

Robust Stability Analysis of Smith Predictor-based Congestion Control Algorithms for Computer Networks[☆]

Luca De Cicco^{a,*}, Saverio Mascolo^a, Silviu-Iulian Niculescu^b

^aPolitecnico di Bari, Dipartimento di Elettrotecnica ed Elettronica, Via Orabona, 4, Bari, Italy

^bLaboratoire des Signaux et Systèmes (L2S), CNRS-SUPELEC, 3 rue Joliot-Curie 91192 Gif-sur-Yvette cedex, France

Abstract

Congestion control is a fundamental building block in packet switching networks such as the Internet due to the fact that communication resources are shared. It has been shown that the plant dynamics is essentially made of an integrator plus time delay and that a proportional controller plus a Smith predictor defines a simple and effective controller. It has been also shown that the today running TCP congestion control can be modelled using a Smith predictor plus a proportional controller. Due to the importance of this control structure in the field of data network congestion control, we analyze the robust stability of the closed-loop system in the face of delay uncertainties that in data networks are present due to queuing. In particular, by applying a geometric approach, we derive a bound on the proportional controller gain which is necessary and sufficient to guarantee the closed-loop stability for a given bound on the delay uncertainty.

Keywords: Time-delay systems, robust stability, congestion control, Smith predictor

1. Introduction

Time delays are often present in feedback control systems due to reasons such as the transport of material or information. From the control theoretic point of view, it is well-accepted that, quite often, an increase of the time delay may lead to instability of the closed-loop system and to performance degradation as well (see, e.g. Niculescu (2001), Niculescu and Michiels (2007), Zhong (2006) and the references therein). In such cases, the design of a finite-dimensional controller, such as the classical PID, is very challenging since the closed-loop system has an infinite number of characteristic roots and the resulting controller could provide an unacceptable sluggish closed-loop dynamics (Astrom and Hagglund, 1995).

The Smith principle (Smith, 1959) is a classical approach which is often employed to design effective infinite-dimensional controllers for time delay systems using an appropriate transformation of the control scheme which takes the delay out of the loop (see, for instance, Niculescu and Michiels (2007) for further details). Such a controller design proved its interest in various applications covering congestion mechanisms in communication networks (Mascolo, 1999), motion synchronization in various network configurations (Cheong et al., 2009a) or collaborative simulations in ring-like networks (Cheong et al., 2009b). It is known that, by assuming the exact knowledge of both the plant model and time delay, controllers

designed using a Smith predictor are very effective in counteracting the effect of time delays. It is worth to mention that such a method is less effective in the presence of modelling errors in the delay terms. Therefore, robustness issues of the Smith predictor with respect to uncertainties in the knowledge of the time delay have been extensively studied since 1980 see, for instance, Palmor (1980), Yamanaka and Shimemura (1987), Gudin and Mirkin (2007).

Indeed, the Internet represents a challenging case study in the context of time delay systems due to the presence of delays that are caused by the propagation of the information, which is sent in form of data packets, from a source to a destination through a series of communication links and router queues.

A cornerstone component of the Internet protocol stack is the end-to-end congestion control which has been implemented in the Transmission Control Protocol (TCP) by Van Jacobson in order to avoid congestion and preserve network stability (Jacobson, 1988).

Basically, the goal of a network congestion control algorithm is to adequately throttle the input rate at each source of a connection so that router buffer overflows are avoided. The congestion control proposed by V. Jacobson, and introduced in the TCP, proposes to address this problem implementing a distributed end-to-end algorithm that closes the loop at the end-point of the connection and does not require any explicit feedback from the routers. The only (implicit) feedback signal considered by TCP is represented by packet loss events that are interpreted by the congestion control algorithm as a synonymous of network congestion.

In the last decade, the issue of modelling the TCP congestion control algorithm has gained a great deal of attention in the scientific community. Indeed, fluid flow models play an important

[☆]This paper is the extended version of a paper that has been presented at the 8th IFAC Workshop on Time-Delay Systems, Sinaia, Romania, September 2009.

*Corresponding author. Tel.: +39 080 5963851; Fax: +39 080 5963410

Email addresses: ldcicco@gmail.com (Luca De Cicco),
mascolo@poliba.it (Saverio Mascolo),
Silviu.Niculescu@lss.supelec.fr (Silviu-Iulian Niculescu)

role in investigating the dynamics of the TCP flows subject to the Van Jacobson algorithm (see Hollot et al. (2002)).

In Mascolo (1999), a simple model of the plant made of an integrator (modelling the bottleneck queue) plus two time delays (modelling forward and backward delays), has been proposed along with a Smith predictor plus a proportional controller to design a congestion control algorithm. An important contribution of the paper is the proof that a Smith predictor controller with a proportional gain models the congestion control law which is employed in the today running TCP congestion control algorithm. Moreover, the model presented in Mascolo (1999) has been employed in Grieco and Mascolo (2004) to design and implement a *rate-based* congestion control algorithm which has been found to produce flows that are TCP-friendly.

A similar plant model is employed in Quet et al. (2002) to design a rate-based congestion control algorithm implemented at the router that is robust to uncertain time-delays by employing the \mathcal{H}^∞ technique. With respect to Mascolo (1999), the controller parameters in Quet et al. (2002) are quite complex to derive and, at the best of authors' knowledge, no real implementations are currently available to assess the effectiveness of the proposed solution.

It is well-known that the measurement of the plant time delay to be used in the Smith predictor can be affected by uncertainties due to the fact that the time delay is made of a constant propagation delay plus time-varying queuing delays. To the purpose, the TCP estimates the Round Trip Time (RTT) in order to set the retransmission timeout (RTO) which is needed for detecting heavy congestion episodes in the network. The RTT is defined as the time that elapses from when a segment is sent until the corresponding acknowledgment segment is received by the sender. In the standard TCP implementation (Postel, 1981), the RTT is measured each RTT seconds, whereas no measurements are taken on retransmitted segments in order to avoid spurious timeouts (see Karn and Partridge (1987)). For these reasons the standard TCP provides a measurement of RTT that may be affected by errors. In order to overcome this issue an optional scheme has been proposed and standardized in Jacobson et al. (1992) which makes use of timestamps in an optional field of the TCP header. However, even if the timestamp option is employed by both peers of the communication, the granularity chosen for TCP timestamps is implementation-dependent and can be itself a source of significant errors. In a recent work, Veal et al. (2005) carried out an extensive measurement campaign on RTTs. Authors used 500 servers and found that 76% of the servers had timestamping option enabled, and out of these servers 37% used a 100 ms granularity, 55% a 10 ms granularity and only 7% of them had a granularity of 1 ms.

A preliminary study on robust stability of the proportional Smith predictor used for congestion control in data networks (Mascolo, 1999) has been carried out by using the Nyquist criterion in Mascolo (2003). It revealed that in order to guarantee asymptotic stability it is sufficient that $\Delta < 1/k$ where Δ represents the delay uncertainty and k is the gain of the proportional controller.

The goal of this paper is to provide a *complete characterization of the robust stability of congestion control model* intro-

duced in Mascolo (1999) by applying the geometric approach idea developed in Gu et al. (2005). In the following we will show that the geometric approach is very simple to apply and is able to give an easy to understand "picture" of the robustness of the system in the case delay uncertainties are present. The parameter-space to be considered is represented by the nominal delay and the corresponding delay error, which can be both positive and negative (Morarescu et al., 2007). It is worth to mention that the particular structure of the closed-loop scheme allows an appropriate re-scaling of the system's parameters that will be explicitly exploited in the robustness analysis improving thus the bound proposed in Mascolo (2003).

The rest of the paper is organized as follows: in Section 2 we briefly review the model of the closed loop congestion control in a generic packet switching network presented in Mascolo (1999); in Section 3 we apply the geometrical approach developed in Morarescu et al. (2007) in order to find the stability crossing curves of the system; in Section 4 we present the robust stability analysis; in Section 5 some simulations are presented to support the theoretical results obtained; finally Section 6 concludes the paper.

2. Congestion Control Model

A network connection is basically made by a set of communication links and store-and-forward nodes (routers) where packets are enqueued before being routed to the destination. Congestion can arise when packets arrive at the router at a rate $r(t)$ which is above the capacity of the output link so that the router queue builds up until it is full and it starts to drop packets.

In Mascolo (1999), a model of the Internet flow and congestion control as a time delay system is provided. In particular, the model consists of a feedback loop in which two time delays are present as it is shown in Figure 1: T_1 models the propagation time of a packet from source to the bottleneck queue and T_2 models the propagation time from the bottleneck to the destination and then back to the sender. The round trip time of the connection is $T = T_1 + T_2$. It has to be noted, that considering a constant RTT is a modelling simplification that is often employed in the literature to make the model more tractable. Throughout the rest of the paper constant time-delays are considered, and in Section 5 simulation results, obtained when time-varying time-delays are present in the feedback loop, are shown to validate the theoretical findings discussed in this paper.

The simple integrator $1/s$ models the bottleneck queue that is filled (or drained) by the rate mismatch $r(t) - b(t)$, where $b(t)$ is the bottleneck available bandwidth.

The controller is a proportional Smith predictor with gain k that computes the rate $r(t)$ to match the available bandwidth $b(t)$ and to produce a stable output. The reason for using a simple proportional controller is that in this way the closed-loop dynamics can be made that of a first-order system with time constant $1/k$ delayed by T_1 . Thus, the step response of the system can be made faster by increasing the proportional gain k providing an always stable system without oscillations or overshoots.

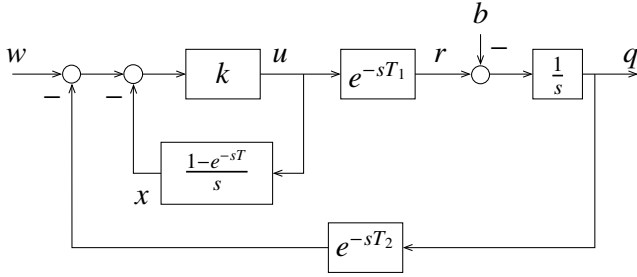


Figure 1: Functional block of the congestion control model

This choice provides a controller in which only one design parameter, i.e. the gain k , has to be tuned having a direct influence on the dynamics of the output. It is worth noting that in Palmor and Blau (1994), by following optimal control arguments, a primary controller that inverts the plant model and that adds a pole in the origin is proposed, with the aim of making the closed-loop transfer function a first order system plus dead time that is stable for any positive value of the gain. In the case of the integral plant that is considered in this paper, this means that the primary controller is a simple gain k . Finally, the input signal $w(t)$ models the congestion window (cwnd) or the advertised window, that is used by the congestion control algorithm to bound in-flight packets (Mascolo, 1999). A remarkable feature of this model is that different variants of TCP congestion control algorithms can be modelled in a unified framework by proper input shaping of the proportional Smith predictor controller (Mascolo, 2006). Moreover, in Mascolo (1999) it has been shown that the Smith predictor models the self-clocking property of TCP congestion control that is essential for the implementation of an effective congestion control algorithm (Bansal et al., 2001). Finally, we remark that a Smith predictor controller is recommended when designing a congestion control algorithm for data networks, since using PID controllers would provide an unacceptable sluggish system due to large delays involved in communication networks (see Astrom and Hagglund (1995) and Mascolo (1999)).

An issue that affects a general system modelled by Figure 1 is that the system is not internally stable. Thus, if the inner loop is affected by a numerical noise or a mismatch between the initial conditions of the integrator modelling the plant and the integrator employed in the Smith predictor exists, then the output of the inner loop would diverge. However, in the considered system, the algorithm that realizes the inner loop is not affected by such issues. In the time domain the output $x(t)$ of the inner loop is given by:

$$x(t) = \int_0^t u(\xi) d\xi - \int_0^t u(\xi - T) d\xi. \quad (1)$$

However, the difference between the two integrals in (1) is not numerically computed in the algorithm, since this quantity is actually the difference between the data sent up to time t and the data sent up to time $t - T$. This quantity is already known without any error (because data are sent as packets of known integer size), whereas the input rate $u(t)$ is the variable that could

be computed as ratio of the data sent over a certain time interval, but this variable is never used or computed in the algorithm.

A look at the algorithm employed to compute the integral of $u(t)$ shows why no computational error can be introduced by the inner loop:

$$\text{At } t_0 = 0 : \text{data_sent} \leftarrow 0$$

$$\text{At time } t_k : \text{data_sent}(t_k) \leftarrow \text{data_sent}(t_{k-1}) + p_k$$

In the algorithm, `data_sent` is the integral computed in the inner loop which is updated each time a packet of size p_k is sent at time t_k . No error can be introduced here since packet sizes are exactly known, i.e. packet size p_k is precisely known and it is stored in an integer variable, so that its machine representation is not affected by any inaccuracy. For what concerns the mismatch on the initial conditions of the integrators, this case cannot occur since the integrator models the connection buffer that, at the time of starting the connection (i.e. $t = 0^-$), is always zero since no data has been sent yet.

Thus, the only practical issue due to the pole in the origin is that step-like disturbances cannot be rejected at steady state (see Astrom et al. (1994)). Nevertheless, in the case of network congestion control rejection of the disturbance is not the primary issue. The goal of the control is to guarantee that the queue reaches a steady state and the input rate $r(t)$, that is the output of the Smith-Predictor, matches the available bandwidth with a zero steady state error.

Another issue of the model depicted in Figure 1 is that, when a Smith predictor controller is employed, model mismatches are known to affect the closed loop dynamics. In this case, the only source of mismatch between the model and the actual plant is the entity of the delay (see Section 1), whereas the model of the bottleneck queue is an integrator and does not add any uncertainty. In the next sections we will give simple tuning rules for the design parameter k in order to retain asymptotic stability when the measurement of time delay T is uncertain.

3. Stability Crossing Curves in the Parameters Space

3.1. Review of the geometrical approach

We start by briefly reviewing the geometrical approach developed in Morescu et al. (2007) which we will employ to analyze the robust stability of the considered system. The reader is advised to refer to Gu et al. (2005) for a complete description of the method. We denote with $a(s; \tau_1, \tau_2)$ the characteristic function of the closed-loop system where τ_1 represents the nominal delay used in the Smith predictor and $\tau_2 = \tau_1 + \Delta$ represents the actual plant delay affected by a bounded mismatch Δ . It is easy to show that the characteristic function in this case is given by:

$$a(s; \tau_1, \tau_2) = 1 - h(s)e^{-\tau_1 s} + h(s)e^{-\tau_2 s} \quad (2)$$

where $h(s)$ is the transfer function of the closed loop system when no delays are present in the loop:

$$h(s) = \frac{C(s)G_0(s)}{1 + C(s)G_0(s)}$$




Figure 2: Triangle formed by the three vectors when the characteristic equation holds

with $G_0(s)$ being the delay free plant and with $C(s)$ being the controller transfer function.

In order to analyze the stability of the system, we look for the solutions of the characteristic equation:

$$a(j\omega; \tau_1, \tau_2) = 0. \quad (3)$$

In this way, we are able to find all the conditions under which the closed-loop system has at least one pole on the imaginary axis. The geometrical approach relies on the observation that the three terms of the characteristic function (2) can be seen as vectors in the complex plane. Therefore, the equality $a(s; \tau_1, \tau_2) = 0$ can be represented in the complex plane via an isosceles triangle as it is shown in Figure 2. Thus, equation (3) is equivalent to the following three conditions:

1. The triangular inequality must hold for the triangle shown in Figure 2, which implies that:

$$|h(j\omega)| \geq \frac{1}{2} \quad (4)$$

2. Equation (3) must satisfy the phase rule;
3. The sum of the internal angles of the isosceles triangle must be equal to π ;

The solution of (4), which does not depend on time delays τ_1 or τ_2 , forms the *frequency crossing set* Ω which is the union of a finite number N of intervals of finite length $\Omega_1, \Omega_2, \dots, \Omega_N$ (Gu et al., 2005). It has to be noted that, when $\lim_{\omega \rightarrow \infty} |h(j\omega)| \geq 1/2$, the closed-loop system would be practically unstable even if $h(s)$ is asymptotically stable, i.e. instability arises for any given delay mismatch Δ (see Palmor (1996)). For an extensive study about the practical stability issues of Smith predictor controllers the reader is advised to refer to Niculescu and Michiels (2007). Finally, if (4) is not satisfied for any ω and $\lim_{\omega \rightarrow \infty} |h(j\omega)| < 1/2$, the frequency crossing set Ω is empty and the closed-loop system would be asymptotically stable independent of the entity of the time delays.

For any $\omega > 0$ which belongs to the frequency crossing set there exists at least a pair (τ_1, τ_2) in the parameters space such that the system has at least one imaginary pole. The conditions 2 and 3 imply that for all $\omega \in \Omega$ all the couples $(\tau_1, \tau_2) \in \mathbb{R}_+^2$

satisfying $a(j\omega; \tau_1, \tau_2) = 0$ can be found using the following equations:

$$\tau_1^{u\pm} = \frac{\angle h(j\omega) + 2u\pi \pm q(\omega)}{\omega}, \quad (5)$$

$$\tau_2^{v\pm} = \frac{\angle h(j\omega) + (2v-1)\pi \mp q(\omega)}{\omega}, \quad (6)$$

where u and v are integers such that the corresponding $\tau_1^{u\pm}, \tau_2^{v\pm}$ are non negative and $q(\omega)$ represents the internal angle of the isosceles triangle:

$$q(\omega) = \arccos\left(\frac{1}{2|h(j\omega)|}\right).$$

In order to understand the meaning of equations (5) and (6) let us fix $u = \bar{u}$, $v = \bar{v}$ and consider the set $\Omega_i \subseteq \Omega$: if ω varies in Ω_i and we evaluate (5) and (6) for both positive and negative signs we obtain two curves in the parameter space (τ_1, τ_2) which we denote $\mathcal{T}_{\bar{u}, \bar{v}}^{i+}$ and $\mathcal{T}_{\bar{u}, \bar{v}}^{i-}$ respectively. It is worth noting that the curves $\mathcal{T}_{\bar{u}, \bar{v}}^i = \mathcal{T}_{\bar{u}, \bar{v}}^{i-} \cup \mathcal{T}_{\bar{u}, \bar{v}}^{i+}$ can be either open curves or closed curves depending on the set Ω_i we are considering. In particular, it is easy to show that if the left end of Ω_i is 0 then the associated curve is an *open curve* with both ends approaching ∞ when $\omega \rightarrow 0$. On the other hand, if the left end of Ω_i is not 0 then $\mathcal{T}_{\bar{u}, \bar{v}}^i$ is a *closed curve* (Gu et al., 2005).

We define the *stability crossing curves* \mathcal{T} in the τ_1, τ_2 plane as the union of all the curves $\mathcal{T}_{u,v}^i$ when $i \in \{1, \dots, N\}$, and u and v vary in the set of integers.

Finally, it is important to point out that when a stability crossing curve is crossed in the τ_1, τ_2 plane, at least one pair of poles cross the imaginary axis of the complex plane Gu et al. (2005).

3.2. Stability crossing curves of the computer network congestion control model

In order to characterize the impact of the delay uncertainty on the stability of the considered feedback system we apply the geometric approach we have reviewed in Section 3.1. It is worth to notice that the delay-free model of the plant is $G_0(s) = 1/s$ and the controller transfer function is $C(s) = k$.

We suppose that the system described in Section 2 is affected by a delay uncertainty Δ which is bounded by $\delta > 0$, i.e. $|\Delta| < \delta$. By considering the delay uncertainty, the characteristic equation of the system can be rewritten as follows:

$$1 + \frac{k}{s} - \frac{k}{s} e^{-\tau_1 s} (1 - e^{-\Delta s}) = 0, \quad (7)$$

where $\tau_1 = T_1 + T_2$ represents the nominal round trip time (RTT) of the considered connection, which is used in the Smith predictor, and $\tau_2 = \tau_1 + \Delta$ is the actual plant time delay.

By multiplying by $s/(s+k)$ both sides of (7) we obtain:

$$1 - \frac{k}{s+k} e^{-\tau_1 s} + \frac{k}{s+k} e^{-(\Delta+\tau_1)s} = 0, \quad (8)$$

so that by considering $h(s) = k/(s+k)$ (8) is in the form of (2). We are interested in characterizing the stability of the system when τ_1, τ_2 and k vary in \mathbb{R}_+ .




Figure 3: Stability crossing curves for the considered system

By making the change of variable $z = s/k$ we obtain:

$$1 - \frac{1}{z+1}e^{-h_1z} + \frac{1}{z+1}e^{-h_2z} = 0, \quad (9)$$

where $h_1 = k\tau_1$ and $h_2 = k\tau_2$, which reduces the free parameters to two. It is worth to notice that the transformation from (8) to (9) simply involves a *scaling* of the closed-loop eigenvalues by $1/k$, thus indicating a natural trade-off between gain and delay since when k increases the closed loop poles approach to the imaginary axis Niculescu and Michiels (2004).

We are now ready to study the stability of the original system in the h_1, h_2 plane regardless the value of the proportional gain k .

First of all, by applying (4) we find that the crossing set is made by the single interval $\Omega = [0, \sqrt{3}]$ which means that the stability crossing curves in the h_1, h_2 plane are open curves which extend to infinity when $\omega \rightarrow 0$. By using (5) and (6) the stability crossing curves of the considered system are parametrized as follows:

$$h_1^{u\pm}(\omega) = \frac{-\arctan \omega + 2u\pi \pm \arccos\left(\frac{\sqrt{1+\omega^2}}{2}\right)}{\omega}, \quad (10)$$

$$h_2^{v\pm}(\omega) = \frac{-\arctan \omega + (2v-1)\pi \mp \arccos\left(\frac{\sqrt{1+\omega^2}}{2}\right)}{\omega}. \quad (11)$$

Figure 3 shows the stability crossing curves of the considered system. We start by considering the h_2 axis ($h_1 = 0$), which means that we are employing a simple proportional controller without the Smith predictor. By starting from the origin and increasing the value of h_2 the first curve is crossed at $h_2 = \pi/2$ which means that the system becomes unstable for $h_2 > \pi/2$ as expected. On the other hand, the axis h_1 represents the system

in which no delay affects the plant, but the Smith predictor is in the controller. Figure 3 shows that the system is stable for all the delays in the Smith predictor.

It is worth to mention that points below (above) the bisector represent the case of negative (positive) delay mismatch, whereas points lying on the positive bisector represent the case of perfect matching of nominal delay τ_1 with the actual delay τ_2 . Indeed, if we move on this line no curves will be crossed since the Smith predictor in this case provides a stable system regardless the value of the proportional gain k .

4. Robust Stability Analysis

In this Section, we will develop an analysis of the robust stability of the considered system by using the stability crossing curves we have shown in the previous Section. We already know that the considered system is always asymptotically stable for any delay τ_1 and any proportional gain k as far as the delay uncertainty is zero due to the perfect compensation of the time delay τ_1 provided by the Smith predictor. In the h_1, h_2 plane this condition simply means that the system is asymptotically stable on all the positive bisector.

In order to characterize the robustness of the system in the face of delay uncertainties, we compute the maximum delay mismatch which still preserves stability. Thus, the problem here is to look for the maximum deviation δ with respect to a generic point (τ_1^*, τ_1^*) with $\tau_1^* \geq 0$ which lies on the positive bisector such that the system is stable for any (τ_1, τ_2) which satisfies:

$$|\tau_2 - \tau_1^*| < \delta$$

We remark that solving the maximum admissible delay uncertainty problem is equivalent to find the minimum distance between the stability crossing curves and a generic point on the positive bisector of the h_1, h_2 plane.

Thus for any $\tau_1^* > 0$ we have to solve:

$$\delta(\tau_1^*) = \min_{u,v} \min_{\tau_2^{\pm} \in \mathcal{F}} |\tau_2^{v\pm} - \tau_1^*| \quad (12)$$

so that the maximum delay to retain stability is:

$$\delta = \min_{\tau_1^* \in \mathbb{R}_+} \delta(\tau_1^*) \quad (13)$$

Proposition 1. *A necessary and sufficient condition for the asymptotic stability of the system independent of the value of the nominal delay τ_1 is:*

$$|\Delta| < \frac{\alpha}{k} \quad (14)$$

where Δ is the delay uncertainty, $\alpha \cong 1.4775$ and k is the proportional gain of the controller.

PROOF. We start by considering the stability crossing curves in the parameters space h_1, h_2 . In order to find the minimum distance between the stability crossing curves and a generic point of positive bisector of the h_1, h_2 plane we evaluate the tangent




Figure 4: Stability regions for the considered system. Stability region found in Proposition 1 is marked in gray.

to the crossing curves with direction parallel to the positive bisector:

$$\frac{dh_2}{dh_1} = 1 \Leftrightarrow \frac{dh_2}{d\omega} \frac{d\omega}{dh_1} = 1 \Leftrightarrow \frac{dh_1}{d\omega} = \frac{dh_2}{d\omega} \quad (15)$$

To the purpose we look for a subset $\bar{\mathcal{T}}$ of the stability crossing curves \mathcal{T} that are the “closest” curves to the positive bisector. By considering a generic curve $\mathcal{T}_{u,v}$ and by applying (10) and (11) it turns out that for all u and v and for all $\omega \in \Omega$ it holds $h_2^{v+} - h_1^{u+} < h_2^{v-} - h_1^{u-}$ so that it is sufficient to consider only the curves $\mathcal{T}_{u,v}^+$ in the region $h_2 > h_1$ and the curves $\mathcal{T}_{u,v}^-$ in the region $h_2 < h_1$, since they are the closest ones to the positive bisector. Thus, we can refer without loss of generality to the generic curves of \mathcal{T} as $\mathcal{T}_{u,u+i}$ for all i and u in the integers. Straightforward computations on (10) and (11) give:

$$h_2^{u+i} - h_1^u > h_2^{u+i-1} - h_1^u$$

which means that when i decreases the curves $\mathcal{T}_{u,u+i}$ will move downwards in the h_1, h_2 plane.

Figure 4 shows the values of u and v for the curves $\mathcal{T}_{u,u}$ and $\mathcal{T}_{u,u+1}$. It is then easy to show that if we set $v = u$ we obtain the closest curves to the positive bisector in the region $h_2 < h_1$ whereas the curves with $v = u + 1$ are those which are closest to the positive bisector in the region $h_2 > h_1$. In conclusion we can restrict our search to the set :

$$\bar{\mathcal{T}} = \mathcal{T}_{u,u}^- \cup \mathcal{T}_{u,u+1}^+$$

for all u in the integers. Let us consider the region $h_2 > h_1$ i.e. we consider the subset $\mathcal{T}_{u,u+1}^+$. By considering (15) after straightforward computations we get the following equation:

$$\arccos\left(\frac{\sqrt{\omega^2 + 1}}{2}\right) + \frac{\omega^2}{\sqrt{\omega^2 + 1}\sqrt{3 - \omega^2}} +$$

$$+\pi\left(v - u - \frac{1}{2}\right) = 0 \quad (16)$$

with $\omega \in \Omega$. By letting $v = u + 1$ the equation (16) can be numerically solved to find the unique solution $\bar{\omega} \cong 1.3483$ rad/s in Ω which is independent of u . If we substitute this value in (10) and (11) we obtain:

$$h_1(\bar{\omega}) = h_1 = 4.6601u - 0.2654$$

$$h_2(\bar{\omega}) = h_2 = 4.6601v - 3.4480$$

Thus, all the points belonging to the the curves $\mathcal{T}_{u,u+1}^+$ having a tangent which is parallel to the positive bisector, lie on the line:

$$h_2 = h_1 + 1.4775 \quad (17)$$

For this reason we can conclude that the maximum uncertainty, in the h_1, h_2 coordinates is 1.4775. The proof is completed by recalling that $h_1 = k\tau_1$ and $h_2 = k\tau_2$ and that $\tau_2 = \tau_1 + \Delta$. Thus, we finally obtain:

$$h_2 - h_1 < 1.4775 \Rightarrow k\Delta < 1.4775 \Rightarrow \Delta < \frac{1.4775}{k} \quad (18)$$

It is worth to notice that the same procedure can be followed in the case $v = u$ which leads to the inequality:

$$h_1 - h_2 < 1.4775 \Rightarrow -k\Delta > 1.4775 \Rightarrow \Delta > -\frac{1.4775}{k} \quad (19)$$

Thus, by considering both (18) and (19) we obtain (14). In order to prove the necessity of the condition (14) let us consider the curves $\mathcal{T}_{u,u+1}$. The points of the curve $\mathcal{T}_{u,u+1}$ that correspond to the frequency $\omega = 1.3483$ rad/s lie on the line described by (17) so that the maximum delay uncertainty admissible for those points is exactly α/k . If we select a larger value for δ the system will become unstable at least on those points. This concludes the proof.

Remark 1. The fact that the maximum uncertainty allowed does not depend on the nominal delay τ_1 is a nice feature of the Smith predictor based controller. This makes the controller effective even with large delays. It is important to notice that Proposition 1 is a necessary and sufficient condition for the robust stability of the system that is valid independent of the value of the nominal delay τ_1 . This does not prevent that, for a specific value of τ_1 , the system is stable for a value of Δ that is greater than α/k . Nevertheless, Proposition 1 guarantees that if the maximum entity of the delay error Δ is known, the controller gain k can be tuned so that for any value of the nominal delay τ_1 the stability of the system is guaranteed.

Figure 5 shows $k\Delta$ as function of $k\tau_1$ and can be obtained by employing equations (10) and (11). The two curves shown in black delimit the stability region in the case $k\tau_1$ is variable, whereas the two black dashed lines represent the stability boundary independent of τ_1 , i.e. the robust stability condition (14). It is worth to notice that the stability bounds as function of $k\tau_1$ shown in Figure 5 recover the results found in Gudin and Mirkin (2007), where a Nyquist approach was employed. Finally, we remark that, although (14) can be conservative for

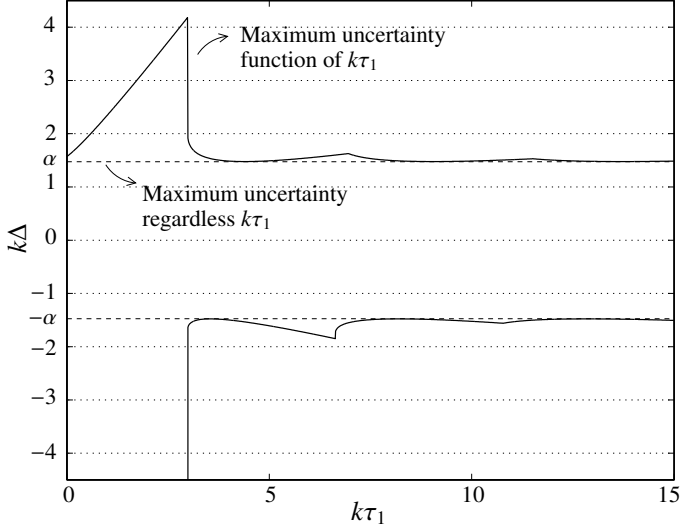


Figure 5: Maximum delay uncertainty to retain asymptotic stability as function of $k\tau_1$ (continuous black curves) or independent of τ_1 (dashed black lines)

some values of the nominal delay (see Figure 5), it represents a simple and practical rule that does not require the designer to make difficult assumptions regarding the nominal delay. In fact, the nominal delay, i.e. the RTT of the connection, can vary in a range between 0.01s in the case of a wired connection up to 1s and more in the case of wireless connections such as in the case of satellite paths.

Remark 2. The condition (14) expresses a trade-off between the maximum delay mismatch δ and the proportional gain that can be used to tune the controller gain k .

Remark 3. This result improves the robust stability condition $|\Delta| < 1/k$ found in Mascolo (2003) and in Morari and Zafiriou (1989), by using different approaches and related analytical arguments.

Proposition 2. *The system is stable, independent of the value of τ_1 , if the delay uncertainty Δ satisfies the following inequality:*

$$-\tau_1 < \Delta < -\tau_1 + \frac{\beta}{k} \quad (20)$$

with $\beta = 1.1188$.

PROOF. The proof follows the same arguments of Proposition 1, therefore it is omitted.

Remark 4. The condition (20) implicitly requires the delay uncertainty Δ to be negative, i.e. the nominal delay τ_1 should be always below the actual delay of the plant τ_2 . Thus, condition (20) has no particular meaning for the characterization of controller robustness, since the sign of the uncertainty is not known a priori.





Figure 6: Queue length (a) and input rate evolution (b) when $\tau_1 = 1$ sec, $k = 4$ sec $^{-1}$, $b(t) = 100$ packets/sec, $w = 150$ packets

5. Simulation results

In this Section we report simulation results obtained by using a SIMULINK model that implements the system depicted in Figure 1. The bottleneck available bandwidth has been set to vary as a step function starting at time $t = 1$ sec and having a final value of $b = 100$ packets/sec. The queue set-point is a step function starting at $t = 0$ sec with a final value of $w = 150$ packets. The gain of the controller has been set to $k = 4$ sec $^{-1}$ corresponding to a maximum delay uncertainty of $\delta \cong 0.37$ sec. The nominal RTT of the connection is 1 sec. Figure 6 reports the queue evolution $q(t)$ and the input rate $r(t)$ when the delay uncertainty is either zero, $\delta/2 \cong 0.185$ sec or $\delta \cong 0.37$ sec.

The figure shows, as expected, that the performance of the closed loop response degrades when the delay uncertainty increases. In particular, oscillations are present when the delay uncertainty is $\delta/2$ still providing an acceptable response, whereas when the delay uncertainty increases to the maximum

allowed value persistent oscillations occur.

In order to check if the bound on the gain k expressed in Proposition remains valid when time-varying round trip times τ_2 are present, we show simulation results obtained when the time delay varies as follows (Shustin and Fridman, 2007): 1) *slowly-varying* delay $\tau_2^s(t) = \tau_1 + \delta(1 + 1/2 \sin 0.1t)$; 2) *moderately-varying* delay $\tau_2^m(t) = \tau_1 + \delta(1 + 1/2 \sin 10t)$; 3) *fast-varying* delay $\tau_2^f(t) = \tau_1 + \delta(1 + 1/2 \sin t^2)$. It is worth to notice that to each of the three delays defined above corresponds a delay mismatch whose average values is δ , maximum value of $3/2\delta$ and minimum value $1/2\delta$. The settings of k , w , and $b(t)$ remain unvaried.

Figure 7 shows the queue dynamics in the three considered cases. The case of slow-varying delay, shown in Figure 7(a) shows the most remarkable oscillations due to the fact that the signal $\tau_2^s(t)$ is above the maximum allowed time delay for $\pi/0.1$ seconds. After a transient that last about 40 s, the queue length eventually reaches the steady state.

The queue dynamics obtained for moderate and fast-varying delays, shown in Figure 7(b) and Figure 7(c), exhibit similar behaviour. In particular, even if the average value of the delay is equal to the critical one, the queue reaches a steady state after around 20 s.

6. Conclusions

In this paper we have analyzed the robust stability of an important class of congestion control algorithms when delay uncertainties are present. By using the geometrical approach developed in Morarescu et al. (2007) a strict stability bound on the parameter of the controller has been found. Such a result suggests that congestion control algorithms that employ controllers made by a Smith predictor plus a proportional gain can be easily tuned in order to be robust with respect to a bounded delay uncertainty.

7. Acknowledgment

We thank the Associate Editor and the anonymous reviewers for the constructive comments that allowed us to improve the overall quality of the paper. Finally, we would like to thank Prof. Z. Palmor for his insightful suggestions and discussion about the paper.

References

- Astrom, K., Hagglund, T., 1995. PID Controllers: Theory, Design, and Tuning. Instrument Society of America.
- Astrom, K., Hang, C., Lim, B., 1994. A new Smith predictor for controlling a process with an integrator and long dead-time. IEEE transactions on Automatic Control 39, 343–345.
- Bansal, D., Balakrishnan, H., Floyd, S., Shenker, S., 2001. Dynamic Behavior of Slowly-Responsive Congestion Control Algorithms, in: Proc. of ACM SIGCOMM 2001, pp. 263–273.
- Cheong, J., Niculescu, S.I., Kim, C., 2009a. Motion synchronization control of distributed multisubsystems with invariant local natural dynamics. IEEE Computer Graphics and Applications 25, 382–398.

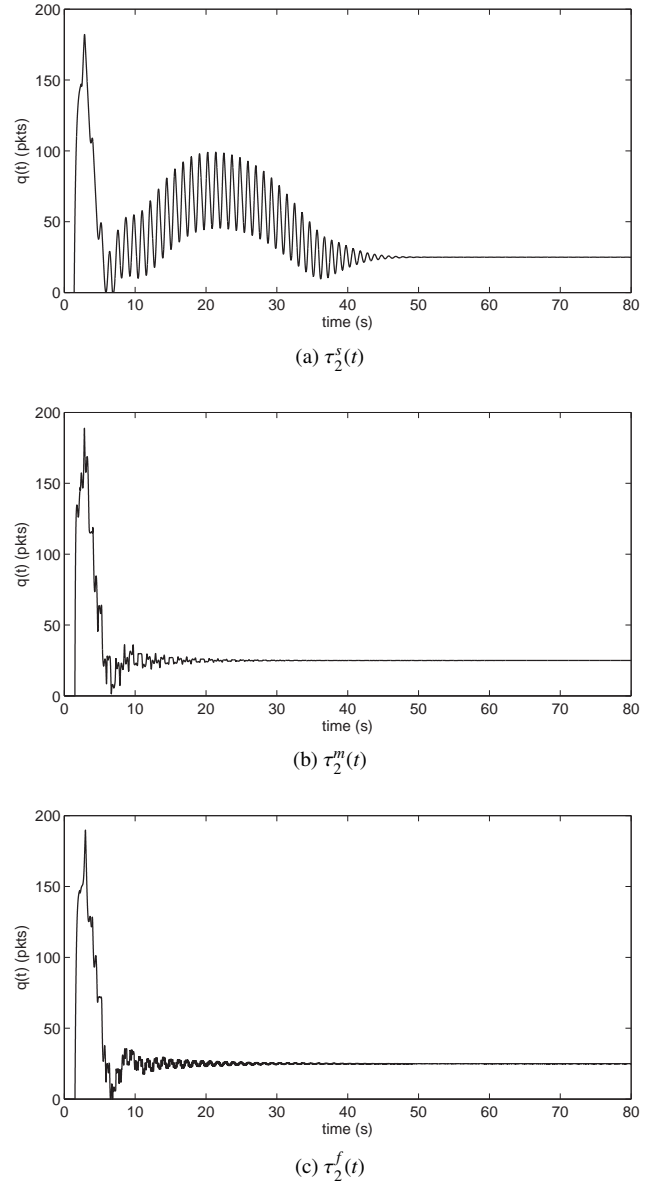


Figure 7: Queue length dynamics in the case of time-varying RTT, $k = 4 \text{ sec}^{-1}$, $b(t) = 100 \text{ packets/sec}$, $w = 150 \text{ packets}$

- Cheong, J., Park, J., Niculescu, S.I., 2009b. Physically based collaborative simulations under ring-like network configurations. *IEEE Computer Graphics and Applications* 29, 64–80.
- Grieco, L.A., Mascolo, S., 2004. Adaptive rate control for streaming flows over the internet. *ACM Multimedia Systems Journal* 9, 517–532.
- Gu, K., Niculescu, S.I., Chen, J., 2005. On stability crossing curves for general systems with two delays. *Journal of mathematical analysis and applications* 311, 231–253.
- Gudin, R., Mirkin, L., 2007. On the delay margin of dead-time compensators. *International Journal of Control* 80, 1316–1332.
- Hollot, C., Misra, V., Towsley, D., 2002. Analysis and design of controllers for AQM routers supporting TCPflows. *Automatic Control, IEEE Transactions on* 47, 945–959.
- Jacobson, V., 1988. Congestion avoidance and control, in: *Proc. of ACM SIGCOMM '88*, pp. 314–329.
- Jacobson, V., Braden, R., Borman, D., 1992. TCP Extensions for High Performance. RFC 1323, standard .
- Karn, P., Partridge, C., 1987. Estimating round-trip times in reliable transport protocols, in: *Proc. of ACM SIGCOMM '87*.
- Mascolo, S., 1999. Congestion control in high-speed communication networks using the Smith principle. Special Issue on “Control methods for communication networks” *Automatica* 35, 1921–1935.
- Mascolo, S., 2003. Modeling the Internet congestion control as a time delay system: a robust stability analysis, in: *Proc. of IFAC workshop on time-delay systems*.
- Mascolo, S., 2006. Modeling the Internet congestion control using a Smith controller with input shaping. *Control engineering practice* 14, 425–435.
- Morarescu, C., Niculescu, S.I., Gu, K., 2007. On the geometry of stability regions of Smith predictors subject to delay uncertainty. *IMA Journal of Mathematical Control and Information* 24, 411–423.
- Morari, M., Zafriou, E., 1989. *Robust Process Control*. Englewood Cliffs, NJ: Prentice Hall.
- Niculescu, S.I., 2001. Delay effects on stability. A robust control approach. volume 269 of *Lecture Notes in Control and Information Sciences*. Springer Verlag.
- Niculescu, S.I., Michiels, W., 2004. Stabilizing a chain of integrators using multiple delays. *Automatic Control, IEEE Transactions on* 49, 802–807.
- Niculescu, S.I., Michiels, W., 2007. Stability and stabilization of time-delay systems. An eigenvalue-based approach. SIAM: Philadelphia, USA.
- Palmor, Z., 1980. Stability properties of Smith dead-time compensator controllers. *International Journal of Control* 32, 937–949.
- Palmor, Z., 1996. Time-delay Compensation - Smith Predictor and its Modifications. in “*The Control Handbook*”, Levine W.S. editor, CRC press Inc., USA.
- Palmor, Z., Blau, M., 1994. An auto-tuner for Smith dead time compensator. *International Journal of Control* 60, 117–135.
- Postel, J., 1981. Transmission Control Protocol. RFC 793, Standard .
- Quet, P., Ataşlar, B., Iftar, A., Özbay, H., Kalyanaraman, S., Kang, T., 2002. Rate-based flow controllers for communication networks in the presence of uncertain time-varying multiple time-delays. *Automatica* 38, 917–928.
- Shustin, E., Fridman, E., 2007. On delay-derivative-dependent stability of systems with fast-varying delays. *Automatica* 43, 1649 – 1655.
- Smith, O., 1959. A controller to overcome dead time. *ISA Journal* 6, 28–33.
- Veal, B., Li, K., Lowenthal, D., 2005. New Methods for Passive Estimation of TCP Round-Trip Times, in: *Proc. of Passive And Active Network Measurement (PAM) '05*, pp. 121–134.
- Yamanaka, K., Shimemura, E., 1987. Effects of mismatched Smith controller on stability in systems with time-delay. *Automatica* 23, 787–791.
- Zhong, Q.C., 2006. *Robust control of time-delay systems*. Springer Verlag: London.