

Stability analysis of a max-min fair Rate Control Protocol (RCP) in a small buffer regime

Thomas Voice (T.Voice@statslab.cam.ac.uk)
 Gaurav Raina (G.Raina@statslab.cam.ac.uk)

Abstract

In this note we analyse various stability properties of the max-min fair Rate Control Protocol (RCP) operating with small buffers. We first tackle the issue of stability for networks with arbitrary topologies. We prove that the max-min fair RCP fluid model is globally stable in the absence of propagation delays, and also derive a set of conditions for local stability when arbitrary heterogeneous propagation delays are present. The network delay stability result assumes that, at equilibrium, there is only one bottleneck link along each route. Lastly, in the simpler setting of a single link, single delay model, we investigate the impact of the loss of local stability via a Hopf bifurcation.

keywords rcp, max-min fairness, small buffers, stability, bifurcation.

1 Introduction

The Rate Control Protocol (RCP) [5, 6] takes a radically different approach towards managing flow and congestion control as compared to the existing Transmission Control Protocol (TCP) congestion avoidance framework. The TCP framework has, imbedded in it, an implicit mechanism for detecting congestion within the network. Loss of a packet, caused by the overflow of a buffer, is intended to provide the necessary feedback information. In sharp contrast, RCP aims to achieve fast flow completion times by communicating *explicit* rate feedback between routers and end-stations [4]. RCP is closely related to the eXplicit Control Protocol, XCP [11]; another algorithm which proposes the use of explicit rate feedback. Both the RCP and XCP algorithms intend to converge to a max-min fair resource allocation [2, 5, 11]; see [16, 17, 18, 22] and references therein for a sample of the literature exploring various issues related to different notions of fairness in a networking context.

The performance of congestion control algorithms is often coupled with the choice of certain parameters in routers in the network. For example, such parameters may correspond to different choices of Active Queue Management (AQM) schemes [3, 10, 14, 15], or the size of buffers in routers [8, 20]. Researchers have begun questioning the design rules for sizing buffers in core routers in the

Internet [1] and some recent literature [7, 20, 25] suggests that buffers should indeed be much smaller as compared to the current design principles. In light of the separate developments made towards the understanding of RCP [2, 4, 5, 6] and buffer sizing [1, 7, 20, 25] it is natural to investigate the two, i.e., RCP and small buffers, together. Such an investigation is also motivated by the observation that previous buffer sizing studies have focussed primarily on the impact of smaller buffers with TCP. In this paper, the focus of our analysis will be on the limiting regime of a max-min fair RCP fluid model operating over a communication network with small buffers.

Previous control theoretic analysis of RCP has focussed on a single bottleneck, where the queue is modelled as a saturated integrator [2]. In our work, we assume that the size of the buffers is small enough so that it is no longer possible to explicitly model the queue. Rather, with such small buffers, at the time scale of operation of the congestion controllers, it is the distribution of queue size that plays the important role [25]. It is worth stating that we do not, in any way, contribute to the question of exactly how small the size of buffers ought to be. The model we analyse is simpler than the one analysed by [2] in that we do not explicitly model the queue, but more involved in the sense that our framework represents a heterogeneous network of arbitrary topology. In our fluid model for RCP, apart from the model for the queue, all other parameters are exactly the same as specified in the original RCP algorithm [5].

The styles of analysis we employ have all been popular in the study of congestion control: for example, global stability without propagation delays [21, 24], local stability with delays [8, 10, 12, 14, 15, 21, 23], and also an analysis of the dynamic system when stability may not be guaranteed [9, 13, 19].

We now outline the essence of our contribution. First, we prove that the fluid model of RCP is globally stable in the absence of propagation delays. Then, we derive a set of conditions for local stability when arbitrary heterogeneous propagation delays are present. The network delay stability result relies upon the assumption that, at equilibrium, there is only one bottleneck link along each route. Furthermore, in the simpler setting of a single link, single delay, model we are able to provide a more in depth study of the RCP algorithm. Using bifurcation analysis, we investigate the impact of the loss of local stability in a special case where we show that the RCP algorithm would always give rise to a super-critical Hopf bifurcation.

An *overview* of this paper follows. In Section 2, we analyse the stability properties of RCP over networks of arbitrary topology. In Section 3, we investigate local instability of RCP in a single link, single delay, model via a Hopf bifurcation analysis. Finally, in Section 4 we summarise our contribution, and discuss some avenues for further research.

2 RCP over arbitrary topology networks

In this section, our objective is to analyse the stability of RCP in a small buffer regime, over networks with an arbitrary topology. We first show global stability

for the RCP fluid model in the absence of propagation delays and then derive a set of conditions for local stability when arbitrary heterogeneous propagation delays are present. The network delay stability result relies upon the assumption that, at equilibrium, there is only one bottleneck link along each route. That is, at equilibrium, for each route there is a unique link with minimal per-flow bandwidth available. We begin with the model description.

2.1 Model description

At the level of theoretical abstraction, our set up for an Internet like communication network follows the commonly adopted framework [21].

We suppose that the communication network comprises of an interconnection of a set of *routes*, S , with a set of *links*, J . Each route $r \in S$ represents a user of the network. Associated with each route is a set of links which represents the path along which that user transmits information through the network. Further, a route r has associated with it a flow rate $x_r(t) \geq 0$, which represents a dynamic fluid approximation to the rate at which the user is sending packets along route r , at time t .

The flow rate, for each $r \in S$ is determined by the links $l \in r$, via the use of *explicit* rate feedback. Each link $l \in J$ has associated with it a flow rate $R_l(t)$, which represents the maximum flow rate allowed for routes which pass through l . Each packet which is sent through the network carries, in its header, an explicit rate feedback variable. This variable is initially set to the maximum desired flow rate for r . As the packet passes through each $l \in r$, if the feedback variable is greater than $R_l(t)$, then link l sets it equal to $R_l(t)$. When the packet reaches its destination, an acknowledgement packet (**ack**), containing the final value of the explicit rate feedback variable, is returned to the origin of r , and the flow rate $x_r(t)$ is updated accordingly.

For each route r and link $l \in r$, we let τ_{rl} denote the propagation delay from the origin of r to l , i.e. the length of time it takes for a packet to travel from the origin to link l along route r . Let τ_{lr} denote the propagation delay from l to the origin of r , i.e. the length of time it takes for the explicit rate feedback information from link l to reach the user who is transmitting along route r . In RCP, a packet must reach its destination before an acknowledgement packet (**ack**) is returned to its source. Furthermore, as we are interested in a network with small buffers, we may safely assume that queuing delays are a negligible component of the end to end delay. Thus for all $l \in r$, $\tau_{rl} + \tau_{lr} = \tau_r$, the round trip time for route r .

We now have the following model for the end-system behaviour of RCP. For each route $r \in S$,

$$x_r(t) = \min_{l \in r} R_l(t - \tau_{lr}). \quad (1)$$

Now, for each link $l \in J$, $R_l(t)$ is updated depending upon the total aggregate flow through link l at time t . Although the update rule is discrete, we can model it via a fluid approximation with the following differential equation [2]. For each

link $l \in J$,

$$\dot{R}_l(t) = R_l(t) \left(\frac{\alpha_l}{d_l C_l} (C_l - y_l(t)) - \frac{\beta_l q_l(t)}{d_l^2 C_l} \right)_{R_l(t)}^+, \quad (2)$$

where α_l, β_l are positive constants, C_l is the capacity at link l , d_l is the average round trip time of the *flows* passing through link l ,

$$y_l(t) = \sum_{r:l \in r} x_r(t - \tau_{rl}), \quad (3)$$

and $q_l(t)$ is the queue size at time t . Here we use the notation $a = (b)_c^+$ to mean that $a = 0$ if $b < 0$ and $c \leq 0$, otherwise $a = b$. In the original RCP model [5], $\alpha_l = \alpha$ and $\beta_l = \beta$ for each $l \in J$ for some α and β . We have allowed these constants to vary between links in order to study their effect on stability.

Model for the queue. The basis of our investigation is to study a regime where the buffers are so small that it is no longer possible to explicitly model the queue as a saturated integrator. This assumption simply expresses the idea that with small enough buffers, we may use the approximation

$$q_l(t) = p_l(y_l(t)), \quad (4)$$

where $p_l(\cdot)$ is a continuously differentiable function representing the *mean queue length* of link l . This is consistent with the observation in [25] that in a small buffer regime, it is the distribution of the queue size that plays the prominent role in the dynamics of the congestion control framework. At the level of interest in this paper, we do not motivate any explicit functional form for the mean queue length. Different functional forms may be suitable candidates, but our primary focus is to investigate some stability properties of the dynamical system defined by (1-4).

2.2 Global stability without propagation delays

In this subsection, we seek to investigate the stability properties of RCP without taking into account the effects of propagation delays. This gives us some insight into the stability of the algorithm in general, and is also a plausible model for the special case when propagation delays are small in comparison to the update step size of the algorithm. In terms of the fluid model, this would mean

$$\max_{l \in J} \frac{\alpha_l}{d_l C_l} + \frac{\beta_l q_l(t)}{d_l^2 C_l} \ll \min_{r \in S} \frac{1}{\tau_r}.$$

To model RCP without the effects of delays, we use (1-4), setting $\tau_r = 0$ for all $r \in S$, but leaving all other parameters fixed.

Ideally we would like to show that our delay-free model of RCP, at equilibrium, is globally stable. Unfortunately, as the following example demonstrates, we cannot expect the equilibrium points of (1-4) to be unique, or even isolated.

Consider a network consisting of two identical links, l, j and one route $r = \{l, j\}$. Then, since $p_l(\cdot)$ is strictly increasing, there exists a unique y such that

$$\frac{\alpha_l}{d_l C_l} (C_l - y) - \frac{\beta_l p_l(y)}{d_l^2 C_l} = 0.$$

By inspection, there are equilibrium points at $R = (y, y')$ and $\bar{R} = (y', y)$ for all $y' \geq y$.

Alternatively, consider the situation where l and j are not completely identical, and $C_j > C_l$. Then, there exists a unique y' such that

$$\frac{\alpha_j}{d_j C_j} (C_j - y') - \frac{\beta_j p_j(y')}{d_j^2 C_j} = 0.$$

Now, $y' > y$, so, if $x_r \leq y$ then \dot{R}_j will be strictly positive. However, if $x_r > y$ then \dot{R}_l will be strictly negative. Thus, for this example, no equilibrium point exists.

So, there may be an entire manifold of equilibrium vectors for R , or there may be no equilibrium point at all. However, the same is not true for the equilibrium vector of flow rates x .

Theorem 2.1. *Suppose that $x(t)$ evolves according to (1-4), with $\tau_r = 0$ for all $r \in S$. Then, there exists a unique vector \bar{x} such that $x(t) \rightarrow \bar{x}$ as $t \rightarrow \infty$.*

Proof. We prove this result recursively for the more general system where, for all $l \in J$, we replace (2) with

$$\dot{R}_l(t) = R_l(t) \left(u_l(t) + \frac{\alpha_l}{d_l C_l} (C_l - y_l(t)) - \frac{\beta_l p_l(y_l(t))}{d_l^2 C_l} \right)_{R_l(t)}^+, \quad (5)$$

where $u_l(t) \rightarrow 0$ as $t \rightarrow \infty$.

Let N_l be the number of $r \in S$ such that $l \in r$. We set \bar{y}_l equal to the unique value such that $\dot{R}_l(t) = 0$ when $y_l(t) = \bar{y}_l$, and set $\bar{R} = \min_{l \in J} \bar{y}_l / N_l$. We shall show that, for all $r \in S$ with $\bar{R} = \bar{y}_l / N_l$ for some $l \in r$, $x_r \rightarrow \bar{R}$ as $t \rightarrow \infty$. This allows us to remove such an $x_r(t)$ from the system, by replacing it with \bar{R} plus a vanishing term which we incorporate into the $u_l(t)$, for each $l \in r$. Since this will always remove at least one r from S , this is sufficient to prove our result.

It remains to show that $x_r \rightarrow \bar{R}$ for $r \in S$ with $\bar{R} = \bar{y}_l / N_l$, for some $l \in r$. Now, for any $\epsilon > 0$ there exists a T such that

$$|u_l(t)| < \epsilon \frac{\alpha_l}{d_l C_l},$$

for all $l \in J$, $t > T$. From (5) we see that, if, for $t > T$, for any $l \in J$, $R_l(t) < (\bar{y}_l / N_l) - 2\epsilon$, then $\dot{R}_l(t) > \delta$ for some $\delta > 0$. Thus, for some T' , for all $t > T'$, $l \in J$, $R_l(t) \geq (\bar{y}_l / N_l) - 2\epsilon$.

Let us assume that ϵ is small enough that, for any $j \in J$ with $(\bar{y}_j / N_j) > \bar{R}$,

$$\frac{\bar{y}_j}{N_j} - \bar{R} > \left(2 + 2 \max_{l \in J} N_l \right) \epsilon.$$

Now, let $r \in S$ be such that $\bar{R} = \bar{y}_l/N_l$ for some $l \in r$. Let r' be the set of $l \in r$ with $\bar{R} = \bar{y}_l/N_l$. Suppose that, for some $t > T'$,

$$\min_{l \in r'} R_l(t) > \bar{R} + 2\epsilon \max_{j \in J} N_j.$$

If $x_r(t) = R_j(t)$ and $j \in r'$ then $x_r(t) > \bar{R} + 2\epsilon \max_{j \in J} N_j$. Otherwise $x_r(t) = R_j(t)$ with $j \notin r'$, but since $t > T'$, $x_r > \bar{R} + 2\epsilon \max_{j \in J} N_j$ still holds. For any $l \in r'$, $x_r(t)$ is sufficiently high that $y_l(t)$ must be greater than $N_l \bar{R} + 2\epsilon$. Thus, $\dot{R}_l(t) < -\delta$ for some $\delta > 0$. Therefore, for some T'' , for all $t > T''$,

$$\min_{l \in r'} R_l(t) \leq \bar{R} + 2\epsilon \max_{j \in J} N_j.$$

Hence, for all $t > T''$,

$$|x_r(t) - \bar{R}| \leq 2\epsilon \max_{j \in J} N_j.$$

Since ϵ was arbitrary, the result follows. \square

Note, the vector \bar{x} will be close to that of the max-min fair allocation of flow rates, because of its construction. However, the presence of the function $p_l(\cdot)$ in (2) means that link capacity will not be fully utilised at equilibrium.

2.3 Local stability with propagation delays

In this subsection, we derive conditions for the local stability of (1-4) when propagation delays are present. Our result relies upon the assumption that there is only one bottleneck link along each route, that is for each $r \in S$ there is only one $l \in r$ such that $x_r(t) = R_l(t)$ at equilibrium.

For each $r \in S$, $l \in J$, we let \bar{x}_r be the equilibrium value of $x_r(t)$ and \bar{R}_l be the maximum of \bar{x}_s for all $s \in S$ such that $l \in s$. For each $l \in J$, we let \bar{y}_l be the unique value such that $\dot{R}_l(t) = 0$ whenever $y_l(t) = \bar{y}_l$. We can assume, without loss of generality, that for all $l \in J$,

$$\bar{y}_l = \sum_{r: l \in r} \bar{x}_r.$$

Otherwise, assuming the system is always local to equilibrium, $R_l(t)$ will simply continually increase. If $R_l(t)$ is initially large enough then $\max_{r \in S} x_r(t) < R_l(t)$ for all time and thus, $R_l(t)$ has no effect on the rest of the system, and can be ignored.

Now, we have assumed that, for all $r \in S$, $l \in r$, if $\bar{x}_r = \bar{R}_l$, then $\bar{R}_j > \bar{R}_l$ for all $j \in r$, $j \neq l$. Furthermore, by definition, if $\bar{x}_r \neq \bar{R}_l$, we must have $\bar{x}_r < \bar{R}_l$. Thus, for all $r \in S$, $l \in r$, whenever the system is close to equilibrium, either $x_r(t) = R_l(t - \tau_{lr})$, or else $\bar{x}_r < \bar{R}_l$ and $R_l(t - \tau_{lr})$ has no effect on $x_r(t)$. So, we can isolate each $l \in J$ and model the effect of the rest of the system on $R_l(t)$ as a vanishing perturbation. This allows us to find conditions for stability for the overall system using a recursive argument.

Accordingly, we initially restrict our attention to the single link case, $J = \{l\}$, and we consider the following generalisation of (2),

$$\dot{R}_l(t) = -R_l(t) \left(f_l \left(y_l(t) - N_l \bar{R}_l \right) + u_l(t) \right)_{R_l(t)}^+, \quad (6)$$

where $y_l(t)$ is defined as in (3), $f_l(\cdot)$ is an increasing differentiable function with $f_l(0) = 0$ and $u_l(t) \rightarrow 0$ as $t \rightarrow \infty$. When the recursive argument is complete, the vector $u(t)$ represents the behaviour of $R_j(t)$ for $j \in J$ such that $\bar{R}_j < \bar{R}_l$.

Theorem 2.2. *Consider the case where the network consists of only a single link, $J = l$, with flow rates following (6). If*

$$f'_l(0) \bar{R}_l \sum_{r \in S} \tau_r < 1,$$

then for all ϵ , there exists a and u such that, if $|R_l(t) - \bar{R}_l| \leq a$ for all $t \leq 0$ and $|u_l(t)| < u$ for all t , then $|R_l(t) - \bar{R}_l| \leq \epsilon$ for all t , and $R_l(t) \rightarrow \bar{R}_l$ as $t \rightarrow \infty$.

Proof. Let τ_l be $\max_{r \in S} \tau_r$, the maximum round trip time.

Suppose for some t , $\gamma \leq 1$, a' and u' ; $|R_l(t) - \bar{R}_l| = \gamma a'$, $|R_l(t') - \bar{R}_l| \leq a'$ and $|u_l(t')| < u' a'$ for all $t' \in [t - 2\tau_l, t]$. Then, for all $t' \in [t - \tau_l, t]$,

$$\left| \dot{R}_l(t') \right| \leq (\bar{R}_l + a') \left(u' a' + f'_l(0) |S| a' + o(a') \right).$$

Thus, for all $r \in S$,

$$|R_l(t - \tau_r) - R_l(t)| \leq \tau_r \bar{R}_l f'_l(0) |S| a' + \tau_r (\bar{R}_l + a') a' u' + o(a'). \quad (7)$$

So,

$$\begin{aligned} \frac{\dot{R}_l(t)}{R_l(t)} &= f'_l(0) (|S| \bar{R}_l - y_l(t)) - u_l(t) + o(a') \\ &= f'_l(0) |S| (\bar{R}_l - R_l(t)) - u_l(t) + f'_l(0) \sum_{r \in S} \left(R_l(t) - R_l(t - \tau_r) \right) + o(a') \\ &= f'_l(0) |S| \gamma a' + \delta, \end{aligned}$$

where

$$\begin{aligned} |\delta| &\leq f'_l(0) \bar{R}_l \sum_{r \in S} \tau_r f'_l(0) |S| a' + |S| \tau_l (\bar{R}_l + a') a' u' + o(a') \\ &= f'_l(0)^2 |S|^2 \bar{R}_l d_l a' + |S| \tau_l (\bar{R}_l + a') a' u' + o(a'). \end{aligned}$$

Since $f'_l(0) |S| \bar{R}_l d_l < 1$, there exists a value of $a < \epsilon$, u' and $\gamma < 1$ such that, if $a' \leq a$, then $|\delta|$ is guaranteed to be less than $\gamma f'_l(0) |S| a'$. In which case, $\dot{R}_l(t)$ must have the same sign as $\bar{R}_l - R_l(t)$.

So, we can take $u = u'a$, with a as given above, and if $|R_l(t) - \bar{R}_l| \leq a$ for all $t \leq 0$ and $|u_l(t)| < u$ for all t , then $|R_l(t) - \bar{R}_l| \leq \epsilon$ for all t .

Furthermore, for all $a' < a$, there exists a time T such that, for all $t > T$, $|u_l(t)| < u'a'$. From the above analysis, we know that after T , if $R_l(t)$ is beyond a' of \bar{R}_l , then it will converge to \bar{R}_l at a rate of γ . Thus, there is some $T' > T$ such that, $|R_l(t) - \bar{R}_l| \leq a'$ for all $t > T'$. Since a' was arbitrary, $R_l(t) \rightarrow \bar{R}_l$ as required. \square

We now consider the general network case.

Theorem 2.3. *Under (1-4), the equilibrium point \bar{R} is locally asymptotically stable provided that, for each $l \in J$,*

$$\left(\frac{\alpha_l}{d_l C_l} + \frac{\beta_l p_l'(\bar{y}_l)}{d_l^2 C_l} \right) \bar{R}_l \sum_{r: l \in r, \bar{R}_l = \bar{x}_r} \tau_r < 1. \quad (8)$$

Proof. Let δ be the minimum of $|\bar{R}_l - \bar{R}_j|/2$ for $l \neq j \in J$. Now, if, for all t , all $l \in J$, $|R_l(t) - \bar{R}_l| < \delta$, then each $l \in J$ will evolve according to

$$\dot{R}_l(t) = -R_l(t) \left(f_l(y_l^{eq}(t) - \bar{y}_l^{eq}) + u_l(t) \right), \quad (9)$$

where $u_l(t)$ represents the effect of $R_j(t)$ for all $\bar{R}_j < \bar{R}_l$, and

$$y_l^{eq}(t) = \sum_{r: l \in r, \bar{x}_r = \bar{R}_l} x_r(t - \tau_{rl}),$$

with \bar{y}_l^{eq} equal to the equilibrium value of $y_l^{eq}(t)$. For each $l \in J$, we have

$$f_l(w) = \frac{\alpha_l}{d_l C_l} (w + \bar{y}_l - C_l) + \frac{\beta_l p_l(\bar{y}_l + w)}{d_l^2 C_l}.$$

By definition, (9) is an example of (6), for the reduced network formed by l and all routes $r \in S$ with $l \in r$ and $\bar{x}_r = \bar{R}_l$. Furthermore, (8) is precisely the condition of Theorem 2.2 for this reduced system.

Suppose $|R_l(t) - \bar{R}_l| < \delta$ for all $l \in J$, for all t . By differentiability of $p_l(\cdot)$, for any u , we can find ϵ such that if $|R_j(t) - \bar{R}_j| < \epsilon$ for all t , for all j such that $\bar{R}_j < \bar{R}_l$, then $|u_l(t)| < u$ for all t . This allows us to prove local stability recursively.

We begin with the $l \in J$ such that \bar{R}_l is maximal. We apply Theorem 2.2 to find conditions under which $R_l(t)$ converges to \bar{R}_l and $|R_l(t) - \bar{R}_l| < \delta$ for all t . This gives us a and u , where the initial conditions of $R_l(t)$ should be within a of \bar{R}_l , and $u_l(t)$ should tend to zero and always be bounded by u . We can find an $\epsilon < \delta$ so that if $|R_j(t) - \bar{R}_j| < \epsilon$ for all t and $R_j(t)$ converges to \bar{R}_j for all $j \neq l$, then $|u_l(t)| < u$ for all t and $u_l(t)$ converges to 0. So, we set $R_l(t) \in [\bar{R}_l - a, \bar{R}_l + a]$ for $t \leq 0$ as our initial condition for l , remove l from the network and repeat this process for $\delta' = \epsilon$. Since we remove one link each time, eventually we will find suitable initial conditions for the entire network. \square

Note, each link does not necessarily need to keep track of which flows are under its control in order to meet (8). The condition holds if, for all $l \in J$,

$$\left(\frac{\alpha_l}{d_l} + \frac{\beta_l p'_l(\bar{y}_l)}{d_l^2} \right) \frac{\bar{y}_l}{C_l} d_l^p < 1,$$

where d_l^p is the average per *packet* round trip time,

$$d_l^p = \frac{1}{\bar{y}_l} \sum_{r:l \in r} \bar{x}_r \tau_r.$$

This suggests that it should be d_l^p not d_l that appears in the RCP controller (2). If, for each $l \in J$ we let γ_l be such that $p'_l(\bar{y}_l) = \gamma_l p_l(\bar{y}_l) / \bar{y}_l$ then,

$$p'_l(\bar{y}_l) = \gamma_l \frac{\alpha_l d_l}{\beta_l} \frac{C_l - \bar{y}_l}{\bar{y}_l} < \frac{\gamma_l \alpha_l d_l}{\beta_l},$$

thus, (8) is ensured if, for all $l \in J$,

$$\alpha_l = \frac{d_l^p}{d_l(1 + \gamma_l)}. \quad (10)$$

This parameter choice scheme is attractive, because it is fairly decentralised and only requires local information. Each link needs only to measure, estimate or be informed of the round trip times for packets passing through that link. However, these results rely on the weak assumption that each route has only *one* bottleneck link. However, any network with multiple bottleneck routes can easily be transformed into one without, for example, by making a small change in the parameters β_l , $l \in J$. Preliminary numerical results suggest that Theorem 2.3 may hold in general; but finding an analytic result to confirm this, or a counter-example to disprove it, remains an open problem.

In the next section we investigate the impact of loss of local stability in RCP.

3 Local bifurcation analysis

A key focus in the choice of parameters for any congestion control proposal is to ensure that they lead to a stable equilibrium. Most proposals for congestion control, for example see [21], lead to the analysis of non-linear time delayed dynamical systems. For such non-linear systems, typically sufficient conditions for local stability guide parameter choices. Following local stability, a local bifurcation theoretic analysis can make us comfortable in running the system close to the *edge* of stability.

We first recapitulate a result about the loss of local stability in a non-linear retarded functional differential equation and then use it to analyse the RCP fluid model. Following the analysis in [19] we outline a local Hopf bifurcation result for the following non-linear delay equation

$$\dot{u}(t) = \eta(-\xi_x u(t - \tau) \pm \xi_{xy} u(t) u(t - \tau)), \quad (11)$$

where $\eta, \tau, \xi_x, \xi_{xy} > 0$. The parameter η has been intentionally introduced to just tip the above equation over the edge of (local) stability. This exogenous, non-dimensional, parameter will *act* as our bifurcation parameter.

We may state the following about equation (11).

Theorem 3.1. *A necessary and sufficient condition for local stability is*

$$\eta\xi_x\tau < \pi/2,$$

and treating η as the bifurcation parameter, the first local Hopf bifurcation occurs with period 4τ at $\eta = \eta_c$, where $\eta_c\xi_x\tau = \pi/2$. Further, as the Hopf condition is just violated, the equation will *always* undergo a super-critical Hopf bifurcation where the amplitude of the stable bifurcating solutions will be proportional to

$$\frac{\xi_x}{\xi_{xy}} \sqrt{\frac{20\pi(\eta - \eta_c)}{3\pi - 2}}.$$

3.1 RCP: single link, single delay model

In our bifurcation analysis, for the sake of simplicity, we shall leave the queuing term out of the model by taking $\beta = 0$. Consider the following single link, single delay, RCP fluid model

$$\dot{R}(t) = \eta R(t) \left(\frac{\alpha}{C\tau} (C - y(t)) \right), \quad (12)$$

where $y(t) = \sum_s R(t - \tau)$ is the aggregate load at the link and η is the non-dimensional bifurcation parameter. Define $u(t) = R(t) - \bar{R}$, and take a Taylor expansion of (12) to obtain

$$\dot{u}(t) = -\eta \frac{\alpha}{\tau} u(t - \tau) - \eta \frac{\alpha}{\tau \bar{R}} u(t) u(t - \tau).$$

Now using Theorem 3.1, we may state the following about equation (12). A necessary and sufficient condition for local stability is

$$\eta\alpha < \pi/2,$$

and treating η as the bifurcation parameter, the first local Hopf bifurcation occurs with period 4τ at $\eta = \eta_c$, where $\eta_c\alpha = \pi/2$. If the Hopf condition is just violated, the equation will *always* undergo a super-critical Hopf bifurcation where the amplitude of the stable bifurcating solutions will be proportional to

$$\bar{R} \sqrt{\frac{20\pi(\eta - \eta_c)}{3\pi - 2}}.$$

We highlight that equation (12) cannot produce a sub-critical Hopf bifurcation. However, in the bifurcation analysis we have omitted any additional non-linear effects that would arise if $\beta > 0$ in the RCP model.

We now summarise our contribution in this paper, and outline some avenues for further research.

4 Conclusion

It is noteworthy to observe that the small buffer regime has allowed us to tackle the question of stability for a large network with an arbitrary topology. First, we proved that the RCP fluid model is globally stable in the absence of propagation delays. Then, we derive a set of conditions for local stability when arbitrary heterogeneous propagation delays are present. The network delay stability result relies upon the weak assumption that, at equilibrium, there is only one bottleneck link along each route. An interesting avenue for research would be to show that either this result may hold in general, or provide a counter-example to disprove it. Finally in a single link, single delay, model we investigate the impact of the loss of local stability in a special case (setting $\beta = 0$) where we show that the RCP algorithm would *always* give rise to an innocuous looking super-critical Hopf bifurcation.

As RCP aims for max-min fairness it is appropriate to first consider a model that embodies the original formulation, as we did in this paper. We now outline some natural avenues for further research.

Fairness and Stability. A conveniently parameterized family of α -fair rate allocations was introduced in [17]. The parameter α lies in the range $(0, \infty)$, and the cases $\alpha \rightarrow 0$, $\alpha = 1$ and $\alpha \rightarrow \infty$ correspond respectively to an allocation which achieves maximum throughput, is proportionally fair or is max-min fair [17]. In this paper we have only considered a max-min fair allocation mechanism, as has been originally specified [5]. An immediate direction of further research would be to incorporate different notions of fairness into the RCP framework, and analyse the stability of such networks.

Impact of the RCP parameters: α_l and β_l . We have only considered the limiting case of very small buffers. The choice of β impacts the rate at which the queue is drained. Small values of β drain the queue slowly; so with small β and for large enough buffers it is appropriate to model the queue as a saturated integrator. However, for large values of β , the queue may drain fast enough so that at the timescale of operation of the congestion control protocols, finer queueing theoretic models may have to be developed. Such queueing models would certainly contribute to the non-linearity in the RCP dynamical system, and hence impact the stability of the time delayed network. Further, the results from local bifurcation theory could also be subtle. In our bifurcation theoretic analysis we set $\beta = 0$, thus removing any non-linearity that may arise from the queue.

Acknowledgements

The authors are grateful to Frank Kelly for comments on earlier drafts and acknowledge funding provided by the EPSRC grant GR/S86266/01. The usual caveat applies.

References

- [1] G. Appenzeller, I. Keslassy, and N. McKeown, “Sizing router buffers,” *Computer Communication Review*, vol. 34, no. 4, pp. 281-292, 2004.
- [2] H. Balakrishnan, N. Dukkupati, N. McKeown, and C. Tomlin, “Stability Analysis of Explicit Congestion Control Protocols,” Stanford University Department of Aeronautics and Astronautics Report: SUDAAR 776, 2005. [online] <http://yuba.stanford.edu/rcp/>
- [3] S. Deb and R. Srikant, “Rate-based versus Queue-based models of congestion control,” *Proceedings of ACM Sigmetrics*, 2004.
- [4] N. Dukkupati and N. McKeown, “Why Flow-Completion Time is the Right Metric for Congestion Control,” *Computer Communication Review*, vol. 36, no. 1, pp. 59-62, 2006.
- [5] N. Dukkupati, M. Kobayashi, R. Zhang-Shen, and N. McKeown, “Processor Sharing Flows in the Internet,” *Thirteenth International Workshop on Quality of Service (IWQoS)*, 2005. [online] <http://yuba.stanford.edu/rcp/>
- [6] N. Dukkupati, N. McKeown, and A. G. Fraser, “RCP-AC: Congestion Control to make flows complete quickly in any environment,” *High-Speed Networking Workshop: The Terabits Challenge*, (in Conjunction with IEEE Infocom), 2006.
- [7] M. Enachescu, Y. Ganjali, A. Goel, N. McKeown, and T. Roughgarden, “Part III: routers with very small buffers,” *Computer Communication Review*, vol. 35, no. 3, pp. 83-90, 2005.
- [8] H. Han, C. V. Hollot, Y. Chait, and V. Misra, “TCP Networks Stabilized by Buffer-Based AQMs,” *Proceedings of IEEE Infocom*, 2004.
- [9] H. Han, C. V. Hollot, D. Towsley, and Y. Chait, “Synchronization of TCP Flows in Networks with Small DropTail Buffers,” *Proceedings of IEEE Conference on Decision and Control*, 2005.
- [10] C.V. Hollot, V. Misra, D. Towsley, and W. Gong, “Analysis and design of controllers for AQM routers supporting TCP flows,” *IEEE Transactions on Automatic Control*, vol. 47, no. 6, pp. 945-959, 2002.
- [11] D. Katabi, M. Handley, and C. Rohrs, “Internet Congestion Control for Future High Bandwidth-Delay Product Environments,” *Proceedings of ACM Sigcomm*, 2002.
- [12] S.S. Kunniyur, and R. Srikant, “Stable, scalable, fair congestion control and AQM schemes that achieve high utilization in the Internet,” *IEEE Transactions on Automatic Control*, vol. 48, no. 11, pp. 2024-2028, 2003.

- [13] R.J. La, "Instability of a Tandem network and its propagation under RED," *IEEE Transactions on Automatic Control*, vol. 49, no. 6, pp. 1006-1011, 2004.
- [14] A. Lakshmikantha, C. Beck, and R. Srikant, "Robustness of Real and Virtual Queue based Active Queue Management Schemes," *IEEE Transactions on Networking*, vol. 13, no. 1, pp. 81-93, 2005.
- [15] S. Liu, T. Basar, and R. Srikant, "Exponential-RED: a stabilizing AQM scheme for low- and high-speed TCP protocols," *IEEE Transactions on Networking*, vol. 13, no. 5, pp. 1068 - 1081, 2005.
- [16] P. Marbach, "Priority Service and Max-Min Fairness," *IEEE Transactions on Networking*, vol. 11, pp. 733-746, 2003.
- [17] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE Transactions on Networking*, vol. 8, no. 5, pp. 556-567, 2000.
- [18] B. Radunovic and J.-Y. Le Boudec, "Rate performance objectives of multihop wireless networks," *IEEE Transactions on Mobile Computing*, vol. 3, no. 4, pp. 334-349, 2004.
- [19] G. Raina, "Local Bifurcation Analysis of Some Dual Congestion Control Algorithms," *IEEE Transactions on Automatic Control*, vol. 50, no. 8, pp. 1135 - 1146, 2005.
- [20] G. Raina and D. Wischik, "Buffer sizes for large multiplexers: TCP queueing theory and instability analysis," *Proceedings of EuroNGI*, 2005.
- [21] R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.
- [22] L. Tan, X. Zhang, L. Andrew, and M. Zukerman, "Price-based Max-Min Fair Rate Allocation in Wireless Multi-hop Networks," *IEEE Communications Letters*, vol. 10, no.1, pp. 31-33, 2006.
- [23] G. Vinnicombe, "On the stability of networks operating TCP-like congestion control," *Proceedings of International Federation of Automatic Control*, 2002.
- [24] T. Voice, "A Global Stability Result For Primal-Dual Congestion Control Algorithms With Routing," *Computer Communication Review*, vol. 34, no. 3, pp. 35-41, 2004.
- [25] D. Wischik and N. McKeown, "Part I: buffer sizes for core routers," *Computer Communication Review*, vol. 35, no. 3, pp. 75-78, 2005.