Impact of queueing delay estimation error on equilibrium and its stability

Corentin Briat†, Emre A. Yavuz, and Gunnar Karlsson

ACCESS Linnaeus Center, KTH, SE-100 44 Stockholm, SWEDEN
http://www.access.kth.se/
{cbriat, emreya, gk}@kth.se
† corresponding author

Abstract. Delay-based transmission control protocols need to separate round-trip time (RTT) measurements into their constituting parts: the propagation and the queueing delays. We consider two means for this; the first is to take the propagation delay as the minimum observed RTT value, and the second is to measure the queueing delay at the routers and feed it back to the sources. We choose FAST-TCP as a representative delay-based transmission control protocol for analysis and study the impact of delay knowledge errors on its performance. We have shown that while the first method destroys fairness and the uniqueness of the equilibrium, the stability of the protocol can easily be obtained through tuning the protocol terms appropriately. Even though the second technique is shown to preserve fairness and uniqueness of the equilibrium point, we have presented that unavoidable oscillations can occur around the equilibrium point.

Keywords: Congestion Control; FAST-TCP; Time-Delay Systems; fairness; stability

1 Introduction

Most recent developments for the internet have concerned the development of delay-based congestion control and its instantiation in the form of FAST-TCP [22]. It is possible to get better performance in terms of shorter queues and lower losses, both resulting in lower end-to-end delay for a transfer, when control is based on a continuous state variable rather than the binary signal of a packet drop. We have in a previous work shown that knowledge of both queuing and propagation delays is necessary and sufficient to obtain stability, fairness, and efficiency in distributed congestion control, such as TCP [18]. In this paper, we study how these measures are obtained and the impact any imperfection could have on the control performance.

Motivation The motivation to study congestion control relates to its importance for network operation. Mission critical information systems, for instance supervisory control and data acquisition systems used for controlling the power grid, are increasingly relying on internet communication for both system status data and control commands. The trend to port such systems from proprietary low speed data networks to the internet is based on ready availability of higher capacity and low delay that improve system operation (such as smart grid); the cost to provide similar connection quality by proprietary networks might not be justifiable for competitive reasons. A network *without* congestion control may still be efficient given perfect error handling in the form of selective repeat ARQ or erasure coding [1] as shown in a recent work by Bonald et al. The proviso is that one of two possible conditions must be met: The first is that the capacities of access links are small with respect to the shared links, in which case the efficiency is high for drop-tail FIFO queues. The second is by using a fair-drop policy in routers for which the sharing will also be fair. Noting the increases in access links owing to commercial offering of 100 Mb/s and higher-rate DSL connections, and of fiber to homes that are increasingly being installed, we conclude that the fair-drop mechanism appears necessary before daring to dismantle congestion control. An additional disadvantage with unregulated elastic flows is that the sharing with inelastic flows is obliterated; while it is possible to share resources between the two types by qualitative differentiation of traffic controls (rate control for elastic flow and admission control of inelastic flows) [13]. Hence, we believe congestion control remains a vital function for sound internet operation for a foreseeable future.

Contributions This paper reports new analytical results on the performance of delay-based congestion control. We present how the estimation of queueing delay from samples of round-trip delay may cause unfairness among competing flows and instability if the protocol tuning parameters are not chosen accordingly. The alternative is to measure the queueing delay in routers and feed back those measures to all senders; hence ensuring fairness since all act on the same information. In such a case, the error at equilibrium will be the same for all users leading to a unique and fair equilibrium point. However, due to the quantization of the measurements, we have proved that the network will inevitably exhibit an oscillating behavior. Our analysis is based on time-delay systems theory and it involves large polynomials and complex transcendental equation analysis, which is typically unsolvable analytically except in very simple cases. Thus, we have considered only the single-user/single-buffer case to obtain the results for the stability analysis. It is however important to stress that our results may be scaled to a more general context, at least qualitatively, since observed issues that occur for the single-user/single-buffer topology, are expected to appear also for cases with complex topologies.

Outline of paper The paper has the following outline. Section 2 covers the related work and in section 3 we give a brief overview of FAST-TCP. Section 4 presents the model for the delay based congestion avoidance protocols. In section 5 we analyze the impact of the queuing delay estimation error on equilibrium and its stability. Section 6 provides a similar analysis, yet considers the quantization error when measuring the queuing delay. We conclude the work in section 7.

2 Related Work

The congestion control algorithm implemented in TCP Reno [8] has performed well and gone through several enhancements since then. It is however well-known that as bandwidth-delay product continues to grow, it will become a performance bottleneck itself. The poor performance of TCP Reno in such networks is due to slowness of linear increase by one packet per round-trip time, severeness of multiplicative decrease per loss event, the difficulty in maintaining large average congestion windows, which requires an extremely small equilibrium loss probability and using a binary congestion signal based on packet loss which causes oscillations.

Delay-based congestion control has been proposed in [10, 21, 3] to overcome these difficulties. Control protocols established on delay-based congestion avoidance (DCA) algorithms are shown to achieve better performance than protocols established on congestion avoidance algorithms based on packet loss as network congestion indicator, with respect to network efficiency, stability and latency [22, 10, 2]. Queuing delay can be estimated more accurately compared to loss since packet losses are rare events, and information on queuing delay has finer resolution than what loss samples provide. In [18], several objectives (stability, fairness, efficiency) are considered and necessary conditions for a delay-based protocol to achieve these are provided. It is shown that knowledge of both the (aggregate) queuing delay and the constant propagation delay are necessary to ensure both fairness and efficiency.

DCA protocols react to shifts in measured RTT with the assumption that it is either an indication of congestion or of excess capacity in the bottleneck link. Since the protocols try to sustain a buffer occupancy at approximately a constant level, a new flow could easily overestimate the propagation delay by mistakenly including a constant queueing delay in the measured minimum RTT when it starts. This overestimation eventually causes unfairness, known as the *persistent congestion problem*, among contending flows as the sender assumes the link being less congested than what it actually is. Many methods have been proposed to correct this problem [11, 12, 5, 20] yet almost all of them rely on intermediate routers to be upgraded and thus redeployed. A few other attempts to solve the problem without support from the queuing management mechanisms in routers, have been presented in [6, 17] under certain limitations such as having large enough and homogeneous propagation delays, not too many simultaneous flows, and no uncontrolled cross traffic.

3 A DCA protocol: FAST-TCP

FAST-TCP is the most representative DCA protocol that achieves higher throughput, lower latency and fewer packet drops compared to previous versions of TCP. Loss based protocols such as TCP Reno and its variants drive the network to congestion so that they can receive the feedback they need in order to adjust the size of their congestion windows. In contrast, TCP protocols with delay based congestion avoidance mechanisms, such as TCP Vegas and FAST-TCP, keep an estimate of the round trip propagation delay, which is the minimum RTT observed throughout a single connection, to track changes in the queuing delay by measuring the RTT continuously. The queue length is estimated by measuring the difference between the observed RTT and the estimated propagation delay. FAST-TCP is an enhanced version of TCP Vegas that does not penalize the flows with large bandwidth-delay products and that has good convergence properties. FAST-TCP also differs from TCP Vegas the way the rate is adjusted when the number of packets stored is too small or large. TCP Vegas makes fixed size adjustments, independently of how far it is from the equilibrium point. FAST-TCP takes larger steps when the system is far away from the equilibrium and smaller steps otherwise to improve the speed of convergence and stability.

FAST-TCP implements a learning procedure consisting of estimating the propagation by the minimum observed RTT (plus some filters) and this has been modelled accurately in [9]. One additional feature of FAST-TCP is independence of the propagation delays in equilibrium; hence sources with large propagation delays are not penalized as is usually the case for classical implementations of TCP. The congestion window update law explicitly depends on both the queuing and the propagation delays [22] in order to achieve fairness even in the case of heterogeneous propagation delays.

4 Analysis of The Ideal Case

In this section some preliminary results from [4] are recalled. We assume that both propagation and queuing delays are perfectly known (ideal case).

We consider here the following continuous-time fluid-flow nonlinear model, very often used to analyze the behavior of a single-bottleneck network (with unlimited queue length) with N users using FAST-TCP [9, 4]:

$$\dot{\tau}(t) = \frac{1}{c}\sum_{i=1}^{N} \frac{x_i(t-T_i^f)}{\tau(t)+T_i} + \delta(t) - 1 \tag{1}$$

$$\dot{x}_i(t) = \gamma\left(-\frac{\tau(g(t-T_i^b))}{T_i+\tau(g(t-T_i^b))}x_i(t) + \alpha_i\right),\ i = 1, \ldots, N \tag{2}$$

where x_i, τ, c, $T_i = T_i^f + T_i^b$, T_i^f and T_i^b are the window size of user i, the queuing delay, the buffer output capacity, the total constant propagation delay of source i and the forward and backward constant propagation delays of source i respectively. The disturbance term δ represents the normalized cross-traffic modeling unregulated flows and non-FAST-TCP flows such as the usual TCP. The terms γ and α_i are protocol tuning parameters, the first one acts on the speed of reaction of the protocol (the bandwidth, in the control theoretic sense) and α_i is the desired number of enqueued packets at equilibrium. The function $g(\cdot)$ is the inverse function of $f(t) = t + \tau(t)$ as defined in [4] and allows to write the overall network rigorously as a nonlinear time-delay system with state-dependent delay. The unique equilibrium point of this a model is given by:

$$\tau^* = \frac{\sum_i \alpha_i}{c(1-\delta^*)},\ x_i^* = \alpha_i\left(1 + \frac{T_i}{\tau^*}\right) \text{ and } \phi_i^* = = \frac{\alpha_i}{\tau^*} \tag{3}$$

where ϕ_i^* denotes the flow of source i at equilibrium, $i = 1, \ldots, N$. From the above result we can see that the equilibrium is both proportionally-fair and efficient [19, 18]. The equilibrium flows do not depend on the propagation delays, hence sources with large propagation delays are not penalized.

Since the theoretical local/global stability analysis of the nonlinear model (1)-(2) is a very difficult still open problem, we will restrict the analysis to the single source problem (i.e. $N = 1$). Using results from time-delay systems theory [7, 16, 15, 4] and robust analysis [23, 7], the following theorem is proved in [4]:

Theorem 1 (Network Local Stability [4]). *Let us consider the network model (1)-(2) with single user. Then the equilibrium point (x^*, τ^*) is*

1. *locally delay-independent stable if and only if $\tau^* > T$.*
2. *locally delay-dependent stable if $\tau^* < T$ and $\tau^*(T-1) + T^2 \leq 0$.*
3. *locally delay-dependent stable if $\tau^* < T$, $\tau^*(T-1) + T^2 > 0$ and $\gamma < \frac{1}{\tau^*(T-1)+T^2}$.*

□

In the subsequent sections, the impact of the unperfect knowledge of the queuing delay value on the overall network stability will be studied.

5 Learning the queuing delay

In the FAST-TCP protocol implementation, the propagation delays are estimated as the minimal observed RTTs: $\hat{T}_i(t) := \inf_{s\in[t_i,t]}\{RTT_i(s)\}$ where the

RTT is $RTT_i(t) := T_i + \tau(g(t - T_i^b))$, $\hat{T}_i(t)$ is the estimated propagation delay for source i at time t and t_i is the arrival time of source i in the network. So, the sources learn their propagation delays through an iterative process using the RTT measurements. The actual advantage of such a procedure lies in its simplicity: only the senders' routines need to be modified, the entire infrastructure (routers/servers) remains unchanged. However, according to the scenario, the sources may be unable to estimate their propagation delay accurately, resulting in an underestimation of the queueing delays and a loss of fairness. The only way to observe the actual propagation delay is to communicate when the queueing delay is 0. However, DCA protocols have an antagonist effect since they strive to maintain a non-zero queue for efficiency and flow fairness.

In order to take into account the learning process in the analysis, the protocol model (2) is refined to

$$\dot{x}_i(t) = \gamma\left(-x_i(t) + \frac{T_i+\varepsilon_i(t)}{\tau(g(t-T_i^b))+T_i}x_i(t) + \alpha_i\right) \tag{4}$$

where the learning errors $\varepsilon_i(t)$ are defined by $\varepsilon_i(t) = \hat{T}_i(t) - T_i \leq \tau(g(s - T_i^b))$. Since $\varepsilon_i(t)$ can be anything according to the scenario, we will rather focus on asymptotic properties (equilibrium points and local stability) of the above protocol model for any possible values for the errors at equilibrium.

5.1 Impact on the equilibrium point

Solving for the equilibrium points for (1)-(4) yields the expressions:

$$x_i^* = \frac{\alpha_i(T_i+\tau^*)}{\tau^*-\varepsilon_i^*},\ \ \phi_i^* = \frac{\alpha_i}{\tau^*-\varepsilon_i^*} \text{ and } \varepsilon_i^* < \tau^* \tag{5}$$

where the ε_i^*'s are the estimation errors at equilibrium and the delay at equilibrium τ^* solves

$$\sum_{i=1}^N \phi_i^* - c(1-\delta^*) = 0. \tag{6}$$

From the above equations, we can clearly see that when the errors at equilibrium for each source are different, then proportional-fairness cannot be achieved. Hence, fairness is only reachable under the very strong assumption of exact knowledge of all the propagation delays. The equation (6) defining the equilibrium delay can be rewritten as the polynomial equation $P_N(\tau) = 0$ where

$$P_N(\tau) := \sum_{i=1}^N \left[\alpha_i \prod_{j\neq i}^N(\tau^* - \varepsilon_j^*)\right] - c(1-\delta^*)\prod_{i=1}^N(\tau^* - \varepsilon_i^*) \tag{7}$$

and for which only the nonnegative solutions must be considered.It is well known that, in general, there is no analytical solutions to $P_N(\tau) = 0$ when N is large. Hence we will restrict the analysis to simple cases where analytical results can be obtained.

Equilibrium points analysis - Single source case In the single source case, the equilibrium queuing delay is given by $\tau^* = \frac{\alpha}{c(1-\delta^*)} + \varepsilon^*$ which is nothing but a simple shift of the ideal equilibrium (3) leading to an increase of the queuing delay. When $\varepsilon^* \to 0$, we recover the equilibrium point of the ideal case.

Equilibrium points analysis - Two sources case In the 2-sources case, the equilibrium delay is defined by the following polynomial $P_2(\tau) = -\eta_2\tau^2 + \eta_1\tau - \eta_0$ with $\eta_2 = c(1-\delta^*)$, $\eta_1 = \alpha_1 + \alpha_2 + \eta_2(\varepsilon_1^* + \varepsilon_2^*)$ and $\eta_0 = \eta_2\varepsilon_1^*\varepsilon_2^* + \alpha_1\varepsilon_2^* + \alpha_2\varepsilon_1^*$. Since by definition, we have $c, \delta^*, \alpha_i, \varepsilon_i^* > 0$, $i = 1, 2$ then $\eta_2 > 0, \eta_1 > 0$ and $\eta_0 > 0$. Hence, the real part of the roots of P_2 are positive. To see that the solutions are all real, it is easy to show that the discriminant $\Delta := \eta_1^2 - 4\eta_2\eta_0$ is positive. Thus, the network admits 2 positive equilibrium points given by

$$\tau^* = \frac{\tau_n}{2} + \left(\frac{\varepsilon_1+\varepsilon_2}{2} \pm \frac{\sqrt{\Delta}}{2c(1-\delta^*)}\right) \tag{8}$$

where $\tau_n = \frac{\alpha_1+\alpha_2}{c(1-\delta^*)}$ is the equilibrium delay in the ideal case. So, in the two sources problem, an overestimation of the propagation delays leads to the creation of two distinct positive queueing delay equilibrium points. Note also that the solutions $\tau^* \to \tau_n$ as $\varepsilon_1, \varepsilon_2 \to 0$ (the zero limit disappears due to a pole/zero cancellation).

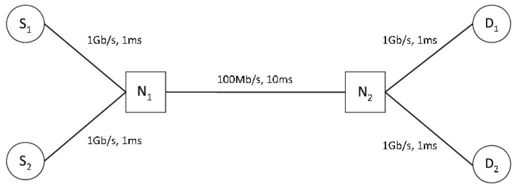

Fig. 1: Network Topology used for Simulation

Example 1. Let us consider the topology depicted in Fig. 1 where we establish a FAST TCP connection between each source S_i, and its corresponding destination D_i. We assume that the packets carry a payload of 1000 bytes and that the maximal queue size is chosen to be sufficiently large to avoid packet dropping. In the considered scenario, we assume that the two flows arrive consecutively. The first one comes at time $t = 10$ seconds, when the queue is empty. Hence the first source can estimate exactly its propagation delay, so $\varepsilon_1^* = 0$, and the queuing delay converge to the equilibrium value $\tau_1^* = \alpha/\eta$ where $\eta = c(1-\delta^*)$. Then the second flow comes at $t = 30$ seconds and makes the queue length increase, hence the minimal measured RTT for the second source is the one measured at $t = 30$ seconds, hence we have $\varepsilon_2^* = \tau_1^*$. Solving for the delay equilibrium points when both flows are active we get $\tau_2^* = \frac{\alpha}{2\eta}\left(3 \pm \sqrt{5}\right)$ which are both positive. Solving now for the flows, we get $\phi_1^* = \frac{2\eta}{3+\sqrt{5}}$ and $\phi_2^* = \frac{2\eta}{1+\sqrt{5}}$ whose sum equals η, showing then efficient but unfair $(\phi_1^* < \phi_2^*)$ equilibrium. The simulation of the topology in Figure 1 is performed with the NS-2 simulator and the obtained results are gathered in Table 1 with a comparison with the theoretical results of Section 5.1. Two scenarios are considered, the first one considers homogeneous propagation delays and no cross-traffic while the second adds a cross-traffic of 20Mb/s.

We can see that, for the considered scenarios, the windows size, the queuing delay and the queue length are quite well predicted. Fig. 2 shows the rate evolution of the sources in the scenario without cross-traffic. As also noticed in [17], the flows fail to converge to a fair equilibrium.

Experiment 1				Experiment 2			
Parameters	variables	Theory	NS-2	Parameters	variables	Theory	NS-2
$\alpha = 200$	ϕ_1^*	5006	4614	$\alpha = 200$	ϕ_1^*	4005	3701
$c = 100\text{Mb/s}$	ϕ_2^*	8100	7413	$c = 100\text{Mb/s}$	ϕ_2^*	6481	5923
$\delta^* = 0$	τ_2^*	40ms	43ms	$\delta^* = 0.2$	τ^*	50ms	54ms
$T_1 = 24\text{ms}$	q_2^*	524	520	$T_1 = 24\text{ms}$	q^*	655	653
$T_2 = 24\text{ms}$	x_1^*	320	310	$T_2 = 24\text{ms}$	x_1^*	296	289
$N_p = 8\text{kb}$	x_2^*	518	499	$N_p = 8\text{kb}$	x_2^*	479	462

Table 1: Comparison of theoretical and simulation results (N_p is packet size)

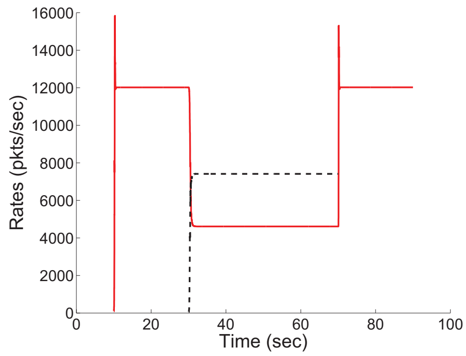

Fig. 2: Rates evolution for Source 1 (plain) and Source 2 (dashed) (no cross-traffic)

5.2 Impact on the stability of the equilibrium point

It is interesting to study the local stability of the equilibrium points to answer the question on the protocol stability at an unfair equilibrium point. To this aim, the following linearized model from (1)-(4) is devised:

$$\begin{aligned} \dot{y}_i(t) &= -\gamma \frac{\varepsilon_i^* - \tau^*}{T_i + \tau^*} y_i(t) - \gamma \frac{(T_i + \varepsilon_i^*) x_i^*}{(T_i + \tau^*)^2} \nu(t - T_i - \tau^*) \\ \dot{\nu}(t) &= \sum_i \frac{1}{c(T_i + \tau^*)} y_i(t) - \sum_i \frac{x_i^*}{c(T_i + \tau^*)^2} \nu(t) + \frac{1}{c} \zeta(t) \end{aligned} \quad (9)$$

where $y_i(t) := x_i(t) - x_i^*$, $\nu(t) := \tau(t) - \tau^*$ and $\zeta(t) := \delta(t) - \delta^*$. Restricting us to the single user case, we get the following theorem for delay-independent stability:

Theorem 2. *The system (9) is locally delay-independent stable if and only if the inequality* $\varepsilon^* < \frac{\alpha}{c(1-\delta^*)} - T$ *holds.*

Proof. The proof is similar to as the one of [4, Theorem 4.3].

Since $\varepsilon^* \geq 0$ then in order to tolerate positive errors, the right-hand side must be at least positive. Based on this, we can conclude that the error has a negative impact on the stability. By extension, the same problem will occur in more complex topologies.

Whenever delay-independent stability is not achieved we have the following theorem:

Theorem 3. *The system (9) is delay-dependent stable if one of the following conditions hold:*

1. $T + \varepsilon^* - \frac{\alpha}{c(1-\delta^*)^2} < 0$*; or*
2. $T + \varepsilon^* - \frac{\alpha}{c(1-\delta^*)^2} \geq 0$ *and* $\gamma < \frac{c(1-\delta^*)^2}{c(T+\varepsilon^*)(1-\delta^*)^2 - \alpha}$.

Proof. The proof is identical as for Theorem 1; see [4].

According to the above results, we can conclude that the stability can be ensured through an appropriate choice of the tuning parameter γ similarly as in the ideal case. Additionally, the error term penalizes the maximal admissible speed of the protocol and has thus an impact on the overall efficiency of the network. Conversely, if the term γ is not chosen in order to consider the eventual error term, stability may be lost.

6 Measuring the queuing delay - Case study

The second solution to the estimation of both queuing and propagation delays consists of an explicit feedback of the queuing delays by the routers. In such a framework, part of the header of the packet is dedicated to contain the sum of all the queueing delays the packet has experienced over its path. Thus, the source gets an explicit value for the total queuing delay and can subtract it from the measured RTT to compute the propagation delay. This solution needs, however, an update of all routers to add this feature. It is hence less simple to deploy than the estimation procedure for the queuing delay that is solely based on RTT measurements. Moreover, since the size of the packet header is constant, the measured aggregate queuing delay is stored with fixed precision and a measurement error is consequently introduced. Therefore, an analysis of the influence of the errors on the equilibrium points and on the stability deserve to be studied. Interestingly, the protocol model incorporating the use of a quantized measure of the queuing delay is very similar to as the one incorporating the learning error:

$$\dot{x}_i(t) = \gamma \left(\frac{-\tau(g(t-T_i^b)) - \varepsilon_i(g(t-T_i^b))}{\tau(g(t-T_i^b)) + T_i} x_i(t) + \alpha_i(t) \right). \tag{10}$$

However, in the present case, the error term satisfies $|\varepsilon_i(\cdot)| \leq q/2$ where q is the resolution of the quantizer. The only differences lie in the sign of the errors which are not restricted to be positive and in the boundedness of the errors. Indeed, the worst case error only depends on the choice of the quantization step while in the previous case, the worst case error was depending on the state of the network.

6.1 Impact on the equilibrium point

At equilibrium, the queuing delay is constant and hence the quantization error is identical for all users, i.e. $\varepsilon_i = \varepsilon^*$, $i = 1, \dots, N$. This an important property of the current approach. Indeed, simple computations yield:

$$x_i^* = \alpha_i \frac{T_i+\tau^*}{\tau^*-\varepsilon^*},\ \ \phi_i^* = \frac{\alpha_i}{\tau^*-\varepsilon^*} \text{ and } \tau^* = \frac{\sum_i \alpha_i}{c(1-\delta^*)} + \varepsilon^*. \tag{11}$$

In such a case, the equilibrium point is unique, efficient and proportionally fair. This was not the case with the learning strategy due to the imbalance between the errors. This is one of the benefits of the approach.

6.2 Impact on the stability

Similarly to as previously, we will consider the single user problem. We will first assume that $\varepsilon^* = 0$ to avoid complex calculations. A discussion will be provided for the case $\varepsilon^* \neq 0$. When the quantization error is 0 at equilibrium, the quantization function $\varphi(\cdot)$ is an odd function which belongs to the sector $(0, 2)$, i.e. we have $0 \leq \frac{\varphi(s)}{s} \leq 2$ for any $s \in \mathbb{R}$. A lot of works have been devoted to the analysis of linear systems interconnected to sector nonlinearities [14]. The problem can be rewritten as the negative feedback interconnection of $\varphi(\cdot)$ and $F(s) = \frac{\mu^3\xi_2\psi_1 e^{-s(\tau^*+T)}}{(s+\mu\xi_1)(s+\mu\xi_3)}$. A very important result for the stability analysis of such interconnections is called the circle criterion and consists of a generalization of the Nyquist criterion:

Theorem 4 (Circle Criterion). *Let us consider an interconnection (with negative feedback) of an asymptotically stable system $F(s)$ and a nonlinear element $\varphi(\cdot)$ satisfying the sector condition (k_1, k_2). The interconnection is asymptotically stable if the graph of $F(j\omega)$ with $\omega \in \mathbb{R}$ does not enter the circle passing through the points $-1/k_1$ and $-1/k_2$ in the complex plane.*

In the considered case, the 'circle' coincides with the vertical plane passing through the point $-1/2$ in the complex plane. It is easy to show that the circle condition is equivalent to show that $F(j\omega)$ does not encircle the point $-1/2$ in the complex plane. This actually consists in a scaling of the Nyquist criterion and is equivalent to the stability of $\hat{F}_2(s) = \frac{2N(s)}{D(s)+2N(s)}$ where $F(s) = N(s)/D(s)$. Since the structure of the quasipolynomial $D(s) + 2N(s)$ is very similar to as in the ideal case, the same approach is used and yields the lemma:

Theorem 5. *Assuming that the quantization error at equilibrium is 0, then the equilibrium point (x^*, τ^*) of system (1)-(10) is*

1. *locally delay-independent stable if $\tau^* > 2T$.*
2. *locally delay-dependent stable if $\tau^* < 2T$ and $2T(T + \tau^*) - \tau^* \leq 0$.*
3. *locally delay-dependent stable if $\tau^* < 2T$ and $2T(T + \tau^*) - \tau^* > 0$ and $\gamma < \frac{1}{2T(T+\tau^*)-\tau^*}$.*

Proof. The proof follows the same lines as for the other results; see [4].

We can conclude on the fact that, in the case $\varepsilon^* = 0$, the tuning term γ can be chosen in order to avoid oscillations. We explain now why this is not possible in the general case $\varepsilon^* \neq 0$.

Indeed, when $\varepsilon^* \neq 0$, the quantization function is not odd anymore since it must be horizontally shifted to be centered around the equilibrium point. Therefore the sector takes the more general form $(0, \theta)$ where θ is defined by $\theta := 2(1-2|\kappa|)^{-1}$ where $\varepsilon^* = \kappa q$, $\kappa \in [-1/2, 1/2]$. Hence, the term θ can reach arbitrarily large values, translating then the 'circle' horizontally to the right. When the error is maximal ($\kappa = \pm 1/2$), the forbidden area of the complex plane is the entire open left-half plane itself. Due to the exponential term of $F(s)$, then the graph of $F(j\omega)$, $\omega > 0$ always enters the open left-half plane and thus limit cycles cannot be avoided.

7 Conclusion

In this paper, we have addressed the fairness and stability properties of delay based transmission control protocols; in particular FAST-TCP. The update law of FAST-TCP congestion windows requires to know both the propagation and queuing delays if fairness and efficiency are to be provided. We have incorporated two approaches to estimate the propagation delay using the model we developed for such protocols.

The first approach, which is also the one implemented in FAST-TCP, employs an iterative learning process to estimate the propagation delay as the minimal observed RTT. The propagation delay is always overestimated in this case unless the queuing delay drops to zero. It is shown that such estimation procedure leads to loss of fairness due to diverse estimation errors at each source along with multiple equilibrium points for the queuing delay. We have developed a model for the delay based congestion avoidance protocols to analyze the impact of the queuing delay estimation error on equilibrium and its stability. Using this model, we have observed that the stability of the equilibrium points, in the single-source case, can be ensured through an appropriate choice of the tuning term γ of the protocol. The developed model, which is able to predict the congestion window size, the queuing delay and the number of enqueued packets quite accurately, is validated by running NS-2 simulations.

The second approach we have analyzed is based on the assumption that the routers feedback a quantized measure of their queuing delay. Hence both propagation and queuing delays (modulo the quantization error) can be estimated easily. We have shown that this strategy manages to preserve the uniqueness of the equilibrium point as well as fairness, yet it can not prevent limit cycles (oscillations) around the equilibrium points for which the quantization error is too large.

References

1. Bonald, T., Feuillet, M., Proutiere, A.: Is the "law of the jungle" sustainable for the internet? In: Proc. IEEE INFOCOM (2009)
2. Brakmo, L.S., O'Malley, S.W., Peterson, L.L.: TCP Vegas: new techniques for congestion detection and avoidance. SIGCOMM Comput. Commun. Rev. 24(4), 24–35 (1994)

3. Brakmo, L.S., Peterson, L.L.: TCP Vegas: End to end congestion avoidance on a global internet. IEEE Journal on selected Areas in communications 13, 1465–1480 (1995)
4. Briat, C., Hjalmarsson, H., Johansson, K., Karlsson, G., Jönsson, U., Sandberg, H.: Nonlinear state-dependent delay modeling and stability analysis of internet congestion control. In: 49th IEEE Conference on Decision and Control, Altlanta, USA (2010)
5. Chan, Y.C., Chan, C.T., Chen, Y.C., Ho, C.Y.: Performance improvement of congestion avoidance mechanism for TCP Vegas. Parallel and Distributed Systems, International Conference on 0, 605 (2004)
6. Cui, T., Andrew, L., Zukerman, M., Tan, L.: Improving the fairness of FAST TCP to new flows. Communications Letters, IEEE 10(5), 414 – 416 (may 2006)
7. Gu, K., Kharitonov, V., Chen, J.: Stability of Time-Delay Systems. Birkhäuser (2003)
8. Jacobson, V.: Congestion avoidance and control. In: SIGCOMM '88: Symposium proceedings on Communications architectures and protocols. pp. 314–329. ACM, New York, NY, USA (1988)
9. Jacobsson, K.: Dynamic Modeling of Internet Congestion Control. Ph.D. thesis, KTH School of Electrical Engineering (2008)
10. Jain, R.: A delay-based approach for congestion avoidance in interconnected heterogeneous computer networks. SIGCOMM Comput. Commun. Rev. 19(5), 56–71 (1989)
11. La, R.J., Walrand, J., Anantharam, V.: Issues in TCP Vegas (2001)
12. Low, S.H., Peterson, L., Wang, L.: Understanding TCP vegas: a duality model. SIGMETRICS Perform. Eval. Rev. 29(1), 226–235 (2001)
13. Lundqvist, H., Mas, I., Karlsson, G.: Edge-based differentiated services. In: In Proc. of International Workshop on Quality of Service (IWQoS) (Jun 2005)
14. Lur'e, A., Postnikov, V.: On the theory of stability of control systems. Prikl. Mat. i Mekh (Applied mathematics and mechanics) 8(3), 3–13 (1944)
15. Michiels, W., Niculescu, S.: Stability and stabilization of time-delay systems. An eigenvalue based approach. SIAM Publication, Philadelphia, USA (2007)
16. Niculescu, S.I.: Delay effects on stability. A robust control approach, vol. 269. Springer-Verlag: Heidelbeg (2001)
17. Rodríguez-Prez, M., Herrería-Alonso, S., Fernández-Veiga, M., Lopez-Garcia, C.: The persistent congestion problem of FAST-TCP: analysis and solutions. European Transactions on Telecommunications 21(6), 504–518 (2010)
18. Sandberg, H., Hjalmarsson, H., Jönsson, U., Karlsson, G., Johansson, K.: On performance limitations of congestion control. In: 48th IEEE Conference on Decision and Control. pp. 5869–5876. Shanghai, China (2009)
19. Srikant, R.: The Mathematics of Internet Congestion Control. Birkhäuser, Boston, USA (2004)
20. Tan, L., Yuan, C., Zukerman, M.: FAST TCP: fairness and queuing issues. Communications Letters, IEEE 9(8), 762 – 764 (aug 2005)
21. Wang, Z., Crowcroft, J.: Eliminating periodic packet losses in the 4.3-tahoe BSD TCP congestion control algorithm. SIGCOMM Comput. Commun. Rev. 22(2), 9–16 (1992)
22. Wei, D.X., Jin, C., Low, S.H., Hegde, S.: FAST TCP: motivation, architecture, algorithms, performance. IEEE/ACM Trans. Netw. 14(6), 1246–1259 (2006)
23. Zhou, K., Doyle, J., Glover, K.: Robust and Optimal Control. Prentice Hall, Upper Saddle River, New Jersey, USA (1996)