Stability Analysis of Explicit Congestion Control Protocols

Hamsa Balakrishnan, Nandita Dukkipati, Nick McKeown and Claire J. Tomlin

***Abstract*—In the context of explicit congestion control protocols like XCP and RCP where the equilibrium queue lengths are zero, we show that the stability region derived from traditional Nyquist analysis is not an accurate representation of the actual stability region, and that the use of switched linear system models with time delay and new Lyapunov tools can provide sound sufficient stability conditions.**

I. Introduction and motivation

To demonstrate the stability of fluid-flow models of protocols (modifications to TCP such as STCP, FAST and HSTCP as well as new congestion control algorithms, such as the eXplicit Control Protocol (XCP) [5] and the Rate Control Protocol (RCP) [2]), researchers have used a combination of control-theoretic analysis and simulation, often relying on control theory to demonstrate soundness. The existing techniques to analyze such models linearize the system equations about the equilibrium and then use linear system analysis tools such as the Nyquist criterion to find parameters that determine the response to congestion, as well as rate increases, for which the system is stable. The success of such linearized analysis depends on how well the system dynamics can be approximated by its first-order behavior about the equilibrium point. In particular, when the equilibrium point lies on a discontinuity in the system dynamics (often caused by physical constraints, such as nonnegative queue lengths), the stability of the linearized system gives *no* guarantees on the stability of the system, even for simple network topologies such as bottleneck links.

The motivation for this work stems from the system dynamics of XCP and RCP. For a single bottleneck link of capacity C traversed by N flows with equal round trip delays d, aggregate flow rate $y(t)$, and queue length $q(t)$, the system can be modeled by the following delay differential equations, where α and β, chosen through stability analysis, determine the aggregate feedback given $y(t)$ and $q(t)$ [5]:

$$\begin{array}{rcl} \dot{y}(t) & = & -\frac{\alpha}{d}(y(t-d)-C) - \frac{\beta}{d^2}q(t-d) \\ \dot{q}(t) & = & \begin{cases} y(t)-C, & q(t)>0 \\ \max(0, y(t)-C), & q(t)=0 \end{cases} \end{array} \qquad (1)$$

We refer to this system as *SYSTEM 1*. Linear analysis, such as that in [5], would analyze the linearization by considering only one possible mode of behavior of the system, namely,

$$\begin{array}{rcl} \dot{y}(t) & = & -\frac{\alpha}{d}(y(t-d)-C) - \frac{\beta}{d^2}q(t-d) \\ \dot{q}(t) & = & y(t)-C. \end{array}$$

When we compare the stability of this system for different system parameters (α and β) obtained through linear (Nyquist) analysis, with the simulated system for $d = 200$ms (Fig. 1, left), the Nyquist analysis suggests that the shaded region of parameters is stable; simulations, however, suggest that a potentially much larger region, that to the left of the dotted line shown, is stable. If we simulate both the linearization and the switched systems for two sets of parameter values ($\alpha = 0.8$, $\beta = 0.55$ and $\alpha = 1.4$, $\beta = 0.3$), we notice that the first set of parameters results in a stable system; while for the second set of parameters, linear analysis predicts a stable system, while simulations indicate that the system is unstable [1]. The considerable difference in the stable region predicted by linearization and the actual stable region motivates a more careful analysis. We notice that *SYSTEM 1* is a switched system with two modes of operation, one when the queue length is positive, and one when the queue length is zero, and that the equilibrium point lies on the line ($q(t) = 0$) on which the switching between the two systems occurs. It is known that for a switched system, linearizing about an equilibrium point at which the system dynamics are discontinuous could lead one to erroneous conclusions, even on its local stability [4].

We advocate caution in the use of linear stability theory in the analysis of explicit congestion control protocols in which the equilibrium queue lengths are zero. Instead, we propose a method for taking discontinuities in the system dynamics into account by modeling the protocol as a switched system, and present a computational technique to analyze the stability of

H. Balakrishnan is with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139. email: hamsa@mit.edu. N. Dukkipati and N. McKeown are with the Department of Electrical Engineering, Stanford University. email: nanditad,nickm@stanford.edu. C. Tomlin is with the Department of Electrical Engineering and Computer Science, University of California at Berkeley. email: tomlin@eecs.berkeley.edu

Manuscript received January 2007; revised July 2007.

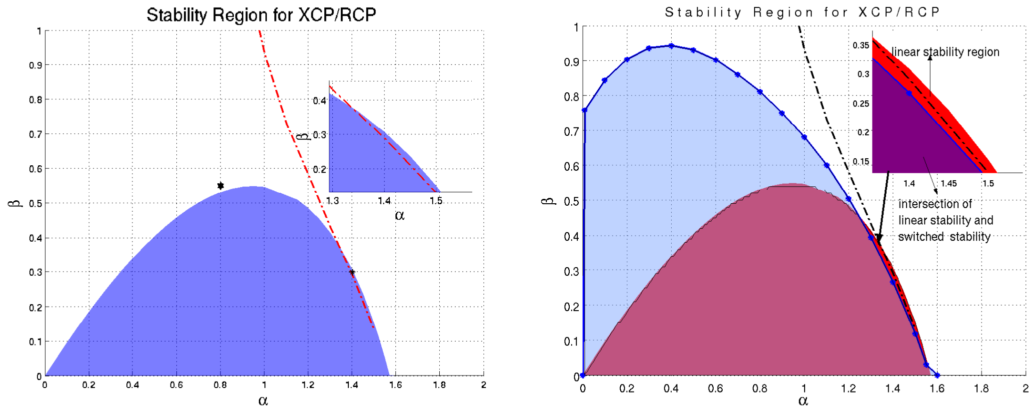

Fig. 1. (Left) Comparison of linearized stability region (shaded area) with simulated stability region (area to the left of the dotted line), for the *SYSTEM 1*. (Right) Provably safe regions of α and β (for $d = 200$ ms).

switched linear time-delay systems, thereby obtaining sound sufficient criteria for the stability of explicit congestion control protocols.

II. Theory: Discretized Lyapunov functionals for switched systems with time-delay

We propose a new method to search for piecewise quadratic Lyapunov functionals for switched linear systems. There has been a recent attempt to solve similar problems using Lyapunov functions of a different form [6]. Following results for switched systems with no time-delay and linear systems with time-delay [1], given a switched time-delay system $\dot{x}(t) = A_i x(t) + A_{d_i} x(t-d)$, state $x(t) \in X_i$, where X_i is a partition of the state space given by the dynamics, x_t is the trajectory, and i indexes the partition in the state space, we search for Lyapunov functionals of the form

$$\begin{aligned} V_i(x_t) \quad = \quad & x(t)^{\mathrm{T}} P_i x(t) + 2x(t)^{\mathrm{T}} \int_{-d}^{0} Q(\zeta) x(t+\zeta) d\zeta \\ & + \int_{-d}^{0} \int_{-d}^{0} x(t+\zeta)^{\mathrm{T}} R(\zeta,\eta) x(t+\eta) d\eta d\zeta \\ & + \int_{-d}^{0} x(t+\zeta)^{\mathrm{T}} S(\zeta) x(t+\zeta) d\zeta \end{aligned}$$

where $P_i = F_i^{\mathrm{T}} T F_i$ for continuity, as in the case of switched linear systems. We tackle this problem by combining time-discretization methods, so far used for linear time-delay systems [3], with a space discretization technique used to analyze switched systems with no time-delay. We try to find the matrix functions of space, P_i, and the matrix functions of (discretized) time, Q, R and S. Details of our approach are presented in [1]. There are several advantages in designing an analysis tool of this form: the time-discretization technique is known to decrease conservatism in proving stability for linear systems in which the stability depends on the values of the delay [3]; partitioning the state space is an efficient way of analyzing the stability of switched hybrid systems [4]. The combination of the two methods reduces the stability analysis to the solution of Linear Matrix Inequalities (done efficiently using standard solvers such as SeDuMi [8]). This method can be extended to systems with heterogeneous delays [1].

III. Results: Finding parameters with provable stability for the XCP equations

We use the methods described above to find Lyapunov functions that prove the stability of the XCP equations for different values of α, β, and d. We embed the XCP system equations (1) in a switched system which is defined for all $x(t) = y(t) - C$ and $q(t)$, given by

$$\begin{array}{ll} \dot{q}(t) = x(t) & \\ \dot{x}(t) = -\frac{\alpha}{d} x(t-d) - \frac{\beta}{d^2} q(t-d) & \text{if } q > 0 \text{ or } x \geq 0 \\ & \\ \dot{q}(t) = -q(t) & \\ \dot{x}(t) = -\frac{\alpha}{d} x(t-d) - \frac{\beta}{d^2} q(t-d) & \text{if } q \leq 0 \text{ and } x \leq 0 \end{array} \qquad (2)$$

We refer to this system as *SYSTEM 2*. *SYSTEM 1* is stable if *SYSTEM 2* is stable. This is because *every* trajectory of (1) is a trajectory of (2). Since a system is stable if and only if every trajectory of it is stable, *SYSTEM 2* is stable implies that *SYSTEM 1* is stable.

The outer boundaries of the provably stable regions of parameters for a round trip delay of 200 ms are plotted in Fig. 1 (right). The smaller (dark) region corresponds to the stable region predicted by linear analysis, which ignores the switch. The inset shows a closer look at the region where the switched Lyapunov results are conservative (which is to be expected, since they are derived from a sufficient condition for stability) – while the linear analysis results predict a stable system, the switched system is unstable. Fig. 1 (right) also shows that the actual stable region is much larger than that predicted by the linearization.

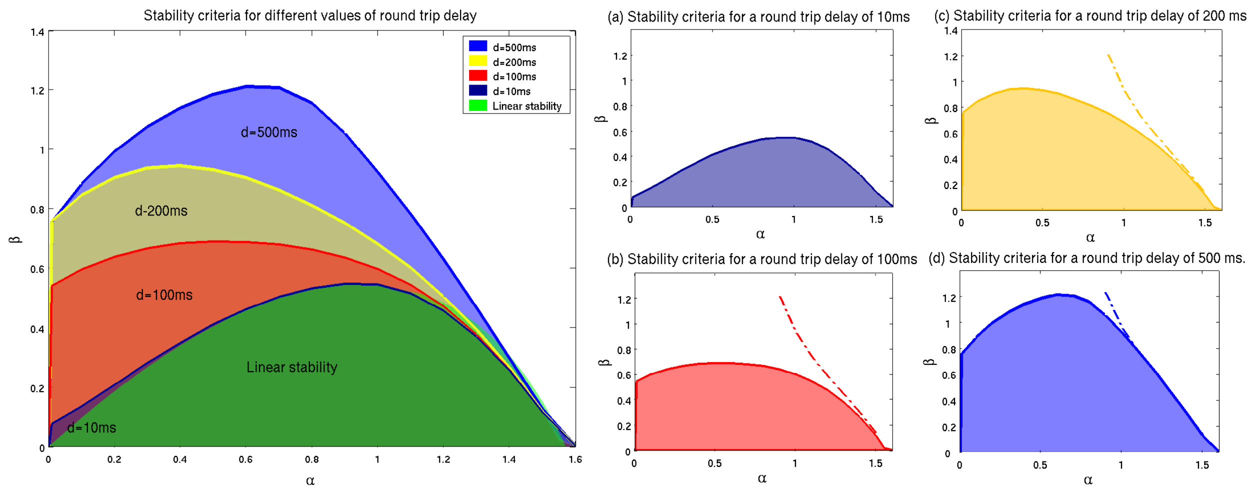

Fig. 2. Provably safe boundaries of α and β (for $d = 10$ms to 200 ms). The dotted lines in the figures on the right correspond to the simulated stability boundaries.

Effect of delay on stability criteria

The proposed Lyapunov functionals for switched systems provide us with sufficient conditions for delay-dependent stability. Since studies have shown that 85% of Internet traffic has round trip delays between 15-500 ms [7], we analyze the stability for this range of round trip delays. The provable stability boundaries, in terms of α and β are shown in Fig. 2. We find that for small delays, it is more difficult to prove the stability of the switched system. We should bear in mind that these results are based on sufficiency conditions, and therefore our not being able to prove stability does not imply instability. For values of delay more than 100 ms, we can prove stability for a substantially large range of parameters. Even for small values of delay, we note that the region stays larger than previously derived using linearization.

In particular, we can prove that the range of parameters recommended for XCP in [5], namely $0 < \alpha < \frac{\pi}{4\sqrt{2}}$ and $\beta = \alpha^2\sqrt{2}$, *is*, in fact, stable for values of delay ranging from 10 ms to 500 ms. While the linear analysis that was used to prove the stability of XCP in [5] was not valid for the XCP model (the equilibrium lies on a line of discontinuity in the dynamics), their results and choice of parameters have been validated in this work. These techniques can also be extended to the case of heterogeneous delays on bottleneck links [1].

IV. Conclusions

We have presented a new computational technique that handles discontinuities and time-delays, and provides less conservative estimates of stable regions of explicit congestion control protocols which do not satisfy linear approximation and analysis tools. This technique also applies to the analysis of traffic engineering protocols such as TeXCP, and networks with heterogeneous delays.

References

[1] H. Balakrishnan, N. Dukkipati, N. McKeown, and C. J. Tomlin. Stability analysis of explicit congestion control protocols. September 2005. Technical Report: SUDAAR 776, `http://yuba.stanford.edu/rcp/SUDAAR_776.pdf`.

[2] N. Dukkipati, M. Kobayashi, R. Zhang-Shen, and N. McKeown. Processor sharing flows in the internet. In *International Workshop on Quality of Service*, 2005.

[3] K. Gu, V. Kharitonov, and J. Chen. *Stability of Time-delay Systems*. Birkhauser, Boston, 2003.

[4] M. Johansson and A. Rantzer. Computation of piecewise quadratic Lyapunov functions for hybrid systems. *IEEE Transactions on Automatic Control*, 43(4):555–559, April 1998.

[5] D. Katabi, M. Handley, and C. Rohrs. Congestion control for high bandwidth-delay product networks. In *SIGCOMM*, 2002.

[6] V. Kulkarni, M. Jun, and J. Hespanha. Piecewise quadratic Lyapunov functions for time-delay hybrid systems. In *American Control Conference*, 2004.

[7] S.H. Low, F. Paganini, J. Wang S. Adlakha, and J.C. Doyle. Dynamics of TCP/RED and a scalable control. In *IEEE INFOCOM*, 2002.

[8] J.F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11–12:625–653, 1999.