The Stability and Convergence Time of Rate Control Protocol (RCP)

I. INTRODUCTION

The purpose of this report is two fold: first, it summarizes what we already know on the stability of the Rate Control Protocol (RCP) [1] and second, it describes simulation results, intuition and hypotheses on the open problems in RCP stability and time to converge. The eventual goal is to solve these open problems using techniques in control theory.

II. BACKGROUND: EQUATIONS DESCRIBING RCP SYSTEM

The following equations describe the RCP system for a single bottleneck link carrying flows with heterogeneous round-trip times.

$$
\begin{aligned}
\dot{R}(t) &= R(t)\left(\frac{\alpha(C-y(t))-\beta\frac{q(t)}{\overline{d}}}{C\overline{d}}\right) && (1)\\
y(t) &= \sum_{i=1}^{N} R(t-d_i) && (2)\\
\overline{d} &= \sum_{i=1}^{N} \frac{d_i}{N} && (3)\\
\dot{q}(t) &= [y(t)-C]; \qquad 0<q(t)<B && (4)\\
&= \max[y(t)-C,0]; \quad q(t)=0 \\
&= \min[y(t)-C,0]; \quad q(t)=B
\end{aligned}
$$

where $R(t)$ is the rate being updated by the router, C is the link-capacity, B is the buffer size, $y(t)$ is the aggregate incoming traffic rate, $q(t)$ is the queue occupancy, $\bar{d}$ is the average round-trip time of the flows, d_i is the RTT of flow i and N is the number of flows.

III. PAST RESULTS ON RCP STABILITY

All results so far on RCP stability are for a single bottleneck link with heterogeneous round-trip times. The shaded region in Figure 1 shows the stable region of the linearized system [2] for homogenuous RTTs. This region is independent of the link-capacity, RTT and the number of flows. The region obtained analytically by taking care of queue non-linearity turns out to be larger than the linearized region [3] — it is the region to the left of the dotted line in Figure 1. In the case of a single link with heterogeneous round-trip times, there is a large enough region of RCP parameters to ensure a stable and robust system independent of link-capacity, RTTs and the number of flows.

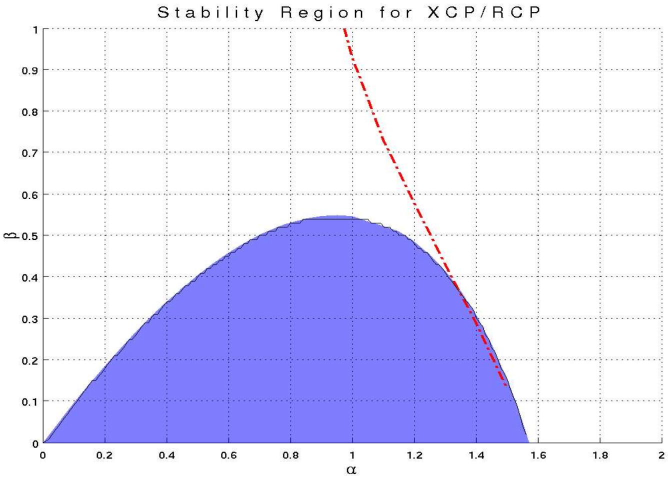

Fig. 1. Region of parameters for which RCP system is stable: 1) shaded region shows the linearized region and 2) region to the left of the dotted line shows the stable region of the non-linear system (i.e. taking care of the queue non-linearity).

IV. WHAT WE DO NOT KNOW YET: RCP STABILITY IN NETWORK SCENARIO

We do not yet have analytical results on RCP stability for the network scenario. Simulations with multiple congested links indicate that the stable region in the network case is not different from that of the single link case of Figure 1.

V. WHAT WE DO NOT KNOW YET: CONVERGENCE TIMES IN RCP

We are interested in bounding the convergence time of RCP when there is a sudden increase in the offered network load. In the ”typical” case, when there is no such sudden load increase, RCP behaves well in the sense that flows complete quickly, link utilization is high, queue occupancy is small and it is inherently fair. A sudden load increase is the worst-case traffic scenario for RCP and is a direct consequence of the fact that RCP gives a high starting rate to the flows. A natural question that arises is:

If the offered network load increases by k times, how quickly does the RCP algorithm converge? How does the convergence time depend on C, N, flow RTTs, the buffer size at congested link and the parameters α and β?

In this section we present simulations and intuition on the effect of α, β, buffer size and other network parameters on the convergence times of RCP. In each case, we vary one variable from one extreme to the other while keeping all else fixed. The base experiment setup is: link-capacity $C = 1$ Gbps, flow RTTs $= 0.1$s, number of flows $N = 20$,

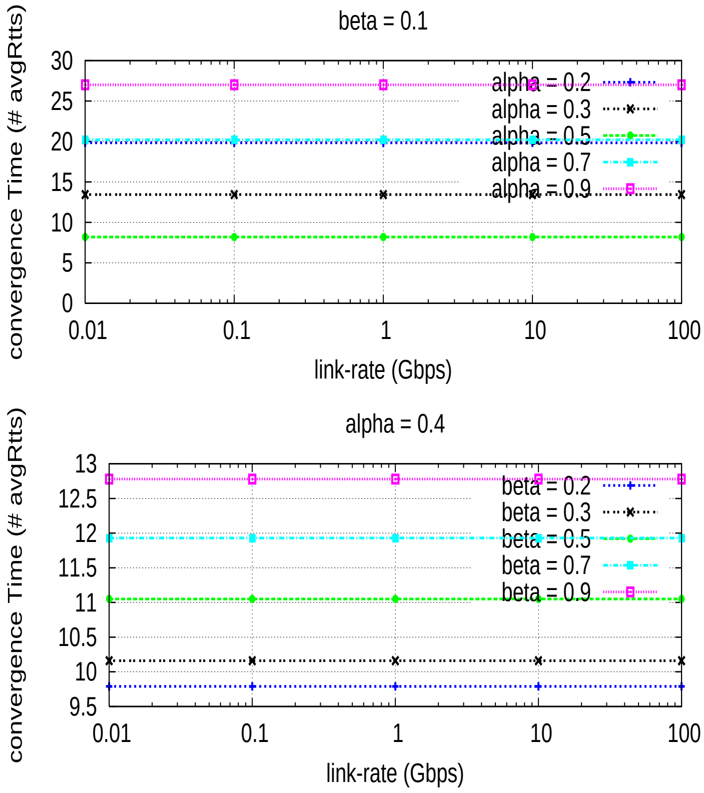

Fig. 2. Figure showing the convergence time (in #RTTs) of RCP versus the link-rate. The convergence time is defined as the time taken for R and q to be within 98% and 99% of their equilibrium values. The top plot is for $\beta = 0.1$ and bottom plot is for $\alpha = 0.4$.

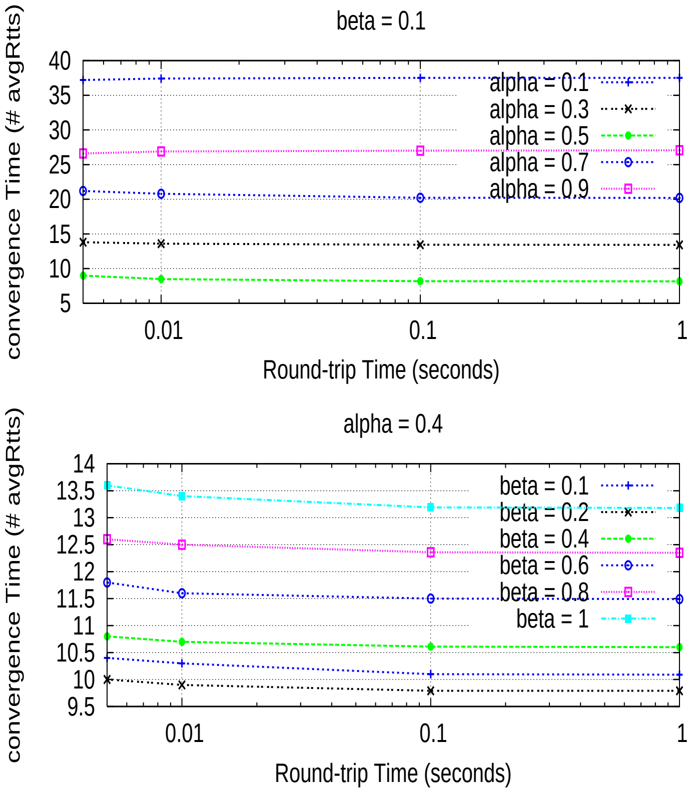

Fig. 3. Figure showing the convergence time (in #RTTs) of RCP versus the link-rate. The convergence time is defined as the time taken for R and q to be within 98% and 99% of their equilibrium values. The top plot is for $\beta = 0.1$ and the bottom plot is for $\alpha = 0.4$.

buffer-size is bandwidth×delay, and the initial conditions are $R(0) = 0.1, q(0) = 0$ corresponding to half the load of the equilibrium (equilibrium is : $R_e/C = 0.05, q_e = 0$). Our observations here are for a single bottleneck link.

Our first observation is that the convergence time — as a measure of the number of round-trip times taken to reach equilibrium — is independent of the link-capacity, the flow round-trip times and the number of flows. Although surprising at first, it is intuitive on a closer look. Let's see why:

A. *As link-rate increases:*

Equations (1-4) were simulated for the base set up with link-rates ranging over four orders of magnitude – from 10 Mbps to 100 Gbps. The simulations were done for $0 \leq \alpha, \beta \leq 1$ in increments of 0.1. Figure 2 shows that the link-rate does not affect the RCP convergence time. This is expected because RCP is designed to scale well with link-rates — note in Equation 1, the increase in R is proportional to the link-rate, C [1].

B. *As round-trip time increases:*

Keeping everything else fixed in the base set up, the flow RTT is now varied from 5 ms to 1 second. The equations were simulated for $0 \leq \alpha, \beta \leq 1$.

Ideally, the number of RTTs to converge to equilibrium should not vary much with the absolute value of the RTT itself. Figure 3 shows an example that this is in fact the case with RCP.

C. *As number of flows increases:*

Does the RCP convergence time depend on the number of flows? To find out, we simulate Equations (1-4) while varying the number of flows from 2 to 20000 flows. In each case, the initial conditions correspond to half the equilibrium load, for example: when $N = 200$, the equilibrium is $R_e/C = 0.005, q_e = 0$ and the initial conditions are chosen as $R(0)/C = 0.01, q(0) = 0$. In other words, the number of flows is different in every experiment but the increase in load is the same in each case (equilibrium load is twice the initial load). Figure 4 shows that RCP's convergence time is almost invariant of N.

D. *As α and β vary:*

Figures 5 and 6 show how long RCP takes to converge (# RTTs) with varying (α, β) when the offered load is doubled. In each of these figures, the top plot shows the time taken for R and q to be within 98% and 99% of their equilibrium values respectively; the middle plot shows the time taken for R and q to be within 95% of the equilibrium; the bottom plot shows the

---

[1] To see this, note that both $R(t)$ and $y(t)$ are $\propto C$, which makes $\dot{R(t)} \propto C$

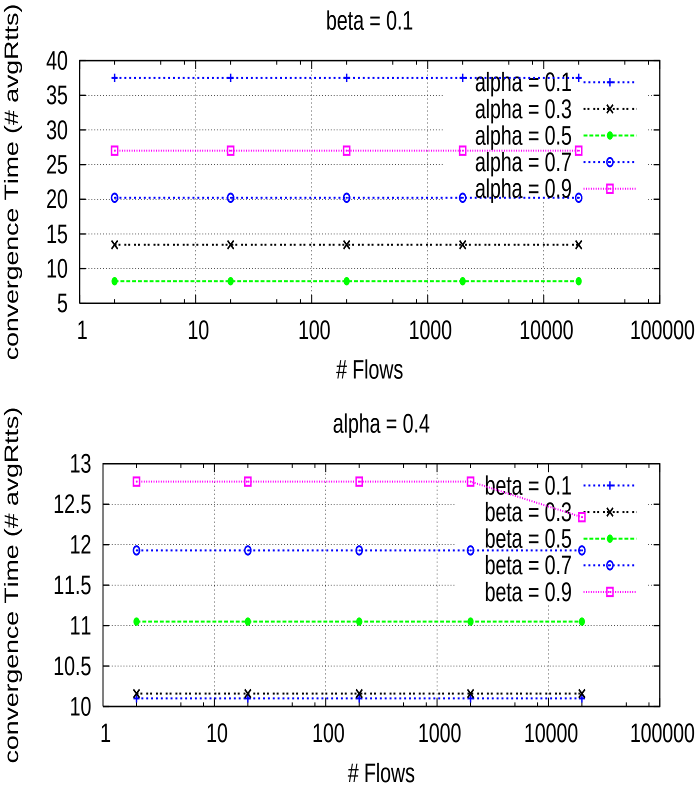

Fig. 4. Figure showing the convergence time (in #RTTs) of RCP versus the number of flows. The convergence time is defined as the time taken for R and q to be within 98% and 99% of their equilibrium values. The top plot is for $\beta = 0.1$ and the bottom plot is for $\alpha = 0.4$.

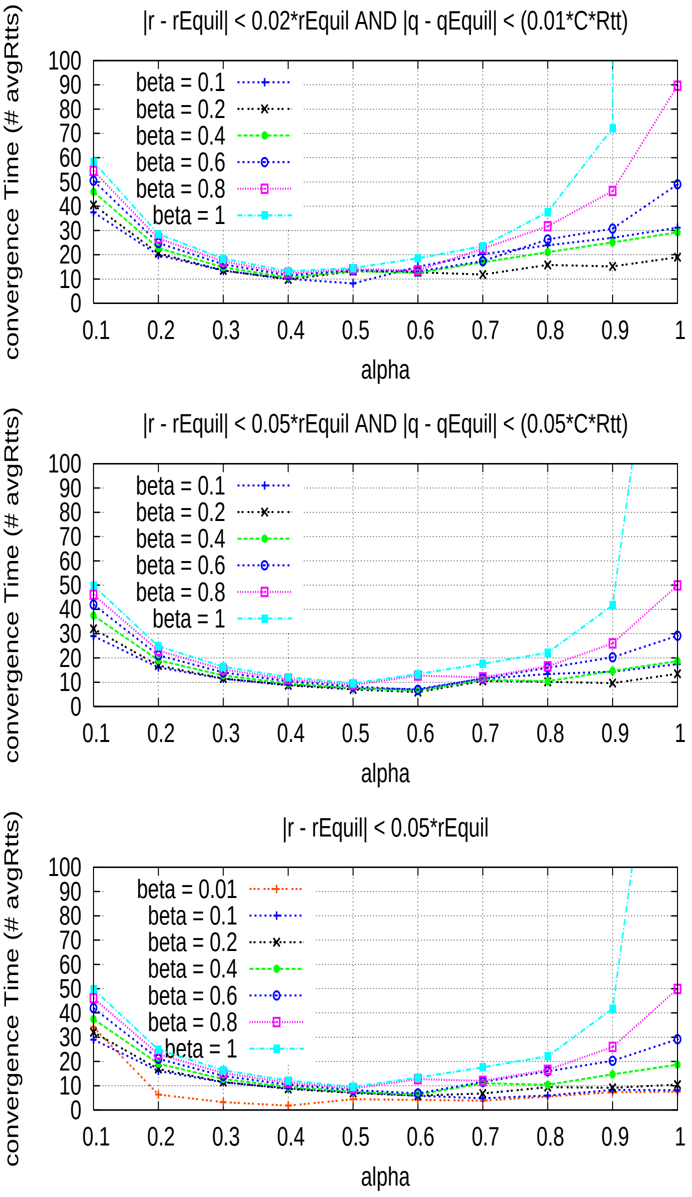

Fig. 5. Figure showing the convergence time (in #RTTs) of RCP versus α. Note that the three plots have slightly different definitions of convergence, as mentioned at the top of each plot.

time taken for R to be within 95% of its equilibrium and we don't care how long the queue takes to drain. The take-away points from these plots are:

1) The border values of alpha and beta, i.e. very small or large values, result in the longest convergence times.
2) As α varies (figure 5): When α is small, the algorithm is slow in catching up with spare capacity, and for large α the rate overshoots multiple times before settling down on the equilibrium. Either way results in long convergence times. Typical R and q trajectories for very small and large α are shown in figures 10, 11 of Appendix I.
3) As beta varies (figure 6): For small β (< 0.1), the queue takes a long time to drain and therefore the long convergence times in top two plots of figure 6. On the other hand if we look at the convergence time exclusively in terms of R (as the third plot does), it decreases monotonically with β.
4) To get a sense of the absolute numbers in convergence time, the minimum observed convergence times are: a) 8.19 RTTs for (R, q) to be within 98% and 99% of equilibrium values when $(\alpha, \beta) = (0.5, 0.1)$ b) 5.92 RTTs for (R, q) to be each within 95% of their equilibrium values when $(\alpha, \beta) = (0.6, 0.2)$. The trajectories for $R(t)$ and $q(t)$ are shown in figure 12 of Appendix I. c) 1.83 RTTs for R to be within 95% of its equilibrium value when $(\alpha, \beta) = (0.4, 0.01)$. $R(t)$ and $q(t)$ are shown in figure 13 of Appendix I.

E. Effect of buffer size

So far the buffer size, B, was fixed at bandwidth×delay product. Figure 7 shows the effect of buffer size on the RCP convergence times under different β values. The convergence times are faster with smaller buffers. With small buffers, $R(t)$, is reduced just the right amount until it reaches the new equilibrium; while with large buffers $R(t)$ is first reduced to below the new equilibrium in order to drain the buffer, and then back up to the equilibrium. In summary, figure 7 gives us the following insights:

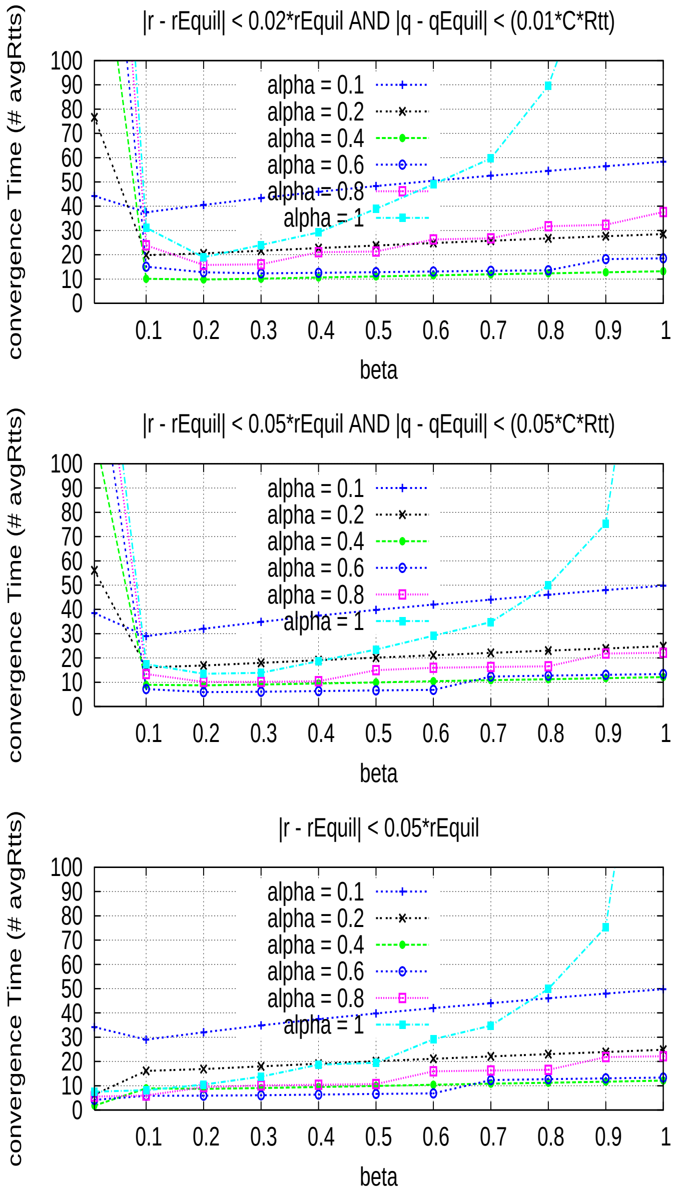

Fig. 6. Figure showing the convergence time (in #RTTs) of RCP versus β. Note that the three plots have slightly different definitions of convergence, as mentioned at the top of each plot.

1) The time to converge is faster with smaller buffers.
2) With small buffers, the β value does not have much of an effect on the convergence time. The points for different β values are all clustered together for smaller buffer sizes.
3) The figure is plotted for $\alpha = 0.4$. The nature of the plot does not change much for other α values — shifts up or down.

F. As the perturbation size increases

So far, we observed the convergence times when the load is doubled. Suppose the load change is an order of magnitude greater than the initial load? In the simulation, this is done by changing the initial conditions relative to the equilibrium of $R_e/C = 0.05$. Figure 8 shows that as the load increases by $2, 5, 10$ and 20 times, the convergence time increases by only a modest amount.

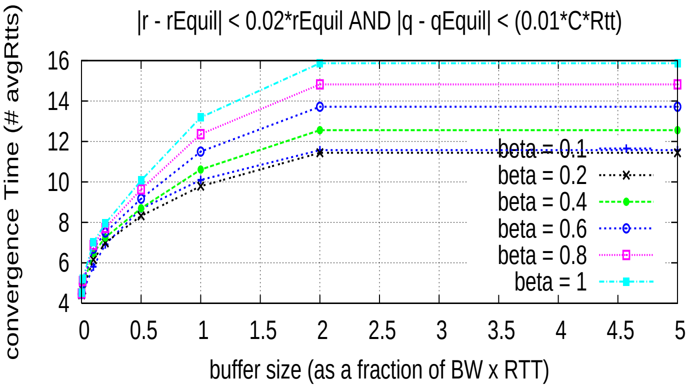

Fig. 7. Convergence time (in #RTTs) versus the buffer size. β varies as shown in the plot while $\alpha = 0.4$.

G. Effect of heterogeneous round-trip times

In this section we observe that the results for the homogeneous RTTs hold true even under heterogeneous flow RTTs. Figure 9 shows the time to converge when 19 flows have RTT of 105 ms while 1 flow has an RTT of 5 ms — the average RTT is equal to that of the base setup i.e. 100 ms. The convergence times are qualitatively similar to the homogeneous RTT case.

APPENDIX I
TRAJECTORIES OF $R(t)$ AND $q(t)$ UNDER VARYING α AND β

REFERENCES

[1] Nandita Dukkipati, Masayoshi Kobayashi, Rui Zhang-Shen, Nick McKeown, ”Processor Sharing Flows in the Internet,” In *Thirteenth International Workshop on Quality of Service (IWQoS)*, Passau, Germany, June 2005.

[2] Nandita Dukkipati, Nick McKeown, ”Processor Sharing Flows in the Internet,” *Stanford University High Performance Networking Group Technical Report TR04-HPNG-061604*, June 2004.

[3] Hamsa Balakrishnan, Nandita Dukkipati, Nick McKeown, Claire J. Tomlin, “Stability Analysis of Explicit Congestion Control Protocols,” Under submission to *IEEE/ACM Transactions on Networking*.

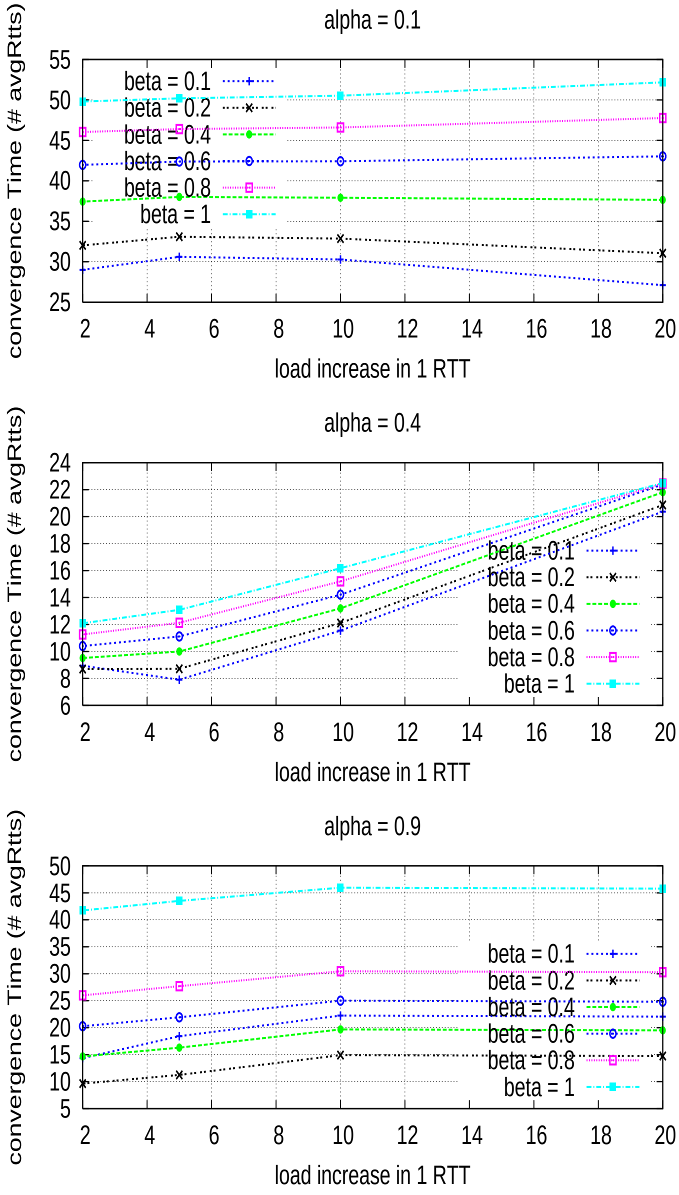

Fig. 8. Convergence time (in #RTTs) versus the the load increase. The three plots show the convergence time for R and q to be within 95% of the equilibrium for different alpha values (as mentioned at the top of each figure).

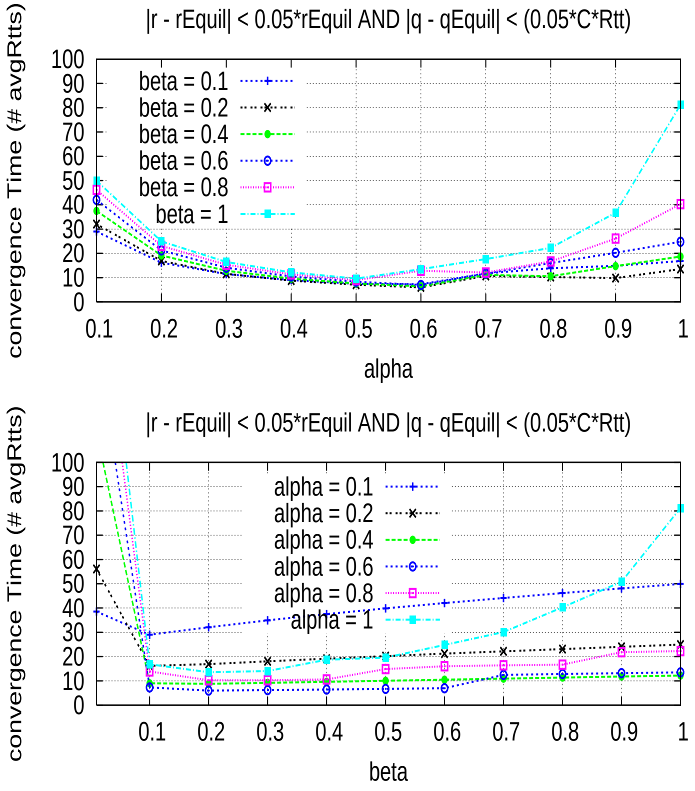

Fig. 9. Convergence time (in #RTTs) versus the the α and β values when flows have heterogeneous round-trip times. The plots show the convergence time for R and q to be within 95% of the equilibrium values.

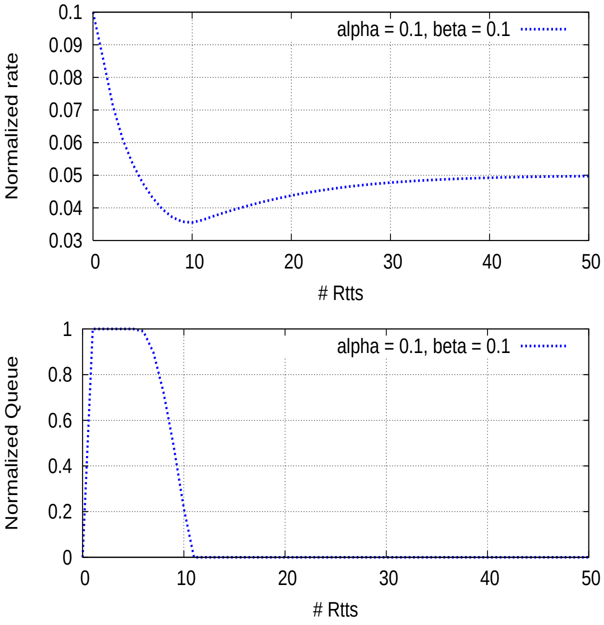

Fig. 10. Figure showing trajectories of R and q for small α values. Choosing a small α takes R a long time to catch up with

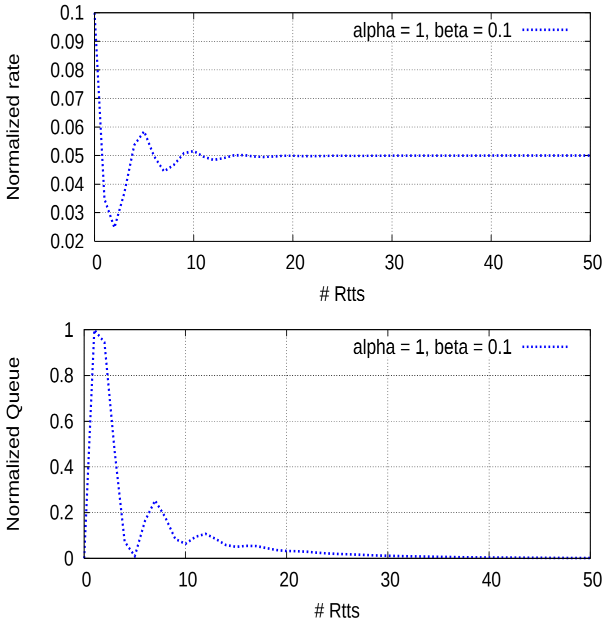

Fig. 11. Figure showing trajectories of R and q for large α values. Choosing a large α overshoots R multiple times before settling down to equilibrium.

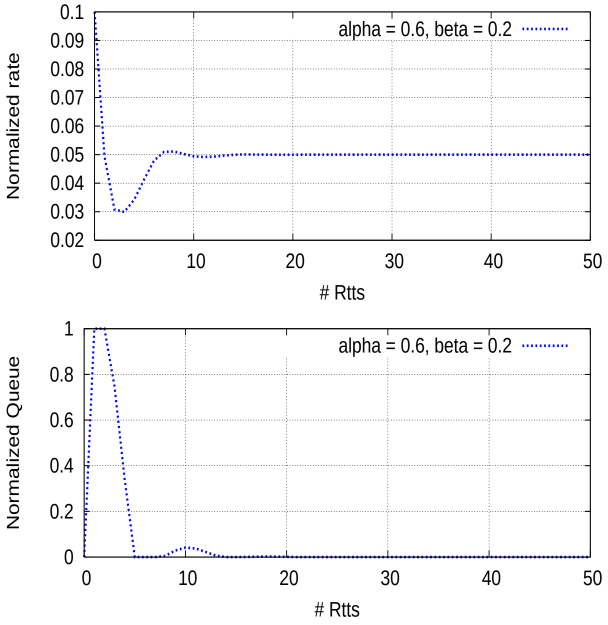

Fig. 12. Figure showing trajectories of R and q for $(\alpha, \beta) = (0.6, 0.2)$. In this case R and q converge to 95% of their equilibrium values within 6 RTTs.

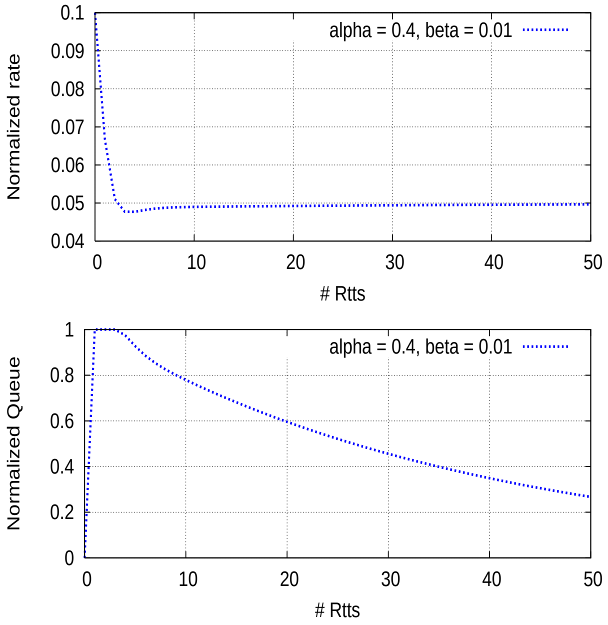

Fig. 13. Figure showing trajectories of R and q for $(\alpha, \beta) = (0.4, 0.01)$. R converges to 95% of its equilibrium value within 2 RTTs. q takes a long time to drain because of small β.