraits). Moreover, the sense of time is also maintained. Hence closed orbits remain closed and trajectories that connect saddle points remain unbroken, etc. More importantly, the stability of these fixed points and closed orbits also remains the same. This is an important observation about phase portraits that we shall use again, later on.

We now present another model used to represent RCP.

2.2 Model B

Model A accounts for the queue term explicitly via the differential equations (2.3), (2.5). 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 (7)

$$\frac{d}{dt}R_j(t) = \frac{aR_j(t)}{C_j\overline{T}_j(t)}\left(C_j - y_j(t) - b_jC_jp_j(y_j(t))\right) \tag{2.23}$$

where

$$y_j(t) = \sum_{r:j\in r} x_r(t - T_{rj}) \tag{2.24}$$

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

$$\overline{T}_j(t) = \frac{\sum_{r:j\in r} x_r(t)\,T_r}{\sum_{r:j\in r} x_r(t)} \tag{2.25}$$

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 equation (2.23), 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

$$x_r(t) = w_r\left(\sum_{j\in r} R_j(t - T_{jr})^{-1}\right)^{-1}, \tag{2.26}$$

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

$$p_j(y_j) = \frac{y_j\sigma_j^2}{2(C_j - y_j)}. \tag{2.27}$$

In essence, the queue size term is being modelled by $p_j(y_j)$ as described in equation (2.27). 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

$$b_j = \frac{\beta}{aC_j\overline{T}_j}. \tag{2.28}$$

The parameter b_j affects the utilisation of resource j at equilibrium. From equation (2.27) and considering the condition for the equilibrium of the nonlinear dynamical system depicted by Model B, we can evaluate the utilisation of resource j, ρ_j as

$$\rho_j \equiv \frac{y_j}{C_j} = 1 - \sigma_j\sqrt{\left(\frac{b_j}{2}.\frac{y_j}{C_j}\right)}, \tag{2.29}$$

which can be simplified as

$$\rho_j = 1 - \sigma_j\sqrt{\left(\frac{b_j}{2}.\rho_j\right)}, \tag{2.30}$$

which simplifies to

$$\rho_j = 1 - \sigma_j\sqrt{\left(\frac{b_j}{2}\right)} + O\left(\sigma_j^2 b_j\right). \tag{2.31}$$

2.2.1 Stability analysis

In (7), a sufficient condition for the local stability of Model B, for heterogeneous propagation delays, has been presented in the presence and absence of queue feedback. It is shown that, for the system to be stable in the presence of queue feedback,

$$a < \frac{\pi}{2}, \tag{2.32}$$

is a sufficient condition. Whereas, for the system to be stable in the absence of queue feedback,

$$a < \frac{\pi}{4}, \tag{2.33}$$

is a sufficient condition.

Here, we present the necessary and sufficient conditions for the local stability of this

model, for a homogeneous propagation delay, both in the presence and in the absence of queue feedback. We note that we get similar sufficient conditions as mentioned in (7). We also prove, analytically, the existence of a Hopf bifurcation at the edge of the stability region, both with and without queue feedback.

With queue feedback

For analysing the stability of Model B, we consider the following scenario. Let the network consist of a single bottleneck link with capacity C, a single route and a common RTT, τ, for all the flows. As there exists only one route and one resource, we shall drop the subscripts, j and r. 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 equation (2.23) becomes

$$\frac{d}{dt}R\left(t\right) = \frac{aR\left(t\right)}{C\tau}\left(C - y\left(t\right) - bCp\left(y\left(t\right)\right)\right) \tag{2.34}$$

where

$$y\left(t\right) = R\left(t - \tau\right) \tag{2.35}$$

and

$$p\left(y\right) = \frac{y}{2\left(C - y\right)}. \tag{2.36}$$

We now linearise equation (2.34) using the Taylor expansion of the right hand side about the equilibrium point. Let us consider $R\left(t\right) = r\left(t\right) + R^{\star}$, where $R^{\star}$ is the equilibrium value of $R\left(t\right)$, and $r\left(t\right)$ is a small perturbation about the equilibrium point. Using this and using equations (2.35), (2.36) we can simplify equation (2.34) as

$$\frac{d}{dt}\left(r\left(t\right) + R^{\star}\right) = \frac{d}{dt}\left(r\left(t\right)\right) = -a\left(\frac{R^{\star} + C}{C\tau}\right)r\left(t - \tau\right). \tag{2.37}$$

At equilibrium, the right hand side of equation (2.34) is zero. Hence, the equilibrium rate is

$$R^{\star} = C\left(\frac{b + 4 - \sqrt{b^2 + 8b}}{4}\right). \tag{2.38}$$

Since there is queue feedback, the parameter b is positive. Using equations (2.37), (2.38), we get

$$\frac{d}{dt}\left(r\left(t\right)\right) = -\kappa r\left(t-\tau\right), \tag{2.39}$$

where $\kappa = a\left(2 + b/4 - \sqrt{b^2/16 + b/2}\right)/\tau$. By taking the Laplace transform of equation (2.39), we get the characteristic equation

$$\lambda + \kappa e^{-\lambda\tau} = 0, \tag{2.40}$$

where λ is the complex argument in the frequency domain that is obtained by taking the Laplace transform of the time domain equation (2.39). Once again, we use the same trick used earlier and introduce a new parameter η as follows:

$$\lambda + \kappa e^{-\lambda\tau\eta} = 0. \tag{2.41}$$

By substituting $\lambda = j\omega$, where $j = \sqrt{-1}$, in the above equation and comparing real and imaginary parts, we get

$$\begin{aligned} \kappa\cos\left(\omega\tau\eta\right) &= 0 \\ \kappa\sin\left(\omega\tau\eta\right) &= \omega. \end{aligned} \tag{2.42}$$

On solving, we get $\kappa\tau\eta = \pi/2$. We note that when $\eta = 0$, the only solution of the characteristic equation (2.41) is $\lambda = -\kappa$, which is stable. Hence, using Rouché's theorem, we note that $\kappa\tau\eta < \pi/2$ is the necessary and sufficient condition for the stability of system represented by the characteristic equation (2.41). We obtain our original characteristic equation by choosing $\eta = 1$ in the characteristic equation (2.41). Hence, the necessary and sufficient condition for the local stability of Model B is $\kappa\tau < \pi/2$, which is

$$a\left(2 + \frac{b}{4} - \sqrt{\frac{b^2}{16} + \frac{b}{2}}\right) < \frac{\pi}{2}. \tag{2.43}$$

If we substitute $b \to 0$ in equation (2.43), we get

$$a < \frac{\pi}{4}. \tag{2.44}$$

This is the sufficient condition for local stability of Model B, with queue feedback since $b \to 0^+$ (where 0^+ refers to approaching zero from the right, in the limit $b \to 0$). This is because, as seen from the stability chart in Fig. 2.4, for all $b > 0$, the region $a < \pi/4$ falls in the stable region.

In equation (2.43), we note that as $b \to \infty$, $a \to \pi/2$. This can be shown by the fact that

$$\lim_{b\to\infty} \left(1 + \frac{b}{4} - \sqrt{\frac{b^2}{16} + \frac{b}{2}}\right) = 0. \tag{2.45}$$

Without queue feedback

Again, we consider the same scenario as with queue feedback. Hence, we consider the rate equation (2.34) and choose $b = 0$ in it and perform the analysis. Once again, 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 point. Using this and using equations (2.35), (2.36), we can linearise equation (2.34), with $b = 0$, as

$$\frac{d}{dt}\left(r(t) + R^\star\right) = \frac{d}{dt}\left(r(t)\right) = -\left(\frac{aR^\star}{C\tau}\right) r(t-\tau). \tag{2.46}$$

At equilibrium, the right hand side of equation (2.34) is zero. However, in this case, $b = 0$ and therefore the equilibrium rate value turns out to be

$$R^\star = C. \tag{2.47}$$

Let us define $\zeta = a/\tau$. So, the right hand side of equation (2.46) is $-\zeta r(t-\tau)$. The characteristic equation, in this case, is

$$\lambda + \zeta e^{-\lambda\tau} = 0, \tag{2.48}$$

where λ is the complex argument in the frequency domain that is obtained by taking the Laplace transform of the time domain equation (2.39). Again, we introduce a new parameter η as follows:

$$\lambda + \zeta e^{-\lambda\tau\eta} = 0. \tag{2.49}$$

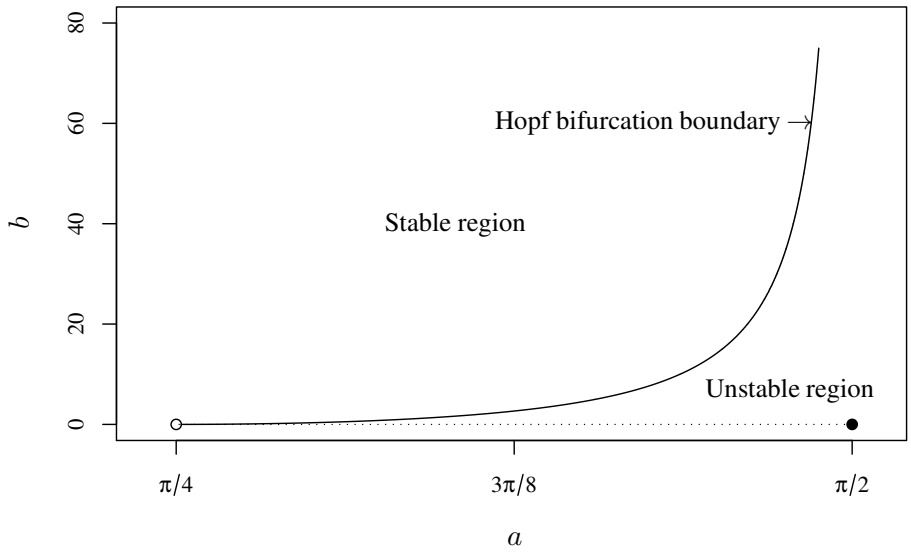

Figure 2.4: Stability chart for Model B

Proceeding similarly as with queue feedback case, we substitute $\lambda = j\omega$, where $j = \sqrt{-1}$, and equate the real and imaginary parts to zero and solve the equations to get,

$$\zeta\tau = a < \frac{\pi}{2} \tag{2.50}$$

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

To summarise, we have obtained the necessary and sufficient condition for the local stability of Model B when $b > 0$ (in the presence of queue feedback) in equation (2.43) and the sufficient condition alone is presented in equation (2.44). The necessary and sufficient condition for local stability of Model B when $b = 0$ (in the absence of queue feedback) is presented in equation (2.50).

We plot the stability chart for this model in Fig. 2.4. The region above the curve is the stability region. We note that there is a discontinuity in the stability chart. This provides further evidence supporting the two distinct necessary and sufficient conditions obtained above, for the two forms of feedback.

Hopf bifurcation

We now show that the system undergoes a Hopf bifurcation as a is varied. Consider the linearised rate equation (2.39), derived with queue feedback. Let us redefine κ as $\kappa = \xi \times a/\tau$, where $\xi = 2 + b/4 - \sqrt{b^2/16 + b/2}$. Then, we get the following characteristic equation:

$$\lambda + \frac{a\xi}{\tau} e^{-\lambda\tau} = 0. \tag{2.51}$$

Differentiating the above equation with respect to a, we get

$$\frac{d\lambda}{da} + \frac{\xi}{\tau} e^{-\lambda\tau} - a\xi e^{-\lambda\tau} \frac{d\lambda}{da} = 0, \tag{2.52}$$

$$\Rightarrow \frac{d\lambda}{da} = \frac{1}{a\tau - \frac{\tau}{\xi} e^{\lambda\tau}}. \tag{2.53}$$

Now, we find the sign of $\mathcal{R}\left(\frac{d\lambda}{da}\right)$ at the boundary region, where $\mathcal{R}(z)$ is defined as the real part of the complex number z.

$$\begin{aligned} \mathcal{R}\left(\frac{d\lambda}{da}\right) &= \mathcal{R}\left(\frac{1}{a\tau - \frac{\tau}{\xi} e^{\lambda\tau}}\right) \\ \Rightarrow \texttt{Sign}\left(\mathcal{R}\left(\frac{d\lambda}{da}\right)\right) &= \texttt{Sign}\left(\mathcal{R}\left(\frac{1}{a\tau - \frac{\tau}{\xi} e^{\lambda\tau}}\right)\right) \\ &= \texttt{Sign}\left(\mathcal{R}\left(\xi a - e^{\lambda\tau}\right)\right) \\ &= \texttt{Sign}\left(\mathcal{R}\left(a + \frac{a}{\lambda\tau}\right)\right) \\ &= \texttt{Sign}\left(a\right) > 0, \end{aligned} \tag{2.54}$$

where the last part was obtained by substituting for $e^{\lambda\tau}$ from equation (2.51). Since the real part of λ is zero, we get the sign of $\mathcal{R}\left(\frac{d\lambda}{da}\right)$ as that of a (since a is real), which is positive. This implies that at the boundary of the stability region, $\mathcal{R}\left(\frac{d\lambda}{da}\right) \neq 0$. However, we note that at this boundary, $\mathcal{R}(\lambda) = 0$, and within the stability region, $\mathcal{R}(\lambda) < 0$, as the system is stable. This implies that as we cross the stability region, $\mathcal{R}(\lambda)$ becomes positive, which means that the fixed point corresponding to this λ becomes unstable. Therefore, as a fixed point changes its stability across this boundary, we say that it

undergoes a Hopf bifurcation. Incidentally, the condition $\mathcal{R}\left(\frac{d\lambda}{da}\right) \neq 0$ is called the transversality condition for a Hopf bifurcation.

We have shown that the system undergoes a Hopf bifurcation at the boundary of the stability region, when there is queue feedback, as we started with equation (2.39). However, the same analysis holds for the case of no queue feedback as the only difference is that, now $\xi = 1$.

CHAPTER 3

Further properties of RCP

So far, we have described two nonlinear dynamical systems models of RCP and analysed their stability. We have also studied, analytically, the existence of a Hopf type bifurcation at the edge of the stability region in Model B. Using these analysis and understanding, we shall uncover certain other aspects of Model B that will aid us in the design process of RCP. To that extent, in this chapter, we shall first study the impact of the multiple time scales on RCP. A question of paramount concern is the effect of queue feedback on the stability of RCP, which we shall subsequently address in this chapter. We also study the impact of utilisation and the effect of parameter a on the stability of RCP.

3.1 Impact of multiple time scales

There exist four key time scales, which impact the dynamics of RCP, in the following ranges relative to queuing delay (abbreviated as qd): $\gg qd$, $> qd$, $\approx qd$ and $< qd$. Queuing delay is obtained as B/C where B is buffer size and C is capacity of the link.

We call the time scale where propagation delay is very high compared to queueing delay ($\gg qd$) as Time Scale 4 and number it downwards. We note that in this time scale, the network behaves deterministically. Model B, presented in Chapter 2, represents the system in this case. In Time Scale 3, the dynamics of the queue play a greater role and certain stochastic effects come into play as a result. In Time Scale 2, the variations in the rate and the queue size are in the same time scale and both the dynamics play a role. The system behaves deterministically and is well represented by Model A of Chapter 2. In Time Scale 1, the fluctuations in the rate are very rapid that there is no meaningful way we can account for its dynamics. Hence, the system behaves more like a stochastic one. There is currently no proper model to represent this time scale.

3.1.1 Analysis of the discrete event simulations

We now analyse the results, taken from (7) with permission, that have been obtained from the discrete event simulator, which is described in the Appendix A. The simulated network consists of a single bottleneck link with a capacity of one packet per unit time and 100 Poisson sources. The RTT is varied from 100 to 100,000 units of time. We choose $a = 0.5$ in Model B and choose b to attain different utilisation values using equation (2.31), when queue feedback is included. When there is no queue feedback, we choose $a = 1$, $b = 0$ and obtain a utilisation of $\gamma \times 100\%$ ($\gamma < 1$) by replacing C with γC. We plot, in Fig. 3.1, the utilisation as the parameter b is varied, with queue feedback, and plot, in Fig. 3.2, the utilisation as γ is varied, without queue feedback.

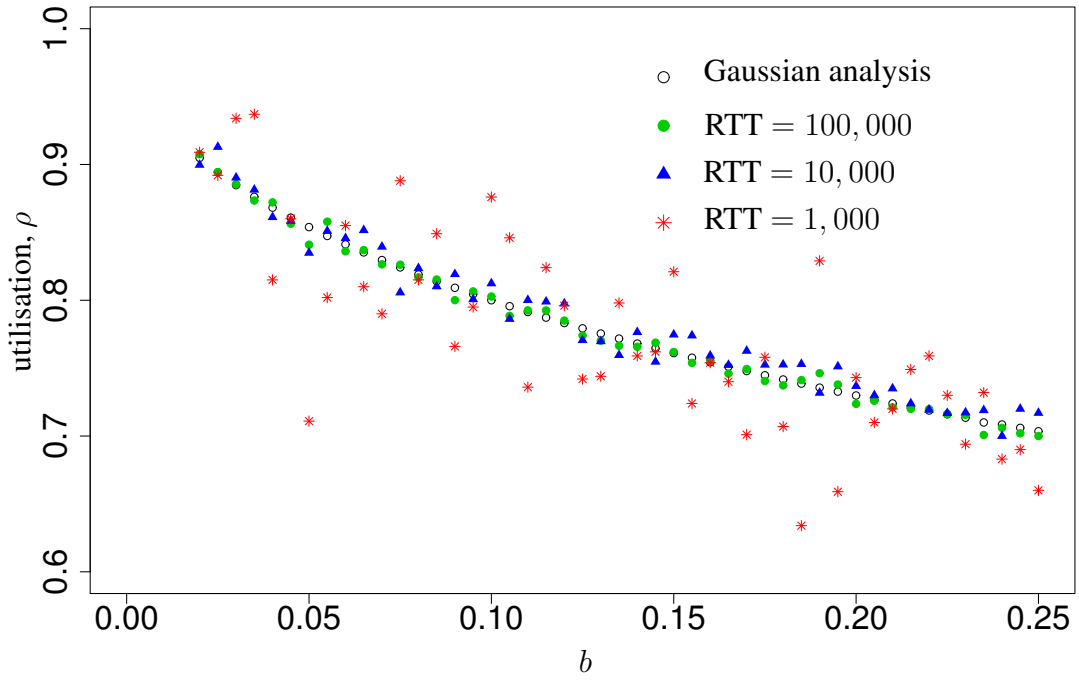

Figure 3.1: Utilisation, ρ, measured over one RTT, for different values of the parameter b with 100 RCP sources sending Poisson traffic.

In both the plots, we note that in the case where RTT = 100,000 units of time, the curves agree very well with the theoretical curve. Therefore, one can't differentiate between the two types of feedback. However, as the RTT is reduced, we note an increased variability in utilisation in both the plots. This clearly shows that RTT indeed impacts the stability and the dynamics of the system, whereas our linear analysis does

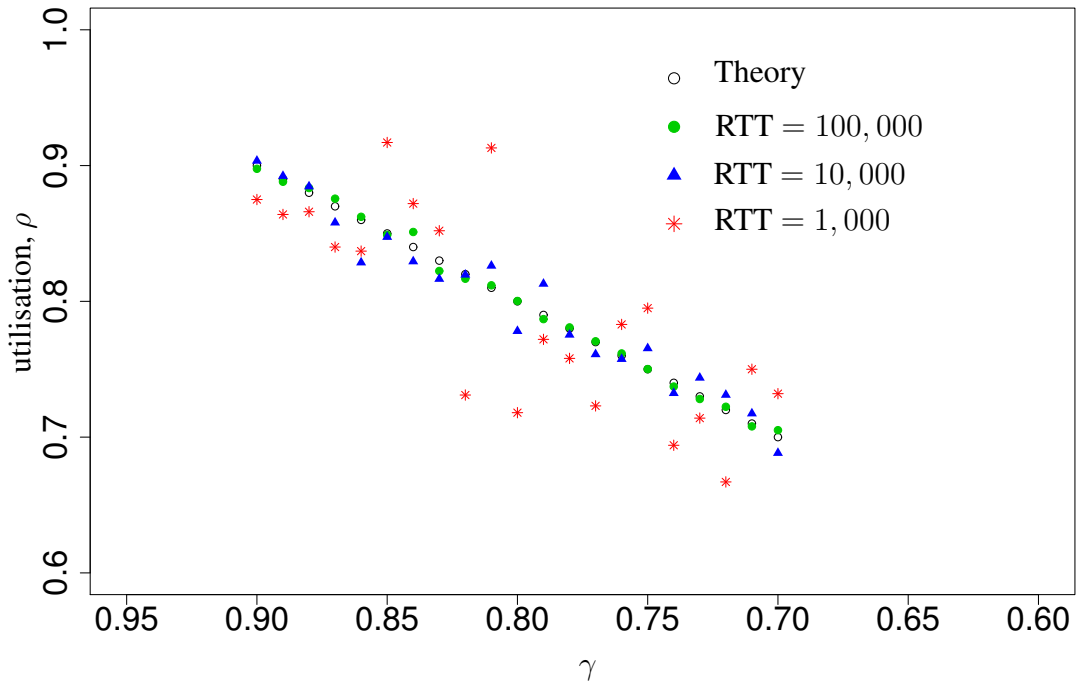

Figure 3.2: Utilisation, ρ, measured over one round trip time, for different values of the parameter γ with 100 RCP sources sending Poisson traffic.

not predict such a dependence. As RTT becomes very small, say 100 units of time, the queuing delay becomes comparable to the propagation delay. Hence, the underlying assumption of Model B breaks down, which relies on the fact that queuing delay is negligible compared to propagation delay.

3.1.2 Numerical Analysis

Using Model B, we have performed further simulations and obtained bifurcation diagrams as a is varied. We plot these bifurcation diagrams in Fig. 3.3 and 3.4. In each case, we have utilised a different value of RTT. We consider a network with a single bottleneck link with capacity $C = 10$ packets per unit time and choose $b = 0.2$ in Model B. For this, we analytically computed the stability criterion on a as $a < 0.935$ and we have obtained a utilisation of 68%. Using equation (2.27), we calculate the average queue size as 1.06. Therefore, we obtain queuing delay, $qd = B/C = 1.06/10 = 0.106$ units of time, where B is the buffer size, which is the average queue size and C is the link capacity. For these bifurcation diagrams, we choose RTT $\gg qd$ in order to consider

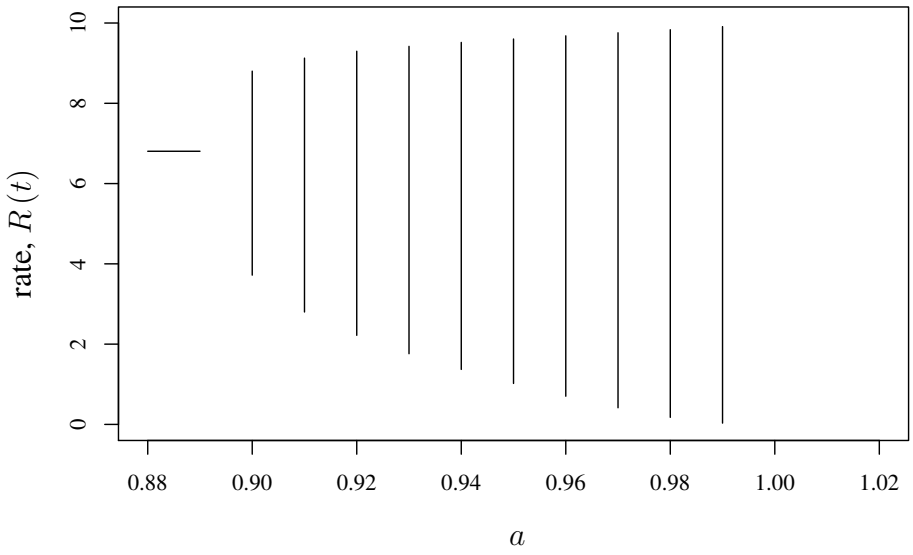

Figure 3.3: Bifurcation Diagram with queue feedback, RTT = 0.6 time units

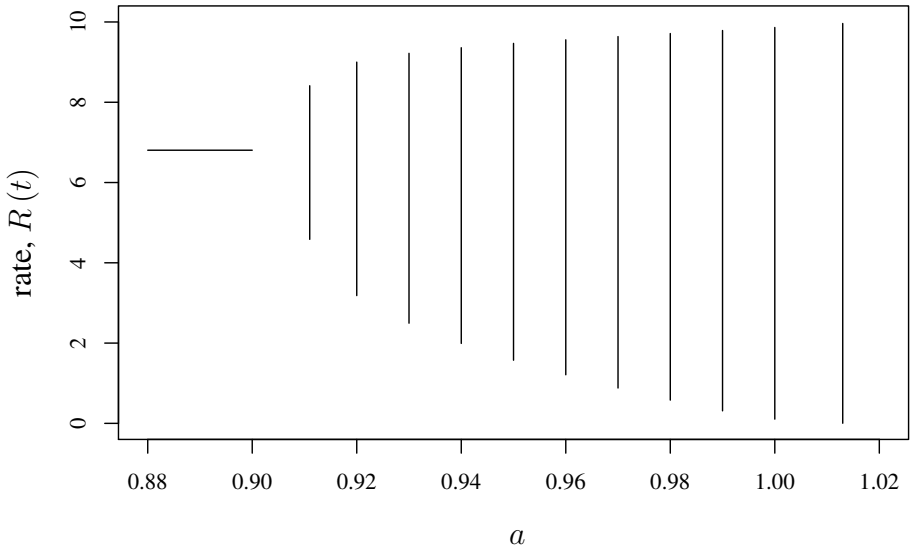

Figure 3.4: Bifurcation Diagram with queue feedback, RTT = 1 time units

cases in Time scale 4. Two convenient values of RTT are 0.6 and 1 time units and these are the values of RTT chosen for the two plots.

From these figures, we note that the bifurcation is affected as RTT is changed. Specifically, we note that as the RTT increases, the value of a at which the bifurcation

occurs increases slightly and the rate at which the amplitude of the limit cycle grows decreases. The overall amplitude of the limit cycle also decreases to a certain extent. Also, the limit cycles start earlier and finally become unstable later in case where RTT = 1 time unit than in the case of RTT = 0.6 time units. This clearly shows that RTT plays a more significant role in the stability of RCP and further analysis is required to comprehend the impact of time scales on RCP fully.

3.2 Impact of queue feedback

Recall that, in Fig. 2.1, we have plotted the traces of queue size from packet level simulations. We aimed at a utilisation of 90% by choosing the parameters $a = 0.5$, $\beta = 1$, with queue feedback. The stability chart in Fig. 2.2 shows that $a = 0.5$ and $\beta = 1$ is outside the provably stable region and indeed we attain deterministic instabilities in our simulations that can be seen in Fig. 2.1(a). In the absence of queue feedback, to target a utilisation of 90%, we choose $a = 1$, $\beta = 0$ and $\gamma = 0.9$ where C is replaced with γC. In this case, the RCP model only reacts to rate mismatch. From Fig. 2.2, we note that $a = 1$ and $\beta = 0$ is within the provably stable region. Indeed, the simulations did not produce any deterministic instabilities as can be evidenced from Fig. 2.1(b). Therefore, in this regime, the presence of queue feedback causes the queue to be *less* accurately controlled. This suggests that there is something fundamentally different between these two forms of feedback and also shows evidence in favour of no queue feedback.

In Model B, we have shown, analytically, 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 thesis. However, the computations performed for Model B, as depicted in Fig. 3.5, suggest that the limit cycles could indeed be stable. In Fig. 3.5, with queue and without queue feedback, we observed stable limit cycles that were plotted in the corresponding bifurcation diagrams. We observe that their corresponding phase portraits are topologically equivalent, implying that the stability of the corresponding

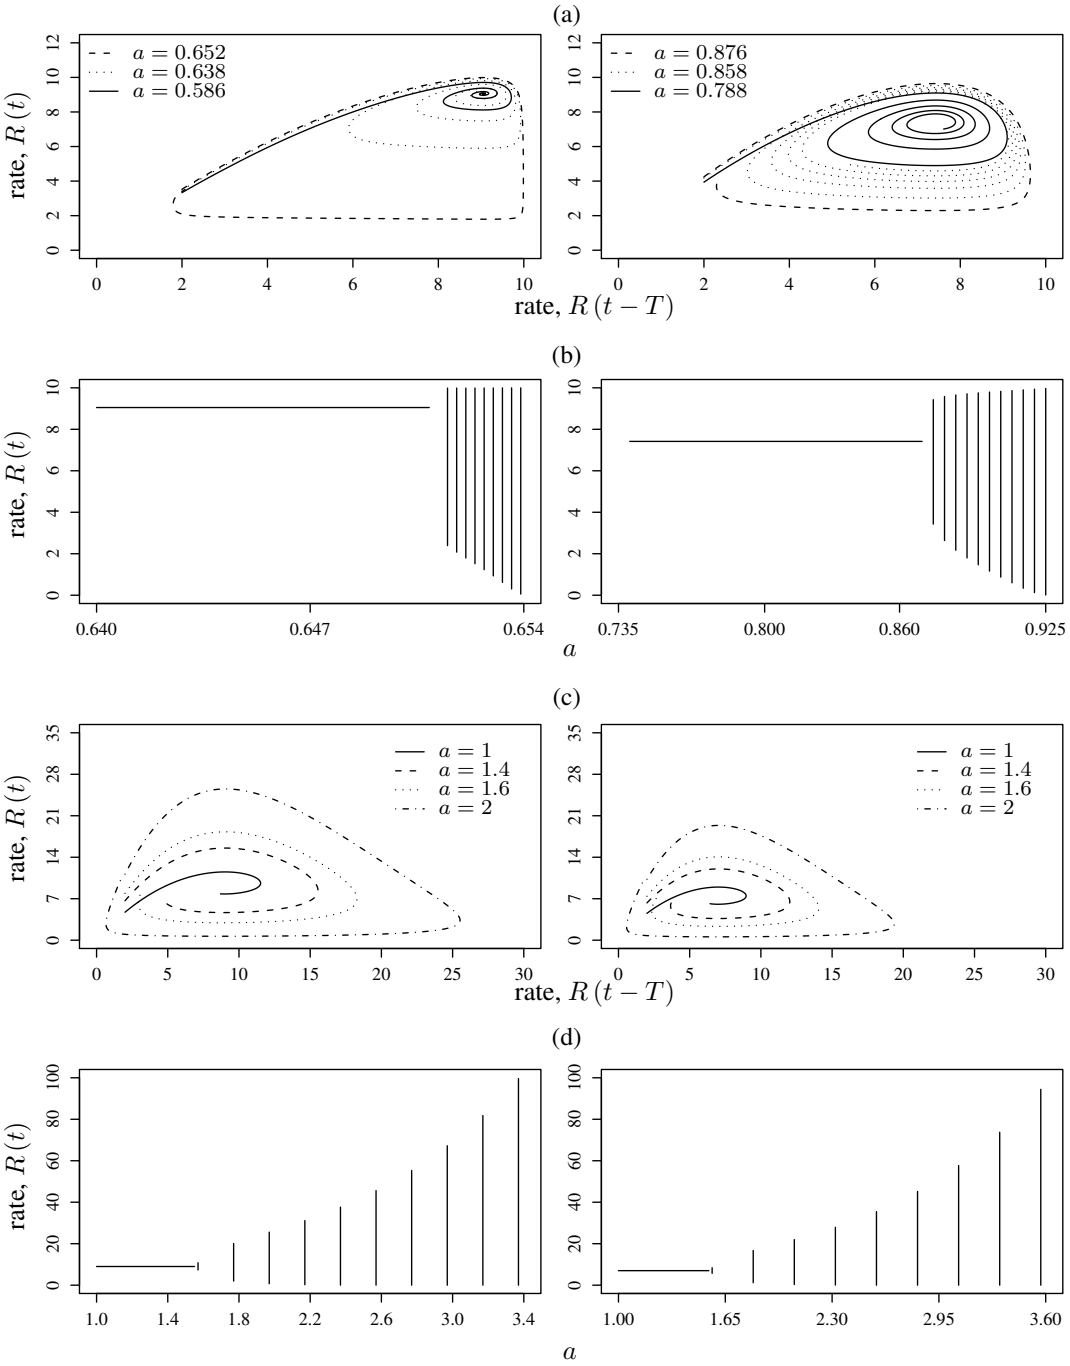

Figure 3.5: Numerical computations for Model B (on left, utilisation = 90%; on right, utilisation = 70%) – With queue feedback (a) Phase portrait, (b) Bifurcation diagram. Without queue feedback (c) Phase portrait, (d) Bifurcation diagram. The values of the parameters are: RTT, $T = 1$ time unit; Capacity, $C = 10$ packets per unit time; with queue feedback – $b = 0.02$ for 90% utilisation , $b = 0.18$ for 70% utilisation; without queue feedback – $b = 0$, $\gamma = 0.9$ for 90% utilisation, $\gamma = 0.7$ for 70% utilisation. The values of parameter a in (a) are chosen to be proportionally spaced from the bifurcation boundary.

fixed points and closed orbits, in the two scenarios, remain the same. We would like to highlight that the computations done for Model A, represented in Fig. 2.3, also provide evidence for the emergence of stable limit cycles as protocol parameters are varied. To that end, the packet level simulations shown in Fig. 2.1, which exhibit nonlinear oscillations appear to be induced via a Hopf type bifurcation. The simulations presented, the computations and the analysis 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 thesis.

3.3 Impact of utilisation

We shall now look at how varying utilisation impacts stability. From Fig. 3.5(b) and 3.5(d), we see that as utilisation 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 us that as utilisation is decreased, the stability of the system increases.

3.4 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.

We find the condition such that $\mathcal{R}\left(\frac{d\lambda}{da}\right) < 0$, where $\mathcal{R}(z)$ is defined previously. If $\mathcal{R}\left(\frac{d\lambda}{da}\right) < 0$ and $\mathcal{R}(\lambda) < 0$, then we realise that the solution of the linearised system in equation (2.39), which is $C_0 e^{\lambda t}$ where C_0 is a constant, decays faster as a increases. This motivates our line of analysis.

Taking the expression for the sign of $\mathcal{R}\left(\frac{d\lambda}{da}\right)$ obtained in section 2.2.1, we get that

$$\texttt{Sign}\left(\mathcal{R}\left(\frac{d\lambda}{da}\right)\right) = \texttt{Sign}\left(\mathcal{R}\left(a + \frac{a}{\lambda\tau}\right)\right). \tag{3.1}$$

On simplifying, the condition obtained for $\mathcal{R}\left(\frac{d\lambda}{da}\right) < 0$ is

$$\mathcal{R}\left(\lambda\right) > -\frac{1}{\tau}. \tag{3.2}$$

We solve the characteristic equation (2.51) for $\lambda = -1/\tau$ and obtain $a = 1/e$. We also note that when $a = 0$, $\lambda = 0$ which is greater than $-1/\tau$. Hence, $\mathcal{R}\left(\frac{d\lambda}{da}\right) < 0$, which implies that as a increases the value of λ decreases. 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 up to

$$a = \frac{1}{e}, \tag{3.3}$$

and then decreases again beyond that. Simulations confirming this analysis have been presented in (2).

Furthermore, we note that parameter a also affects the stability of the system, as is observed from the necessary and sufficient conditions for local stability that have been obtained in our stability analysis, in Chapter 2, for the two models of RCP considered. Fig. 2.3 and 3.5 show the bifurcation diagram depicting the rate value as a is varied. We note that the system enters a limit cycle for a range of values of a.

Considering the above analysis, we conclude that the optimal value of a should be

$$a = \frac{1}{e}, \tag{3.4}$$

corresponding to the fastest rate of convergence.

CHAPTER 4

Conclusions

Rate controlled protocols show potential for developing a fair and stable network architecture for the future high bandwidth-delay product networks. They claim to provide a network with high utilisation, very low loss, a high throughput and short flow completion times. This makes it all the more important to study the properties and design considerations of this class of protocols, one among which is the Rate Control Protocol (RCP). In this thesis, we focussed on some equilibrium properties of some models of RCP.

4.1 Contributions

Currently, the two nonlinear models that have been proposed 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 destabilising 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 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.

4.2 Future Work

First, it would be useful to exhibit that, in fact, both the models exhibit Hopf type bifurcations. Then, it would be imperative to analytically verify the stability of the emerging limit cycles as the Hopf conditions get violated. This would certainly enhance our understanding of the RCP models. Additionally, for the small buffer variant of RCP, it would be useful to ascertain the impact of the protocol parameters on the stochastic effects. A fuller understanding of both the nonlinear and the stochastic effects would help develop a comprehensive understanding of the protocol parameters, and the role of queue feedback.

APPENDIX A

Discrete event simulator - Packet level simulations

The program used for simulating the network running the Rate Control Protocol (RCP) was written in C# on Windows. Here, we describe the main logic behind the program.

After all the necessary variables have been declared, the following is done in the program during each time step:

for $i = 0$ to $maxIndex$ **do**

$PktSendTime[i] \leftarrow ExpRandomV(1/Rate)$

end for

$tempRate \leftarrow Rate$

for $i = 0$ to $maxIndex$ **do**

$PktSendTime[i] \leftarrow tempDelay[i] + PktSendTime[i] * 100000$

end for

In the above code, the program assigns the exponential inter-arrival time, between consecutive packets sent by each of the 100 sources, in the array $PktSendTime$ (implying that $maxIndex = 99$). The previous time that a source sent a packet is stored in $tempDelay$. We add the inter-arrival times to this array, so that the $PktSendTime$ array finally contains the actual time when each source sends the next packet. The multiplication by 100000 is done since 1 second is taken as 100000 time steps.

for $iter = 0$ to RTT **do**

$timestep++$

if $timestep \geq timeStepMax$ **then**

break

end if

if $pipe = 1$ **then**

$pipe \leftarrow 0$

end if

for $i = 0$ to $maxIndex$ **do**

if $PktSendTime[i] \leq timestep$ **then**

$numPktArrCurr++$

$PktSendTime[i] \leftarrow NextPktAt(i, tempRate)$

end if

$SendPacket()$

end for

if noPktSentAbove **then**

$SendPktFromQ()$

end if

$PropagatePkts()$

if newPktsArrivedAtLink **then**

$EvalNewRate()$

end if

$queueSize[timestep] \leftarrow currQSize$

$rate[timestep] \leftarrow currRate$

$EvalUtilisation(timestep)$

end for

Here, we have abstracted several tasks into convenient functions. Firstly, we note that a source sends a packet whenever the current time step is greater than the time step value contained in the array $PktSendTime$, which is the time the source is supposed to send a packet. When this occurs, the function $NextPktAt()$ generates the time at which a particular source needs to send the next packet, which is calculated similar to the way it was done in the first algorithm snippet. It adds the exponential inter-arrival time to the last time a packet was sent by a source. The resultant value is the time step at which the source will send the next packet.

The $SendPacket()$ function checks if the link is idle or not. This is identified via the $pipe$ variable. If it is 1, then a packet is currently being sent, and so the program adds the current packet to the queue, represented by an integer $queue$. This integer stores the number of packets in the queue. The program ensures that the queue size does not exceed 100. If $pipe$ is 0, then the program sends the packet onto the link. It

takes $RTT/2$ time steps for a packet to cross the link. Therefore, the program keeps a track of the number of time steps left before a packet is transmitted completely.

If none of the sources sends a packet in the current time step, then the program sends a packet from the queue using the $SendPktFromQ()$ function, wherein it decrements queue size by 1 and then sends the packet as done in $SendPacket()$.

Subsequently, the program $propagates$ the packets. As noted earlier, it takes $RTT/2$ time steps for a packet to transit the link. The function $PropagatePkts()$ handles this by decreasing the number of time steps remaining to transit a packet. It does this maintenance for each packet in the link. When a packet is completely transmitted, it adds 1 to the total number of packets sent.

The function $EvalNewRate()$ evaluates the updated value of rate as per the equation (2.1). It implements a discrete version of the differential equation and updates the rate whenever a new packet has arrived at the link. While updating the rate, the program may vary several parameters in equation (2.1), such as β or γ, in order to consider several scenarios, during different runs of the program. Subsequently, the current queue size and rate values are stored in corresponding arrays.

Finally, the program evaluates the utilisation values every time step via the function $EvalUtilisation()$. It does this in the following way: let the link have a capacity of sending C packets/second. Let the RTT be T. Then, in T time steps, the link should be able to send $C \times T$ packets if the link where used at 100% utilisation. If the link actually sent x packets in the previous T time steps, then the Utilisation is $\frac{x}{CT} \times 100\%$.

The above algorithm is repeated for each RTT. Once the program executes the above code, an RTT of time steps are over and the program updates the rate used for calculating when the next set of packets are sent. Once the program finishes executing the code for one set of parameters, the program stores relevant data such as Rate versus time, Queue Size versus time, Utilisation versus time, Utilisation versus γ and Utilisation versus β. Then, the program changes the parameters to consider a new scenario and the process continues.

This concludes our explanation of the functioning of our program.

REFERENCES

[1] H. Balakrishnan, N. Dukkipati, N. McKeown, and C. J. Tomlin, "Stability analysis of explicit congestion control protocols," *IEEE Communications Letters*, vol. 11, pp. 823–825, 2007.

[2] N. Dukkipati, "Rate Control Protocol (RCP): Congestion control to make flows complete quickly," Ph.D. Thesis, Dept. Elec. Engg., Stanford Univ., Palo Alto, CA, 2007.

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

[4] N. Dukkipati and N. McKeown, "Why flow-completion time is the right metric for congestion control," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 1, pp. 59–62, 2006.

[5] S. Floyd, "High-speed TCP for Large Congestion Windows," IETF RFC 3649, 2003.

[6] T. Kelly, "Scalable TCP: Improving Performance in High-speed Wide Area Networks," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 2, pp. 83–91, 2003.

[7] F. Kelly, G. Raina, and T. Voice, "Stability and fairness of explicit congestion control with small buffers," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 3, pp. 51–62, 2008.

[8] A. Lakshmikantha, R. Srikant, N. Dukkipati, N. McKeown, and C. Beck, "Buffer sizing results for RCP congestion control under connection arrivals and departures," *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 1, pp. 5–15, 2008.

[9] S. Lang, *Complex Analysis*, 4th ed. New York: Springer, 1998.

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