Stability Analysis of Explicit Congestion Control Protocols

Hamsa Balakrishnan, Nandita Dukkipati, Nick McKeown and Claire J. Tomlin
Stanford University, Stanford, CA 94305
Email: {hamsa, nanditad, nickm, tomlin}@stanford.edu

Abstract

Much recent attention has been devoted to analyzing the stability of congestion control algorithms, in the context of TCP modifications (e.g., TCP/RED [10], [15], FAST [17]) and new protocols (e.g., XCP [21], RCP [8], TeXCP [20]). The control-theoretic framework used in most previous work is linear systems theory. The analyses assume that the system can be well approximated by linearization, and the linearization is then used to derive conditions for stability using techniques based on the Bode or Nyquist criteria.

We show that linearization is *not* a good approximation when the queue lengths are close to zero. Because the goal of several congestion control algorithms is to keep queue lengths small, the linearization turns out to be the most inaccurate precisely in the realm in which a good algorithm would hope to operate. We show, in the context of explicit congestion control protocols like XCP and RCP, that the stability region derived from traditional Nyquist analysis is not an accurate representation of the actual stability region. Using XCP as an example, we then show that modeling the congestion control algorithm as a switched linear control system with time delay, and using new Lyapunov stability conditions can provide sound and more general sufficient conditions for stability than previously derived. For piecewise linear systems with time-delay, the proposed conditions guarantee global stability. We show that the proposed framework can be used to analyze the stability of congestion control protocols in the presence of heterogeneous delays.

Stanford University Department of Aeronautics and Astronautics Report: SUDAAR No. 776, September 9, 2005. This research was supported by an NSF Career award (ECS-9985072). H. Balakrishnan was supported by a Stanford Graduate Fellowship.

I. Introduction

There has been much interest recently in fluid-flow models to analyze the stability of congestion control protocols ([15], [27], [26]). Several researchers have observed that the feedback-based congestion control algorithm TCP [16] is prone to instability as the bandwidth-delay product of the network grows [21], [26]. This observation, as well as the inability of TCP to perform adequately over such network paths, has motivated the development of new congestion control algorithms, such as the eXplicit Control Protocol (XCP) [21] and the Rate Control Protocol (RCP) [8], as well as modifications to TCP (such as STCP [23] , FAST [17] and HSTCP [9]). To demonstrate the stability of these protocols, 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. These stable parameters are then used to tune the protocol. The linear analysis provides local stability results as long as the equilibrium stays bounded away from the physical limits on the system state, such as queue lengths of zero or the maximum value.

The success of linearized analysis depends on how well the system dynamics can be approximated by its first-order behavior about the equilibrium point. If the system equations are "well-behaved" (continuous and differentiable), the stability of the original nonlinear system is generally represented well by the stability of the linearized system. However, when the equilibrium point lies on a discontinuity in the system dynamics, the stability of the linearized system gives *no* guarantees on the stability of the system. In most congestion control protocol models, these discontinuities in system behavior arise from the physical constraints of the system, for example, that queue lengths cannot be negative, resulting in a difference in the system behavior between the case in which the queue lengths are positive and the queue lengths are zero. This system can be thought of as a single system with several (two, in this case) possible modes of behavior, which switches from one mode to the other when the queue length is zero.

In this paper, we advocate caution in the use of linear stability theory in the analysis of congestion control protocols, using explicit congestion control protocols[1] like XCP and RCP, in which the region of stability derived using linear techniques is either too small, or more importantly, too large. Instead, we propose a method for taking discontinuities in the system dynamics into account by modeling the protocol as a switched system. This method is particularly important for the analysis of protocols in which the equilibrium queue lengths are small. We then present a computational technique to analyze the stability of switched linear systems, thereby obtaining less conservative and more sound sufficient criteria for the stability of congestion control protocols. The proposed technique reduces the problem to the solution of Linear Matrix Inequalities (LMIs) [4]. This is a fast computational technique since the stability analysis is now reduced to a convex optimization problem. We show that the method can be extended to heterogeneous round trip delays.

The organization of the remainder of this paper is as follows: In Section II we discuss related work, in particular XCP and RCP. In Section III-A we briefly discuss the stability requirements for a switched system with no time delay, in Section III-B, we present known results on the stability of linear (continuous) systems with time-delay, and in Section III-C, we briefly present our results on the stability of switched time-delay systems. In Section IV, we apply these techniques to the analysis of the stability of XCP, for both uniform and heterogeneous delays.

II. Background and Related Work

A. Motivation

As an example of explicit congestion control, we consider the eXplicit Control Protocol (XCP) [21]. Two key ideas in this protocol are that it generalizes Explicit Congestion Notification [29], so that the routers inform the senders about the degree of congestion at the

[1]Explicit congestion control protocols receive explicit rate/window feedback from the network on how to react to congestion, while in TCP the feedback is implicit, like experiencing packet loss.

bottleneck, and that the utilization control is decoupled from the fairness control. XCP is a window-based congestion control protocol whose dynamics are governed by the following properties: the aggressiveness of the sources is adjusted depending on the delay in the feedback loop, the system slows down as the feedback delay increases, and the system is designed so as to make the dynamics of the aggregate traffic independent of the number of flows [21]. In addition, since the utilization control is decoupled from the fairness control, the efficiency involves only the aggregate traffic behavior. The system dynamics are described in more detail in Appendix A. 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 [21]:

$$\dot{y}(t) = -\frac{\alpha}{d}(y(t-d) - C) - \frac{\beta}{d^2}q(t-d) \tag{1}$$

$$\dot{q}(t) = \begin{cases} y(t) - C & , \quad q(t) > 0 \\ [y(t) - C]^+ & , \quad q(t) = 0 \end{cases} \tag{2}$$

where q is the queue size and y is the aggregate traffic rate. The notation $[y(t) - C]^+$ denotes $\max(0, y(t) - C)$.

A classical linear analysis, such as that in [21], would analyze the linearization by considering only one possible mode of behavior of the system, namely:

$$\dot{y}(t) = -\frac{\alpha}{d}(y(t-d) - C) - \frac{\beta}{d^2}q(t-d) \tag{3}$$

$$\dot{q}(t) = y(t) - C \tag{4}$$

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) in Fig. 1. The Nyquist analysis would suggest 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. Now, consider two sets of parameter values, $\alpha = 0.8$, $\beta = 0.55$ and $\alpha = 1.4$, $\beta = 0.3$.

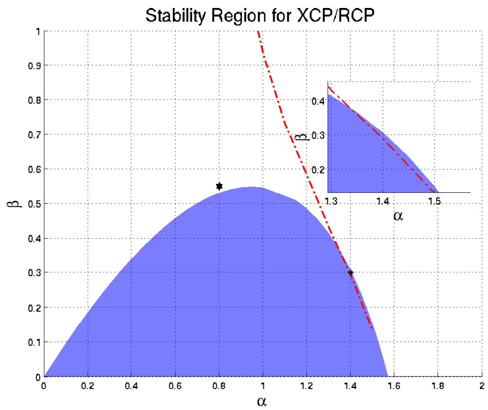

Fig. 1. Comparison of linearized stability region (shaded area) with simulated stability region (area to the left of the dotted line), for the system in (1)-(2).

If we simulate both the linearized and the switched systems for these sets of values, we notice that the first set of parameters results in a stable system (Fig. 2(a)).

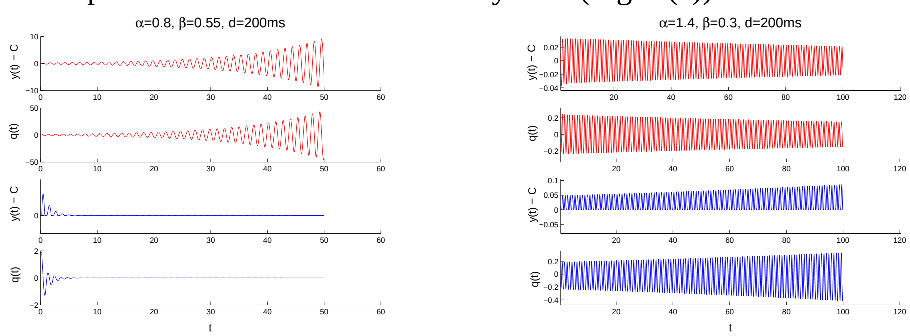

(a) (b)

Fig. 2. Simulated XCP system (rate and queue length vs. time), for linearization [top two] and switched system [bottom two] for (a) $\alpha = 0.8$, $\beta = 0.55$, and for a delay of 200 ms. and (b) $\alpha = 1.4$, $\beta = 0.3$, and for a delay of 200 ms.

For the second set of parameters, linear analysis would predict a stable system, while simulations indicate that the system is unstable (Fig. 2(b)).

The considerable difference in the stable region predicted by linear analysis and the actual stable region suggests that more careful analysis is required. In particular, the equilibrium of the XCP system has zero queue length. If we treat the system as a switched system with two modes of operation, one when the queue length is positive, and one when the queue length is zero, the equilibrium point lies on the line ($q(t) = 0$) on which the switching between the two systems occurs. In fact, it can be shown that for a switched system, linearizing the system about an equilibrium point at which the system dynamics are discontinuous could lead one to erroneous conclusions, even on the local stability of the system [5].

B. The Rate Control Protocol (RCP) [8]

Rate Control Protocol (RCP) is a congestion control algorithm whose key goal is to finish flows as quickly as possible. This is done by explicitly emulating ideal processor sharing (PS) at each link. RCP is an adaptive algorithm that updates the rate assigned to the flows, to approximate processor sharing in the presence of feedback delay without any knowledge of the number of ongoing flows. RCP has three main characteristics that make it simple and practical: (1) the router assigns a *single* rate, $R(t)$, to all flows passing through a link, (2) $R(t)$ is picked based on minimal information such as the aggregate traffic rate and queue occupancy, and (3) there is no per-flow state, queue or per-packet computation.

The router updates $R(t)$ as: $R(t) = R(t-T)[1 + \frac{\frac{T}{d_0}(\alpha(C-y(t))-\beta\frac{q(t)}{d_0})}{C}]$, where d_0 is a moving average of the RTT measured across all flows and T is the rate update interval (*i.e.*, how often $R(t)$ is updated, where $T \leq d_0$). α and β are parameters chosen for stability and performance. If there is spare capacity available (i.e., $C > y(t)$), it is shared equally among all flows. On the other hand, if there is a queue building up ($q(t) > 0$), then the link is oversubscribed and the flow rate is decreased evenly. $\alpha(C-y(t))-\beta\frac{q(t)}{d_0}$ is the desired aggregate change in traffic in the next control interval, and dividing this expression by the number of ongoing flows, gives the change in traffic rate needed per flow. The routers estimate the number of flows as $\hat{N}(t) = \frac{C}{R(t-T)}$.

The RCP control equation bears similarity to the XCP control equation because both protocols try to emulate processor sharing, but the manner in which the individual flows converge to PS differs considerably between XCP and RCP. The dynamics of the aggregate traffic flow in RCP is also modeled by (1)-(2).

C. Related Work

There has been much interest lately in the analysis of the stability of congestion control protocols [6], [27], [15], [19], as well as the design of scalable controls for networks [22], [3], [26], [14], [17]. There has also been progress in the development of new explicit congestion control algorithms [21], [8]. Most of these efforts use tools from linear systems analysis to obtain results on the stability of the protocol [33], [26]. More recently, there have been Lyapunov theory-based approaches to stability analysis, especially to prove the global stability of protocols governed by nonlinear (but continuous) differential equations with delays [7], [31], [34], [28]. The main difference between our paper and the previous work in this area is that we propose a method for the stability analysis of protocols described by delay differential equations with discontinuities. This methodology allows us to analyze systems in which the equilibria have small queue lengths, and the linearization either fails or is a poor approximation of the actual protocol. This analysis is particularly applicable to explicit congestion control protocols like XCP [21] and RCP [8], as well as a technique recently proposed for online traffic engineering called TeXCP [20].

III. Methods and Theory

We prove stability for ranges of parameters of congestion control protocols using Lyapunov theory [30], [13]. These techniques are time-domain based methods, as opposed to Nyquist criteria, which are frequency-domain based methods. One of the advantages of time-domain based methods over frequency-domain based methods is that they extend easily to nonlinear systems, as well as to systems with time-varying delays [12]. In this section, we first present the stability criteria for switched systems without time delays, and note that they can be expressed as Linear Matrix Inequalities (LMIs). Solving LMIs [4] is a convex optimization problem, and there are several efficient toolboxes such as SeDuMi [32] for computing their solutions. We then consider stability criteria for linear time-delay systems (with no switches) and note that techniques exist to prove the stability of these systems using Lyapunov-Krasovskii functionals [12], also reducing the problem to the solution of LMIs. Finally, we consider switched time-delay systems, and combining the two problems, find equivalent LMIs to prove stability.

A. Switched Systems

The challenge of proving the stability of switched systems arises partly from the observation that even if the individual systems are stable, switching between them could result in an unstable system. We illustrate this with an example from [5]. We consider two linear systems given by the equations

$$\begin{aligned} \dot{z}(t) &= -0.1z + y \\ \dot{y}(t) &= -10z - 0.1y \end{aligned} \quad \text{(System A)} \quad \& \quad \begin{aligned} \dot{z}(t) &= -0.1z + 10y \\ \dot{y}(t) &= -z - 0.1y. \end{aligned} \quad \text{(System B)} \tag{5}$$

The two systems are individually globally exponentially stable. We now consider a system that switches between the two, such that it follows the dynamics of system (5-A) if $z \leq 0,\ y \geq 0$ or $z \geq 0,\ y \leq 0$, and the dynamics of system (5-B) if $z \geq 0,\ y \geq 0$ or $z \leq 0,\ y \leq 0$.

Simulations for a range of initial conditions show that the switched system is unstable, although the individual systems are stable. Some trajectories of the switched system are shown in Fig. 3(a).

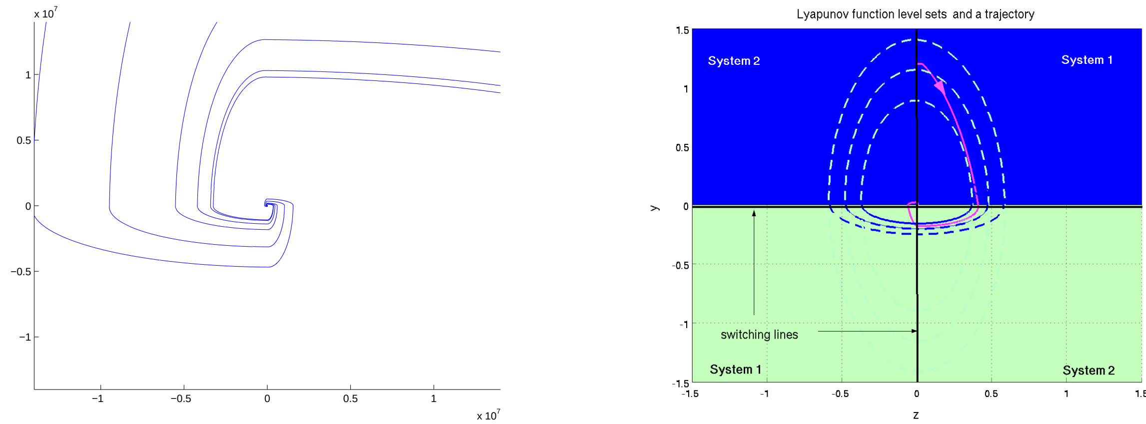

(a) (b)

Fig. 3. (a) Phase portrait for switched system ($y(t)$ vs. $z(t)$), with (5, A - B) in an unstable configuration. (b) Phase portrait for the stable combination of the switched system (solid line), with (5), along with the Lyapunov function level sets (dashed line) that prove stability

Similarly, it is possible to switch between two unstable systems or a stable and unstable system, and obtain a stable switched system. It is not sufficient to ignore the switch, and to consider only the linearization of each system around the equilibrium, while trying to analyze the stability of a switched system – in fact, the stability of the linearized system is *neither necessary nor sufficient* for the stability of the switched system.

Lyapunov stability theory is well-suited for the analysis of such switched systems. The challenge in applying this technique lies in the search for the right form of the Lyapunov function

$V(x,t)$, which in general can be quite complex. For a linear system, a quadratic Lyapunov function is *both necessary and sufficient* for global asymptotic stability. This is a function of the form $V(x,t) = x(t)^{\mathrm{T}} P x(t)$, where P is a positive definite matrix. Since many systems of interest are essentially nonlinear because of "switching" between several linear modes, they can be well-represented by piecewise linear systems. The logical result of this observation is to search for piecewise quadratic Lyapunov functions [18]. This technique involves dividing the state space (the space in which $x(t)$ takes values) into cells, and searching for a quadratic Lyapunov function for each cell. (For the XCP system (1-2) this space is $\mathbb{R}^2$, since the state variables are the rate and the queue length). We then need to enforce the condition that when a trajectory crosses from one cell to another, the Lyapunov function does not increase. A computational method for constructing piecewise quadratic Lyapunov functions to prove stability for switched linear systems was presented in [18]. The problem was formulated as a convex optimization problem based on linear matrix inequalities (LMIs). The results in [18] for piecewise linear systems are described in the following theorem:

Theorem 1: [18] Consider a piecewise linear system of the form

$$\dot{x}(t) = A_i x(t), \text{ for } x(t) \in X_i \tag{6}$$

where $\{X_i\}_{i \in I} \subseteq \mathbb{R}^n,\ I \in \{1, 2, \cdots\}$ is a partition of the state space into a number of closed (possibly unbounded) polyhedral cells with pairwise disjoint interior (*i.e.*, the intersections are restricted to the boundaries between the cells). The index set of cells is denoted by I. We assume that all the cells contain the origin. Then, it is possible to construct matrices E_i and F_i such that

$$\begin{aligned} &E_i x \geq 0\,, && x \in X_i,\ i \in I \\ &F_i x = F_j x\,, && x \in X_i \cap X_j,\ i, j \in I. \end{aligned} \tag{7}$$

If there exist symmetric matrices T, U_i and W_i such that U_i and W_i have nonnegative entries, while $P_i = F_i^{\mathrm{T}} T F_i,\ i \in I$, satisfies

$$\begin{array}{rcl} 0 & > & A_i^{\mathrm{T}} P_i + P_i A_i + E_i^{\mathrm{T}} U_i E_i \\ 0 & < & P_i - E_i^{\mathrm{T}} W_i E_i \end{array},\ i \in I, \tag{8}$$

then every continuous $\mathcal{C}^1$ trajectory $x(t) \in \cup_{i \in I} X_i$ satisfying (6) for $t \geq 0$ tends to the equilibrium exponentially. ■

The conditions given above are linear matrix inequalities in T, U_i and W_i. The parameterization (8) ensures that the level sets of the Lyapunov function from the different regions match at the boundaries. They can be further relaxed so that the Lyapunov function is decreasing on the boundaries [18]. We consider the same system equations (5, A–B), but with the switching scheme reversed from the previous example. We find that this switched system is stable, and in fact, we can find a piecewise quadratic Lyapunov function that proves stability using Theorem 1. The level sets of the Lyapunov function are also plotted in Fig. 3(b).

B. Lyapunov functionals for time delay systems

The difficulty in dealing with time-delay systems is that the future state of the system depends, not only on its current state, but on its past trajectory too. We need to develop means of not only dealing with finite-dimensional systems, but infinite-dimensional systems. These systems are analyzed using delay differential equations [13]. In general, such equations, also known as functional differential equations, are given by $\dot{x}(t) = f(x_t, t)$, where f is a *functional*, and x_t is the graph of x on $[t_0 - d, t]$, where t_0 is the initial time. Unlike ordinary differential equations, where the initial conditions are prescribed at one instant in time, we note that for delay differential equations, the initial conditions are a function $\phi(t)$, defined in the interval $[-d, 0]$.

We first define the notion of stability for a time-delay system.

Definition 1: The *continuous norm* of the initial condition is defined as $\| \phi \|_c = \max_{-d \leq \theta \leq 0} \| \phi(\theta) \|$. The trivial solution $x(t) = 0$ is said to be *stable* if for any $t_0 \in \mathbb{R}$ and any $\epsilon > 0$, there exists a $\delta(t_0, \epsilon)$ such that $\| \phi \|_c < \delta$ implies that $\| x(t) \| < \epsilon$ for $t > t_0$. $x(t) = 0$ is said to be *asymptotically stable* if it is stable, and for any $t \in \mathbb{R}$ and $\epsilon > 0$ there exists a $\delta = \delta(t_0, \epsilon) > 0$ such that $\parallel \phi \parallel_c < \delta$ implies $\lim_{t \to \infty} x(t) = 0$.

We consider the linear time-delay system

$\dot{x}(t) = Ax(t) + A_d x(t-d);\ \ x(t) = \phi(t),\ t \in [-d, 0]$.

Let $\mathcal{C}$ be the set of $\mathbb{R}^n-$valued continuous functions in $[-d, 0]$. Then $x_t \in \mathcal{C}$ is a segment of a system trajectory, $x_t(\theta) = x(t+\theta),\ -d \leq \theta \leq 0$.

There are several candidates for Lyapunov functionals ([13], [12]). We consider the counterpart of the quadratic Lyapunov function, namely the quadratic Lyapunov-Krasovskii functional. A **quadratic Lyapunov-Krasovskii functional** is given by a function $V : \mathcal{C} \longrightarrow \mathbb{R}$, such that

$$\begin{aligned} V(x_t) \;\; = \;\; & \frac{1}{2} x(t)^{\mathrm{T}} P x(t) + x(t)^{\mathrm{T}} \int_{-d}^{0} Q(\zeta) x(t+\zeta) d\zeta + \frac{1}{2} \int_{-d}^{0} \int_{-d}^{0} x(t+\zeta)^{\mathrm{T}} R(\zeta, \eta) x(t+\eta) d\eta d\zeta \\ & + \frac{1}{2} \int_{-d}^{0} x(t+\zeta)^{\mathrm{T}} S(\zeta) x(t+\zeta) d\zeta \end{aligned} \tag{9}$$

where $P = P^{\mathrm{T}} \in \mathbb{R}^{n \times n}$, $Q(\zeta) \in \mathbb{R}^{n \times n}$, $R(\zeta, \eta) = R^{\mathrm{T}}(\eta, \zeta) \in \mathbb{R}^{n \times n}$, and $S(\zeta) = S^{\mathrm{T}}(\zeta) \in \mathbb{R}^{n \times n}$.

The following theorem is a Lyapunov stability theorem for time-delay systems.

Theorem 2: ([13], [11]) The system

$$\dot{x}(t) = Ax(t) + A_d x(t-d);\ \ x(t) = \phi(t),\ t \in [-d, 0]. \tag{10}$$

is asymptotically stable if there exists a Lyapunov functional V of form (9) such that for some $\epsilon > 0$, it satisfies

$$V(x_t) \geq \epsilon x(t)^{\mathrm{T}} x(t) \tag{11}$$

and its derivative along the solution of (10) satisfies

$$\dot{V}(x_t) \leq -\epsilon x(t)^{\mathrm{T}} x(t) \tag{12}$$

for any $x_t \in \mathcal{C}$, where $\dot{V}(\phi) \triangleq \frac{d}{dt} V(x_t) \mid_{x_t = \phi}$. ■

While such a Lyapunov functional might exist, it could be difficult to compute. The discretized functional method [11] tries to compute $Q(\zeta)$, $R(\zeta,\eta)$ and $S(\zeta)$ as piecewise constant functions, on $[-d,0]$. These conditions can also be formulated as LMIs. The method is described in more detail in [11].

C. Discretized Lyapunov functionals for switched hybrid 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 [24]. Following results for switched systems with no time-delay (Section III-A) and linear systems with time-delay (Section III-B), we search for Lyapunov functionals of the form

$$\begin{aligned} V_i(x_t) \;&=\; \frac{1}{2}x(t)^{\mathrm{T}}P_i x(t) + x(t)^{\mathrm{T}}\int_{-d}^{0} Q(\zeta)x(t+\zeta)d\zeta \\ &\quad + \frac{1}{2}\int_{-d}^{0}\int_{-d}^{0} x(t+\zeta)^{\mathrm{T}}R(\zeta,\eta)x(t+\eta)d\eta d\zeta + \frac{1}{2}\int_{-d}^{0} x(t+\zeta)^{\mathrm{T}}S(\zeta)x(t+\zeta)d\zeta \end{aligned} \tag{13}$$

where $P_i = F_i^{\mathrm{T}}TF_i$ for continuity, as in the case of switched linear systems. Our approach to tackle this problem is to combine the time-discretization methods, so far used for linear time-delay systems [11], with the space discretization technique, so far used to analyze switched systems with no time-delay. Details of our approach are presented in Appendix B. In short, given a switched time-delay system

$$\dot{x}(t) = A_i x(t) + A_{d_i} x(t),\ x(t) \in X_i \tag{14}$$

where X_i is a partition of the state space given by the dynamics, we try to find the matrix functions of space, P_i, and the matrix functions of time, Q_p, R_{pq} and S_p (as described in Appendix B), all on their respective discretized grids. 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 [12]; the space-discretization is an efficient way of analyzing the stability of switched hybrid systems [18]. The combination of the two methods reduces the problem to the solution of Linear Matrix

Inequalities,which is a convex optimization problem. Several very efficient toolboxes have been developed for solving LMIs [32], making the method proposed in this paper computationally fast and easy to implement. In addition, similar LMI formulations are possible for the stability analysis of systems with heterogeneous and time-varying delays.

IV. Results

A. 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 (1-2) in a switched system which is defined for all $x(t) = y(t) - C$ and $q(t)$, given by

$$\begin{aligned} &\dot{q}(t) = x(t) \\ &\dot{x}(t) = -\tfrac{\alpha}{d}x(t-d) - \tfrac{\beta}{d^2}q(t-d), \end{aligned} \quad q > 0 \text{ or } x \geq 0 \tag{15}$$

$$\begin{aligned} &\dot{q}(t) = -q(t) \\ &\dot{x}(t) = -\tfrac{\alpha}{d}x(t-d) - \tfrac{\beta}{d^2}q(t-d), \end{aligned} \quad q \leq 0 \text{ and } x \leq 0 \tag{16}$$

The system (1-2) is stable if the switched system(15-16) is stable. This is because *every* trajectory of system (1-2) is a trajectory of (15-16). Since the stability of the latter implies that every trajectory is stable, stability of (15-16) implies the stability of (1-2).

A sample trajectory of the stable subsystem (15) as well as the stable switched system (15-16) are shown in Fig. 4. These plots correspond to the parameter values $d = 200$ ms, $\alpha = 0.8$ and $\beta = 0.5$.

However, it is possible that one of the individual subsystems is unstable, and yet the switched system is stable, implying that XCP is stable. One such example is shown in Fig. 5, corresponding to $d = 200$ ms, $\alpha = 0.8$ and $\beta = 0.55$.

The outer boundaries of the provably stable regions of parameters for a round trip delay of 200 ms are plotted in Fig. 6. The smaller (dark) region corresponds to the stable region predicted by linear analysis, which ignores the switch.

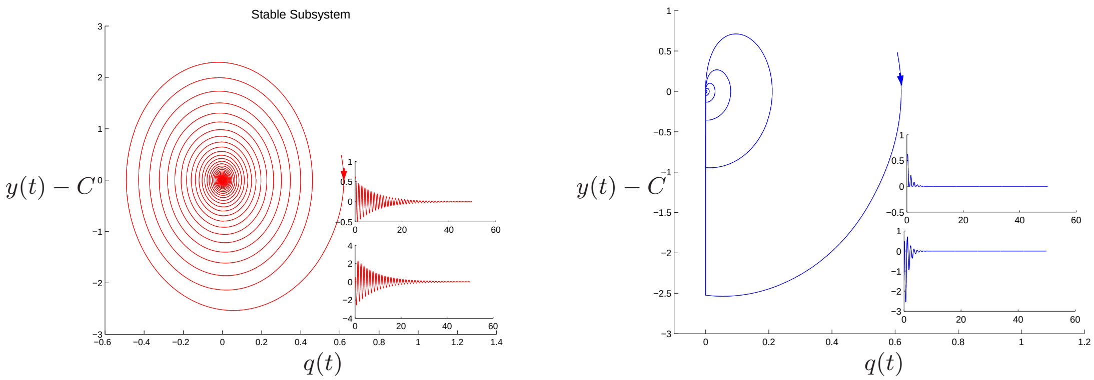

Fig. 4. Phase portrait ($y(t) - C$ vs. $q(t)$) of the stable subsystem [left], with stable switched system [right]

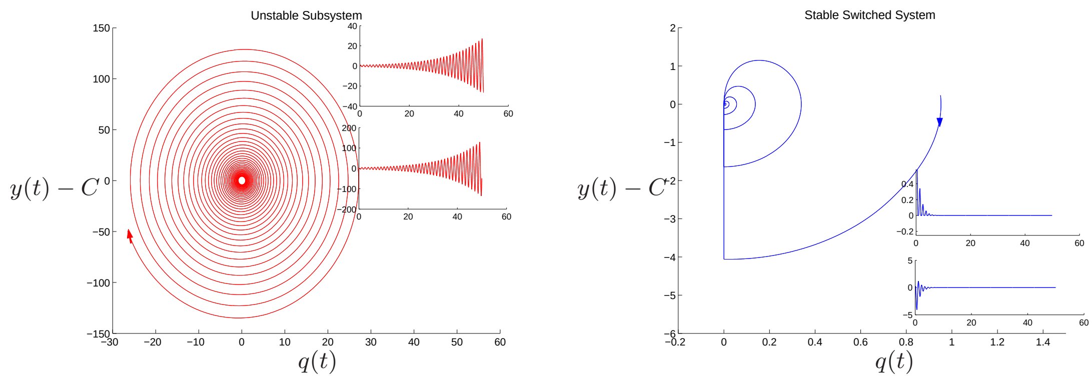

Fig. 5. Phase portrait ($y(t) - C$ vs. $q(t)$) of unstable subsystem [left], with stable switched system [right]

There are two things to note here. 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. 6 also shows that the actual stable region is much larger than that predicted by the linearization.

B. Effect of delay on stability criteria

The proposed Lyapunov-Krasovskii 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 times between 15-500 ms ([1], [26]), we analyze the stability for this range of round trip times. The provable stability boundaries, in terms of α and β are shown in Fig. 7. We find that for small delays, it is more difficult to prove the stability of the switched system.

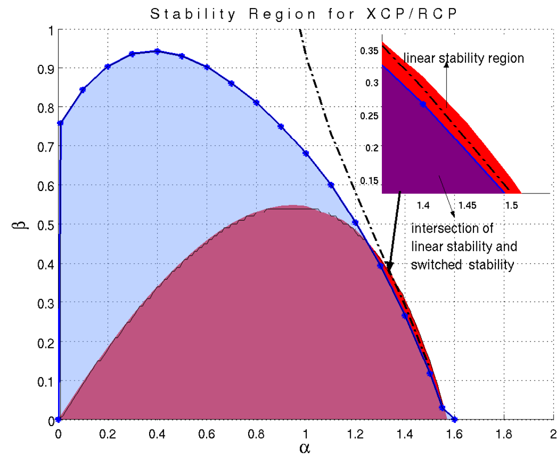

Fig. 6. Provably safe regions of α and β (for $d = 200$ ms)

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. However 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 prove that the range of parameters recommended for XCP in [21], namely

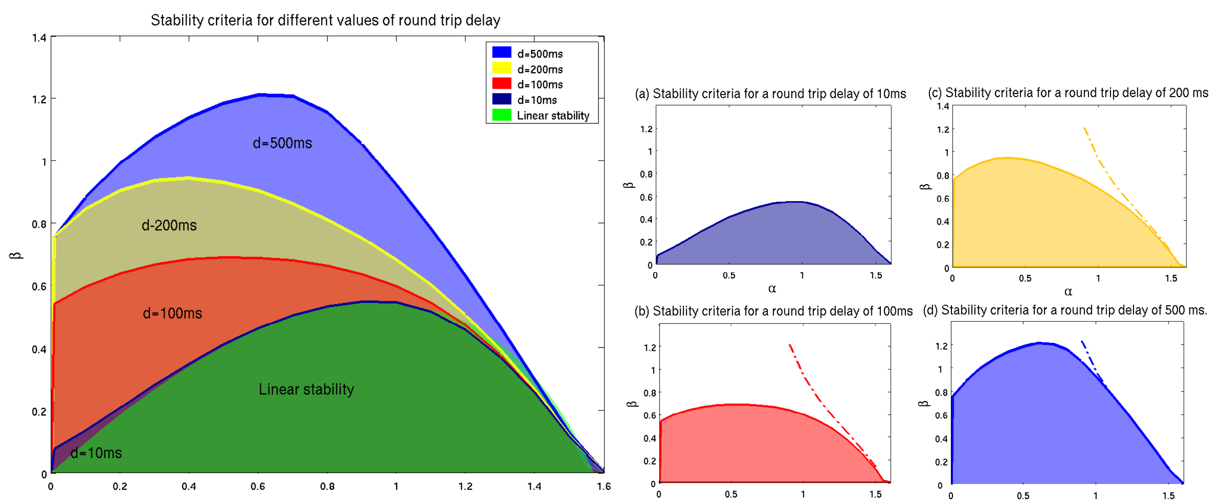

Fig. 7. Provably safe boundaries of α and β (for $d = 10$ms to 200 ms). The dotted lines in the figures on the right correspond to the simulated stability boundaries.

$0 < \alpha < \frac{\pi}{4\sqrt{2}}$ and $\beta = \alpha^2\sqrt{2}$, *is*, in fact, stable for values of delay ranging from 10 ms to 500 ms. In [25], the authors suggest, with reference to XCP [21], that the practice of linearization around the equilibrium requires caution at the tightest bottlenecks that have zero queue yet full utilization, since the equilibrium queues are zero. While the linear analysis that was used to prove the stability of XCP in [21] was not valid for their system (the equilibrium queue length is zero, which is on a point of discontinuity in the dynamics), the results and choice of parameters in [21] are validated in the present work.

C. Robustness

In [21], experimental data was presented to demonstrate the robustness of XCP to high variance in the round trip time. The robustness was measured in terms of the final throughput of two XCP flows with differing RTTs (20 ms and 200 ms), sharing a bottleneck. The theoretical analysis of this behavior, which to our knowledge has not been carried before, requires the stability analysis of a switched linear system with heterogeneous time delays. The system equations, for a single bottleneck link in which the sources have different round trip times, are described in Appendix C. The region of stable parameters is shown in Fig. 8 for RTTs of 100 ms, 200 ms and 400 ms. The results presented there further demonstrate the robustness of XCP to high variance in RTT, with respect to stability. Another important notion is robustness to the variation in system

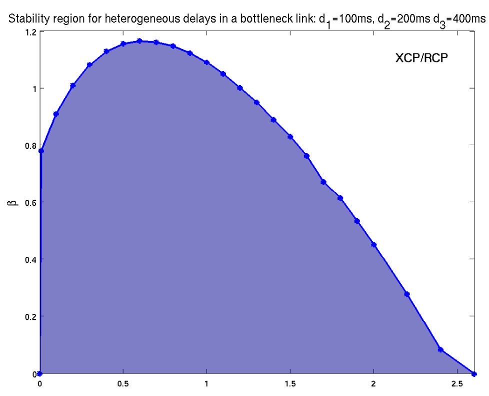

Fig. 8. Values of α and β for which the XCP system is stable, for a single bottleneck link with heterogeneous delays of 100 ms, 200 ms and 400 ms.

parameters, or model uncertainty. In particular, let us consider the case when we can find a Lyapunov functional (consisting of the P_i's, Q, R, S, W_i's and U_i's, as described in Appendix B, Theorem 5) that proves stability for some set of models $\{A_i^{(j)}, A_{d_i}^{(j)}\}$, $j = 1, \cdots, M$, for the system in (14). Then, because of the convex nature of the LMIs, the system is stable for any (possibly time-varying) model in the convex hull of this set of models, *i.e.*, for any system model $\{A_i(t), A_{d_i}(t)\} \in \mathbf{co}\{(A_i^{(j)}, A_{d_i}^{(j)}),\ j = 1, \cdots, M\}$, providing us with a stability result that is robust to model uncertainty.

V. TeXCP: An XCP-Like Protocol for Traffic Engineering

Another context in which the concepts studied in this paper prove useful is in the analysis of online traffic engineering protocols such as TeXCP [20]. TeXCP is a online distributed traffic engineering protocol for balancing loads in realtime. TeXCP recognizes the connection between load balancing and congestion control, and applies ideas from the design of XCP to traffic engineering. As part of the design process, it was necessary to model and analyze the stability of TeXCP [20]. The details of the design process can be found in [20].

As in the case of XCP, TeXCP uses explicit feedback from the routers. The form of the feedback controller is similar to that of XCP. The core router makes the increase proportional to the spare bandwidth (S = Capacity $-$ Load) and the decrease proportional to the queue size (Q). The feedback (Φ) is given by $\Phi = \alpha \cdot T_p \cdot S - \beta \cdot Q$, where the `probe timer` maintained by TeXCP agent fires every T_p seconds. T_p is larger than the maximum round trip time in the network, and is generally set to 100 ms. α and β are parameters chosen to ensure feedback stability. For a single bottleneck along each path, and for a constant RTT of d, if the ingress-egress flows have infinite demands, the system can be modeled by the equations [20]

$$\dot{\phi}(t) = -\frac{\alpha}{T_p}(\phi(t-d) - C) - \frac{\beta}{T_p^2} q(t-d) \tag{17}$$

$$\dot{q}(t) = \begin{cases} \phi(t) - C & , \quad q(t) > 0 \\ [\phi(t) - C]^+ & , \quad q(t) = 0 \end{cases} \tag{18}$$

where q is the queue size and ϕ is the aggregate traffic rate. The notation $[\phi(t) - C]^+$ denotes $\max(0, \phi(t) - C)$. In addition, $T_p > d$.

We notice the similarity to the XCP model (1-2). The techniques we have applied to the stability analysis of XCP still hold for the stability of TeXCP. Since in the TeXCP model the time-dependent gain of the feedback is not the same as the network round trip time, we can comment on the robustness of stability to uncertainty in the RTTs.

Theorem 3: Consider the system given by (17-18). If, for a constant value of T_p, the system is stable for two (constant) values of the RTT, d_{min} and d_{max}, then the TeXCP system is stable for all constant values of d such that $d_{min} \leq d \leq d_{max}$.

Proof: The result follows from the convexity of the LMI conditions for switched time-delay systems (Appendix B) in the delay d. ■

In particular, for the default value of $T_p = 100$ ms, we find that the switched linear system with no time delay ($d = 0$) is stable for all values of α, β such that the system is stable for $d = T_p$. Therefore, we prove that the TeXCP model is stable for values of α, β in the shaded region in Fig. 7(b), for *all* constant RTTs such that $d < T_p$.

VI. Conclusions

We have demonstrated the ability of switched linear system analysis to provide more general results on the stability of explicit congestion control protocols, for cases which do not satisfy the linear approximation and analysis tools. We have shown through several examples that linear systems tools alone are in most cases too conservative, and yet there are regions in the parameter space in which they predict stability when the actual system is unstable. We have proposed a new computational technique that handles discontinuities (like saturation) as well as time-delays in the dynamics. We have shown that this technique also applies to the analysis of traffic engineering protocols such as TeXCP. We are currently working on the extension of this technique to time-varying round trip delays. We have shown through examples from the analysis of explicit congestion control like XCP and RCP that the proposed technique leads to sufficient conditions with less conservative estimates of stable regions, and can be easily applied to the analysis of protocols with discontinuous dynamics, in networks with both uniform and heterogeneous delays.

APPENDIX

A. eXplicit Control Protocol (XCP) Equations

For the sake of completeness, we present the XCP equations for a single link with identical flows [21]. We consider a single link of capacity c in a bottleneck topology, traversed by N XCP flows. $r_i(t)$ is the sending rate of user i at time t, resulting in an aggregate flow $y(t) = \sum r_i(t)$. The router sends the aggregate feedback every control interval d, which reaches the sources after a round trip delay. On receiving this feedback, the senders change the sum of their congestion windows, $\sum w_i(t)$. The aggregate feedback sent per unit time is
$\sum \dot{w}_i(t) = \frac{1}{d}\left(-\alpha d(y(t-d) - c) - \beta q(t-d)\right)$.

The aggregate behavior of the system can be described by

1) Feedback: $\dot{y}(t) = -\frac{\alpha}{d}(y(t-d) - c) - \frac{\beta}{d^2}q(t-d)$.
2) The evolution of the queue size is described by

$$\dot{q}(t) = \begin{cases} y(t) - c, & q(t) > 0 \\ [y(t) - c]^+, & q(t) = 0 \end{cases}$$

3) Fairness controller: The AIMD policy is given by $\dot{r}_i(t) = \frac{1}{N}\left([\dot{y}(t)]^+ + h(t-d)\right) - \frac{r_i(t-d)}{y(t-d)}\left([-\dot{y}(t)]^+ + h(t-d)\right)$ where $[\dot{y}(t)]^+$ denotes $\max(0, \dot{y}(t))$.

As in [21], this system can be analyzed by considering a delayed feedback system with the open loop transfer function $G(s) = \frac{\alpha ds + \beta}{d^2 s^2}$.

For simplicity, [21] selects values of α and β such that the zero-break frequency (ω_z) and the cross-over frequencies of the system are the same. They then prove stability of the linearized system

$$\begin{aligned} \dot{y}(t) &= -\frac{\alpha}{d}(y(t-d) - C) - \frac{\beta}{d^2}q(t-d) \\ \dot{q}(t) &= y(t) - C \end{aligned} \tag{19}$$

independent of the delay d, capacity C and number of sources N, for $\beta = \alpha^2\sqrt{2}$, $\alpha < \frac{\pi}{4\sqrt{2}}$. It is possible to similarly analyze (19) using linear (Bode) analysis [8], and prove that the linearized system is stable independent of d, C and N, in the shaded region in Fig. 1.

The Rate Control Protocol (RCP), while being very different from XCP in its behavior for individual flows, has similar average aggregate behavior, and can be modeled by the same equations [8]. In both cases, since the equilibrium queue lengths are zero, analysis of the linearized system does not give any guarantees on the stability of the actual system.

B. Discretized Lyapunov-Krasovskii functionals

For the sake of completeness, we first describe the methodology of computing a Lyapunov functional for a linear time delay system [11], [12]

$$\dot{x} = Ax(t) + A_d x(t-d) \tag{20}$$

The interval $\mathcal{I} = [-d, 0]$ is divided into N segments.

We define $\mathcal{I}_p = [\theta_p, \theta_{p-1}]$ for $p = 1 \cdots N$, $h = d/N$, and $\theta_p = -ph$, for $p = 0 \cdots N$.

The continuous matrix functions $Q(\zeta)$ and $S(\zeta)$ are chosen to be linear within each segment $\mathcal{I}(p)$, as

$$\begin{aligned} &Q_p = Q(\theta_p);\ S_p = S(\theta_p);\ Q(\theta_p + \alpha h) = (1-\alpha)Q_p + \alpha Q_{p-1}; \\ &S(\theta_p + \alpha h) = (1-\alpha)S_p + \alpha S_{p-1} \end{aligned}$$

for $0 \le \alpha \le 1$, $p = 1, 2, \cdots, N$. As in [11], the square $\mathcal{S} = [-d, 0] \times [-d, 0]$ is divided into $N \times N$ squares $\mathcal{S}_{pq} = [\theta_p, \theta_{p-1}] \times [\theta_q, \theta_{q-1}]$. Then, for $0 \le \alpha \le 1$, $0 \le \beta \le 1$, $p = 1, 2, \cdots, N$, $q = 1, 2, \cdots, N$, the continuous matrix function R is given by

$$\begin{aligned} &R(\theta_p + \alpha h, \theta_q + \beta h) \\ &= \begin{cases} (1-\alpha)R_{pq} + \beta R_{p-1,q-1} + (\alpha - \beta)R_{p-1,q}, & \alpha \ge \beta, \\ (1-\beta)R_{pq} + \alpha R_{p-1,q-1} + (\beta - \alpha)R_{p,q-1}, & \alpha < \beta. \end{cases} \end{aligned}$$

The quadratic Lyapunov-Krasovskii functional given in (9) is therefore determined by the matrices P, Q_p, S_p, R_{pq}, $p, q = 0, 1, \cdots, N$. The discretization is shown in Fig. 9.

Theorem 4: [11] If there exist $P = P^{\mathrm{T}}$, Q_p, $S_p = S_p^{\mathrm{T}}$, $R_{pq} = R_{pq}^{\mathrm{T}}$, $p, q = 0, \cdots N$ such that $\begin{pmatrix} P & \tilde{Q} \\ \tilde{Q}^{\mathrm{T}} & \tilde{R} + \tilde{S} \end{pmatrix} > 0;\ \begin{pmatrix} \Delta & -D^s & -D^a \\ -D^{s^{\mathrm{T}}} & R_d + S_d & 0 \\ -D^{a^{\mathrm{T}}} & 0 & 3S_d \end{pmatrix} > 0$ where Δ, D^a, D^s, $\tilde{R}$, R_d, $\tilde{S}$, and S_d are as defined below, the system (20) is asymptotically stable. ■

$$\begin{aligned} D^a_{1_p} &= -\frac{h}{2}A_d^{\mathrm{T}}(Q_{p-1} - Q_p) + \frac{h}{2}(R_{N,p-1} - R_{Np});\ \ D^a_{0_p} = -\frac{h}{2}A^{\mathrm{T}}(Q_{p-1} - Q_p) - \frac{h}{2}(R_{0,p-1} - R_{0p}) \\ D^a_p &= \begin{pmatrix} D^a_{0p} \\ D^a_{1p} \end{pmatrix}, p = 1, \cdots N \\ D^s_{1_p} &= \frac{h}{2}A_d^{\mathrm{T}}(Q_{p-1} + Q_p) - \frac{h}{2}(R_{N,p-1} + R_{Np});\ D^s_{0p} = \frac{h}{2}A^{\mathrm{T}}(Q_{p-1} + Q_p) + \frac{h}{2}(R_{0,p-1} + R_{0p}) - (Q_{p-1} - Q_p) \\ D^s_p &= \begin{pmatrix} D^s_{0p} \\ D^s_{1p} \end{pmatrix}, p = 1, \cdots N;\ D^s = (D^s_1\ D^s_2\ \cdots\ D^s_N) \end{aligned}$$

$$
\begin{aligned}
R_{d_{pq}} &= h(R_{p-1,q-1}-R_{pq});\;\; R_d = \begin{pmatrix} R_{d_{11}} & \cdots & R_{d_{1N}} \\ \vdots & \ddots & \vdots \\ R_{d_{N1}} & \cdots & R_{d_{NN}} \end{pmatrix};\;\; \tilde{R} = \begin{pmatrix} R_{00} & R_{01} & \cdots & R_{0N} \\ \vdots & \dots & \dots & \vdots \\ R_{N0} & R_{N1} & \cdots & R_{NN} \end{pmatrix} \\
S_{d_p} &= S_{p-1}-S_p;\; S_d = \mathrm{diag}\,(S_{d_1}\; S_{d_2}\; \cdots\; S_{d_N});\; \tilde{S} = \mathrm{diag}\left(\frac{1}{h}S_0\; \frac{1}{h}S_1\; \cdots\; \frac{1}{h}S_N\right);\; \tilde{Q} = (Q_0\; Q_1\; \cdots\; Q_N) \\
\Delta &= \begin{pmatrix} -PA - A^{\mathrm{T}}P - Q_0 - Q_0^{\mathrm{T}} - S_0 & Q_N - PA_d \\ Q_N^{\mathrm{T}} - A_d^{\mathrm{T}}P & S_N \end{pmatrix}
\end{aligned}
$$

We now consider piecewise linear time-delay systems of the form

$$
\dot{x}(t) = A_i x(t) + A_{d_i} x(t-d), \text{ for } x(t) \in X_i;\; x(t) = \phi(t),\, t \in [-d, 0] \tag{21}
$$

where $\{X_i\}_{i \in I} \subseteq \mathbb{R}^n$ is a partition of the state space into a number of closed (possibly unbounded) polyhedral cells with pairwise disjoint interior. In this section, we find Lyapunov functionals for this system, of the form (13), given by

$$
\begin{aligned}
V_i(x_t) &= \frac{1}{2}x(t)^{\mathrm{T}}P_i x(t) + x(t)^{\mathrm{T}}\int_{-d}^{0} Q(\zeta)x(t+\zeta)d\zeta \\
&+ \frac{1}{2}\int_{-d}^{0}\int_{-d}^{0} x(t+\zeta)^{\mathrm{T}}R(\zeta,\eta)x(t+\eta)d\eta d\zeta + \frac{1}{2}\int_{-d}^{0} x(t+\zeta)^{\mathrm{T}}S(\zeta)x(t+\zeta)d\zeta.
\end{aligned}
$$

for $x(t) \in X_i,\; X_{i\in I} \subseteq \mathbb{R}^n,\; I = \{1, 2, \cdots\}$ is a partition of the state space into a number of closed (possibly unbounded) cells.

For this Lyapunov functional to satisfy (11), for all i, $V_i(x_t) \geq \epsilon x^{\mathrm{T}}(t)x(t)$, for some $\epsilon > 0$ we get, as for discretized Lyapunov functions for linear systems with time-delay [11],

$$
\begin{pmatrix} P_i & \tilde{Q} \\ \tilde{Q}^{\mathrm{T}} & \tilde{R} + \tilde{S} \end{pmatrix} > 0. \tag{22}
$$

where $\tilde{Q}$, $\tilde{R}$ and $\tilde{S}$ are as defined above. For this Lyapunov functional, the time-derivative along a trajectory is given by

$$
\begin{aligned}
\dot{V}_i(x_t) &= x^{\mathrm{T}}(t)P_i[A_i x(t) + A_{d_i}x(t-d)] + x^{\mathrm{T}}(t)\int_{-d}^{0} Q(\zeta)\dot{x}(t+\zeta)d\zeta + \dot{x}^{\mathrm{T}}(t)\int_{-d}^{0} Q(\zeta)x(t+\zeta)d\zeta \\
&+ \int_{-d}^{0} d\zeta \int_{-d}^{0} x(t+\zeta)^{\mathrm{T}}R(\zeta,\eta)\dot{x}(t+\eta)d\eta + \int_{-d}^{0} x^{\mathrm{T}}(t+\zeta)S(\zeta)\dot{x}(t+\zeta)d\zeta
\end{aligned}
$$

Condition (12), that $\dot{V}(x_t) \leq -\epsilon x^{\mathrm{T}}(t)x(t)$ for some $\epsilon > 0$, is satisfied if

$$
\begin{pmatrix} \Delta_i & -D_i^s & -D_i^a \\ -D_i^{s^{\mathrm{T}}} & R_d + S_d & 0 \\ -D_i^{a^{\mathrm{T}}} & 0 & 3S_d \end{pmatrix} > 0 \tag{23}
$$

where R_d and S_d are as defined above, and

$$
\begin{aligned}
D^a_{1p_i} &= -\frac{h}{2}A^{\mathrm{T}}_{d_i}(Q_{p-1}-Q_p)+\frac{h}{2}(R_{N,p-1}-R_{Np});\ D^a_{0p_i}=-\frac{h}{2}A^{\mathrm{T}}_{i}(Q_{p-1}-Q_p)-\frac{h}{2}(R_{0,p-1}-R_{0p})\\
D^a_{p_i} &= \begin{pmatrix} D^a_{0p_i} \\ D^a_{1p_i} \end{pmatrix},\ p=1,\cdots N;\ D^a_i=\begin{pmatrix} D^a_{1_i} & D^a_{2_i} & \cdots & D^a_{N_i}\end{pmatrix}
\end{aligned}
$$

$$
\begin{aligned}
D^s_{1p_i} &= \frac{h}{2}A^{\mathrm{T}}_{d_i}(Q_{p-1}+Q_p)-\frac{h}{2}(R_{N,p-1}+R_{Np});\ D^s_{0p_i}=\frac{h}{2}A^{\mathrm{T}}_{i}(Q_{p-1}+Q_p)+\frac{h}{2}(R_{0,p-1}+R_{0p})-(Q_{p-1}-Q_p)\\
D^s_{p_i} &= \begin{pmatrix} D^s_{0p_i} \\ D^s_{1p_i} \end{pmatrix},\ p=1,\cdots N;\ D^s_i=\begin{pmatrix} D^s_{1_i} & D^s_{2_i} & \cdots & D^s_{N_i}\end{pmatrix}
\end{aligned}
$$

$$
\Delta_i = \begin{pmatrix} -P_iA_i-A_i^{\mathrm{T}}P_i-Q_0-Q_0^{\mathrm{T}}-S_0 & Q_N-P_iA_{d_i} \\ Q_N^{\mathrm{T}}-A^{\mathrm{T}}_{d_i}P_i & S_N \end{pmatrix}.
$$

For the Lyapunov functional $V(x_t)$ to be continuous at the cell partitions, we can parametrize the matrix P_i as $P_i=F_iTF_i$, $i\in I$, such that $F_ix=F_jx$, $x\in X_i\cap X_j$, $i,j\in I$. Since the dynamics given by A_i and A_{d_i} is only valid in the cell X_i, we can, as in [18] use the S-procedure [4] to obtain relaxed conditions for quadratic stability. Since the cells X_i are polyhedrons, we can construct matrices E_i, such that $E_ix\geq 0,\ x\in X_i,\ i\in I$. Then, for (11) and (12) to hold, we require

$$
\begin{pmatrix} P_i-E_i^{\mathrm{T}}W_iE_i & \tilde{Q} \\ \tilde{Q}^{\mathrm{T}} & \tilde{R}+\tilde{S}\end{pmatrix} > 0 \tag{24}
$$

$$
\begin{pmatrix} \Delta_i^* & -D_i^s & -D_i^a \\ -{D_i^s}^{\mathrm{T}} & R_d+S_d & 0 \\ -{D_i^a}^{\mathrm{T}} & 0 & 3S_d \end{pmatrix} > 0 \tag{25}
$$

where $\Delta_i^*=\Delta_i-\begin{pmatrix} E_i^{\mathrm{T}}U_iE_i & 0 \\ 0 & 0\end{pmatrix},\ i\in I$ for some matrices $U_i,\ W_i\in\mathbb{R}^n$, with nonnegative entries.

Theorem 5: If there exist matrices $T=T^{\mathrm{T}}$, $P_i=P_i^{\mathrm{T}}=F_i^{\mathrm{T}}TF_i$, U_i and W_i with nonnegative entries, $i\in I$, and Q_p, $S_p=S_p^{\mathrm{T}}$, $R_{pq}=R_{pq}^{\mathrm{T}}$, $p,q=0,\cdots N$, such that (24) and (25) are satisfied, with the notation described above, then the equilibrium $x^*(t)=0$ is asymptotically stable. ∎

As in [18], we could also relax the condition that the Lyapunov functional should be continuous across cell boundaries. It is only necessary that the value of the functional decreases at switching instants. Suppose the condition for the discrete mode to switch from mode j to k is given by $f_{jk}^{\mathrm{T}}=0$ (switching plane). Then, we require $P_j-P_k+f_{jk}t_{jk}^{\mathrm{T}}+t_{jk}f_{jk}^{\mathrm{T}}\geq 0$ for matrices P_j and P_k, and vectors t_{jk}, $j,k\in I$.

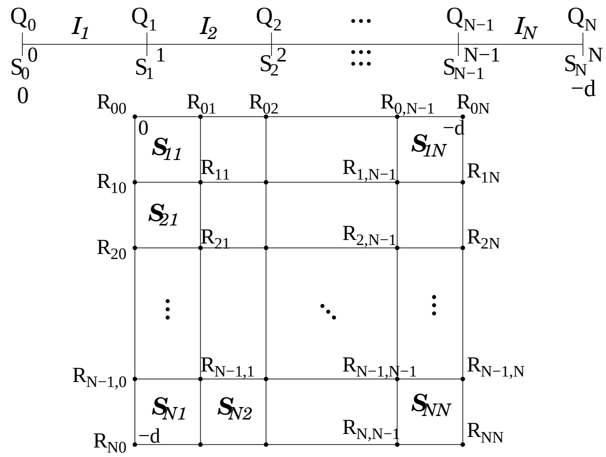

Fig. 9. Time-discretized Lyapunov functionals

C. Stability of XCP/RCP in bottleneck links with heterogeneous round trip delays

We now consider the stability of XCP in a single bottleneck topology, as considered in [21], but consider the case in which the users have different round trip delays.

We consider a single link of capacity C, traversed by N flows. Each sender i has a round trip delay of d_i, which we assume is entirely in the backward flow (from the router back to the sender). The aggregate traffic rate is $y(t) = \sum r_i(t)$,where $r_i(t)$ is the sending rate of user i at time t. In addition, there is a shuffled traffic rate, given by $h(t) = 0.1y(t)$. The router sends some aggregate feedback every control interval $\bar{d}$, where $\bar{d} = \sum d_i/N$, which reaches the sender i after a delay d_i.The router aggregate feedback is given by $\phi(t) = \frac{1}{\bar{d}^2}\left(-\alpha\bar{d}(y(t) - C) - \beta q(t)\right)$. When it receives the acknowledgment, each sender updates its window $w_i(t)$, where $w_i(t) = r_i(t)d_i$, using an Additive Increase Multiplicative Decrease policy. This implies

$$\dot{r}_i(t) \quad = \quad \tfrac{1}{N}\left([\phi(t-d_i)]^+ + h(t-d_i)\right) - \tfrac{r_i(t-d_i)}{y(t-d_i)}\left([-\phi(t-d_i)]^+ + h(t-d_i)\right)$$

At the equilibrium for a bottlenecked link in the XCP system, all flows i have equal rates [25]. Using this to simplify the above equations, we find that for a single bottleneck topology with heterogeneous delays, the system is described by the following equations.

$$\begin{aligned} \dot{y}(t) &= -\tfrac{\alpha}{N\bar{d}}\textstyle\sum_i (y(t-d_i) - C) - \tfrac{\beta}{N\bar{d}^2}\textstyle\sum_i q(t-d_i) \\ \dot{q}(t) &= \begin{cases} y(t) - C & , \quad q(t) > 0 \\ [y(t) - C]^+ & , \quad q(t) = 0 \end{cases} \end{aligned} \tag{26}$$

As an example, we consider a link with 3 flows, with delays of 100 ms, 200 ms and 400 ms. We then find the range of values of the parameters α and β for which the protocol can be proved to be stable for these heterogeneous delays. For the sake of brevity, the details of the LMI formulation are not given in this paper, and can be found in [2]. The region is plotted in Fig. 8.

REFERENCES

[1] M. Allman. A web server's view of the transport layer. *ACM Computer Communication Review*, 30(5), October 2000.

[2] H. Balakrishnan, N. Dukkipati, N. McKeown, and C. J. Tomlin. Stability analysis of switched hybrid time-delay systems - analysis of the Rate Control Protocol. July 2004. Working paper. `http://sun-valley.stanford.edu/hybrid/`.

[3] D. Bansal and H. Balakrishnan. Binomial Congestion Control Algorithms. In *Proceedings of INFOCOM 2001*, Anchorage, AK, April 2001.

[4] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. SIAM, Philadelphia, 1994.

[5] M.S. Branicky. Multiple Lyapunov functions and other analysis tools for switched and hybrid systems. *IEEE Transactions on Automatic Control*, 43(4):475–482, April 1998.

[6] Y. Chait, C.V. Hollot, V. Misra, H. Han, and Y. Halevi. Dynamic Analysis of Congested TCP Networks. In *Proceedings of the American Control Conference 2002*, , 2002.

[7] S. Deb and R. Srikant. Global stability of congestion controllers for the internet. *IEEE Transactions on Automatic Control*, 48(6):1055–1060, June 2003.

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

[9] S. Floyd. HighSpeed TCP for large congestion windows. December 2003. RFC 3649.

[10] S. Floyd and V. Jacobson. Random early detection gateways for congestion control. *IEEE/ACM Transactions on Networking*, 1(4):397–413, August 1993.

[11] K. Gu. A further refinement of discretized Lyapunov functional method for the stability of time-delay systems. *International Journal of Control*, 74(10):967–976, 2001.

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

[13] J.K. Hale and S.M.V. Lunel. *Introduction to Functional Differential Equations*. Springer-Verlag, New York, 1993.

[14] H. Han, C. V. Hollot, Y. Chait, and V. Misra. TCP Networks Stabilized by Buffer-Based AQMs. In *Proceedings of INFOCOM 2004*, , April 2004.

[15] C.V. Hollot, V. Misra, D. Towsley, and W-B. Gong. A control theoretic analysis of RED. In *Proceedings of IEEE INFOCOM 2001*, Anchorage, Alaska, April 2001.

[16] V. Jacobson. Congestion avoidance and control. In *Proceedings of ACM SIGCOMM 1988*, August 1988.

[17] C. Jin, D.X. Wei, and S.H. Low. FAST TCP: motivation, architecture, algorithms, performance. In *Proceedings of IEEE INFOCOM 2004*, March 2004.

[18] 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.

[19] R. Johari and D. Tan. End-to-end congestion control for the internet: delays and stability. *IEEE/ACM Transactions on Networking*, 9(6):818–832, 2001.

[20] Srikanth Kandula, Dina Katabi, Bruce Davie, and Anna Charny. Walking the Tightrope: Responsive Yet Stable Traffic Engineering. In *Proceedings of ACM SIGCOMM 2005*, Philadelphia, PA, August 2005.

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

[22] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control in communication networks: shadow prices, proportional fairness, and stability. *Journal of the Operations Research Society*, 49:237–252, 1998.

[23] T. Kelly. Scalable TCP: Improving performance in highspeed wide area networks. *Computer Communication Review*, 32(2), April 2003.

[24] V. Kulkarni, M. Jun, and J. Hespanha. Piecewise quadratic Lyapunov functions for time-delay hybrid systems. In *Proceedings of the American Control Conference 2004*, Boston, Massachusetts, June 2004.

[25] S.H. Low, L.L.H. Andrew, and B.P. Wydrowski. Understanding XCP: Equilibrium and fairness. In *Proceedings of IEEE INFOCOM 2005*, 2005.

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

[27] V. Misra, W-B. Gong, and D. Towsley. A fluid-based analysis in a network of AQM routers supporting TCP flows with an application to RED. In *Proceedings of ACM SIGCOMM 2000*, Stockholm, Sweden, September 2000.

[28] A. Papachristodoulou. Global Stability Analysis of a TCP/AQM Protocol for Arbitrary Networks with Delay. In *Proceedings of IEEE Conference on Decision and Control*, Paradise Island, Bahamas, December 2004.

[29] K.K. Ramakrishnan and S. Floyd. The addition of Explicit Congestion Notification (ECN) to IP. September 2001. RFC 3168.

[30] S. Sastry. *Nonlinear Systems: Analysis, Stability and Control*. Springer-Verlag, New York, 1999.

[31] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhauser, Boston, 2004.

[32] J.F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11–12:625–653, 1999. Version 1.05 available from `http://sedumi.mcmaster.ca/`.

[33] G. Vinnicombe. On the stability of networks operating TCP-like congestion control. In *Proceedings of IFAC 2002*, Barcelona, Spain, 2002.

[34] L. Ying, G. Dullerud, and R. Srikant. Global Stability of Internet Congestion Controllers with Heterogeneous Delays. In *Proceedings of the American Control Conference 2004*, Boston, Massachusetts, June 2004.