

On Performance Limitations of Congestion Control

H. Sandberg, H. Hjalmarsson, U. T. Jönsson[†], G. Karlsson, and K. H. Johansson

Abstract—Fundamental performance limitations on congestion control is discussed in relation to the information that is available in the controller. Three control architectures that all use bottle-neck buffer delay information of various form as inputs and sending rate as outputs are considered. It is shown that feedback delays from buffer to senders set limits on the achievable performance measured through fairness, efficiency and stability. We apply our findings to TCP FAST and can make some new interesting control theoretic interpretations of that protocol.

I. INTRODUCTION

Congestion control is a vital function for the operation of packet-switched networks. Most traffic in the Internet is controlled by the transmission control protocol (TCP), which uses congestion feedback to regulate sending hosts. The purpose is to reduce the data traffic injected into the network when there is indication of congestion and to increase the traffic when capacity is available. Despite a large volume of work on the analysis of TCP, there remain open issues. A major part of the research is on various versions of TCP and how they can be improved. A long-standing problem is how to achieve a fair, stable and robust congestion control mechanism; this problem still achieves substantial attention [3]. The fundamentals of congestion control have been studied in [5], [12], [7], discussing issues such as when a decentralized and stabilizing control law can be achieved. We are interested in similar issues in this paper. Our approach is different, however, because we take queue dynamics, propagation delays and cross-traffic into account.

Fairness is a central performance criterion for distributed congestion control and concerns the sharing of resources in the network [10]. Present TCP versions, such as New Reno, yield a capacity share for each session that is inversely proportional to its round-trip time (RTT), and inversely proportional to the square root of the loss probability. This sharing results from the control mechanism and cannot be tuned to any desired operational policy (for a reflection of this that is relevant here, see [9]). In this paper, we impose a sharing policy and study the stability and efficiency of the system when the sessions have different round-trip delays. We limit the study to uniform sharing of the bottleneck capacity over the backlogged sessions (i.e., a share of c/n for a capacity c and n sessions).

We consider controllers with various available feedback information. Two external variables influence the dynamics

of the closed-loop system: exogenous uncontrolled cross traffic, y_c (e.g., UDP flows in the Internet) and the number of ongoing congestion-controlled sessions, n . The task of the congestion controllers is to track variations in these two variables under a stated sharing policy. We provide explicit information to the controllers in the form of the RTT, the queuing delay at a bottleneck link, τ , and the tuple (RTT, τ) and we study the achievable performance for the resulting three control architectures. The explicit information is supposed to be ideal (continuous in time and amplitude) but delayed. Each controller (one per session) continuously adjust its fluid sending rate over the interval $[0, c]$. The system is purposely idealized to get tractable models that allow us to make observations on the achievable performance. We remark, however, that the performance will be exacerbated when complexities are added back to the model in order to mimic real systems (sampled and quantized information of the queuing delay for the explicit congestion notification; estimation time and errors for implicit congestion signals; loss of data segments in the forward path and acknowledgments in the return path; other stochastic delay variations than those at the bottleneck link; multiple bottlenecks; non-smooth rate control).

The observations we report pertain to limits on the achievable performance for distributed congestion control with explicit feedback. We are interested in a triad of properties: stability, fairness and efficiency. A short justification might be warranted: by considering efficiency, we exclude the trivial solution for stability and fairness of zero rates to all; by considering fairness, stability is not certain. Our contribution is to establish what explicit information brings to the control problem and what properties the controllers must have in order to be fair, stable and efficient. We exemplify with an analysis of TCP FAST [15], [14].

The paper is structured in seven sections. Section II describes the model of the system and control architectures. Section III holds a feasibility analysis of fairness and efficiency for the architectures, while Section IV focuses on the control limitations. Section V presents an analysis of the attainable bandwidth of the closed-loop control system under propagation delay. Section VI applies the model to TCP FAST in order to exemplify the observations made throughout the paper.

II. NETWORK MODEL AND OBJECTIVES

A general traffic control system can be viewed as a distributed control problem where several user dynamics communicate through a set of link dynamics, see Figure 1. Each user determines a flow rate x based on an aggregate

The authors are with the ACCESS Linnaeus Centre at the Royal Institute of Technology (KTH), Stockholm, Sweden.

The work is supported by the Swedish Research Council.

[†] Corresponding author.




Fig. 1. A model for Internet congestion control (adopted from [8]).

congestion measure q and the link dynamics determine their congestion measures based on the aggregate flow rates y . The aggregation of congestion information and rates are modeled using the forward and backward routing matrices $R_f(s)$ and $R_b(s)$, which also includes forward and backward propagation delays in each path from a source to a link and vice-versa. These matrices have identical structure but the various delays may be different. The cross traffic y_c can be viewed as a disturbance which contains traffic flows that do not obey any congestion control protocol.

The literature suggests a large number of congestion control schemes based on different control mechanisms and congestion indicators, see e.g. [13], [8]. In this paper we consider a rate-based control scheme where all users use the same control protocol. We restrict attention to the case with a single bottleneck link in order to clarify how the achievable performance of such control schemes depends on the number of users, the available delay information, and the cross traffic.

Figure 2 is a block diagram description of the flow model (single-link-multiple-users) considered in the paper, where the following notation is used¹

- d_k , $k = 1, \dots, n$ are constant transport delays,
- τ is the buffer delay (time-varying delay),
- τ_k , $k = 1, \dots, n$ are delayed buffer delays,
- x_k is the flow-rate of user k ,
- y is the aggregate flow-rate,
- y_c is the cross traffic,
- n denotes the number of users.

The link is assumed to operate in a first in, first out fashion that can be modeled as a composition of a saturated integrator, $\left[\frac{1}{cs} \right]_{\tau_{\min}}^{\tau_{\max}}$, defined by

$$\dot{\tau}(t) = \begin{cases} \frac{1}{c} \min(y(t) + y_c(t) - c, 0), & \tau(t) = \tau_{\max}, \\ \frac{1}{c} (y(t) + y_c(t) - c), & \tau \in (\tau_{\min}, \tau_{\max}), \\ \frac{1}{c} \max(y(t) + y_c(t) - c, 0), & \tau(t) = \tau_{\min}, \end{cases}$$

¹It is possible to include delays in the forward path from the users to the links and this would not affect any of the observations in the paper.

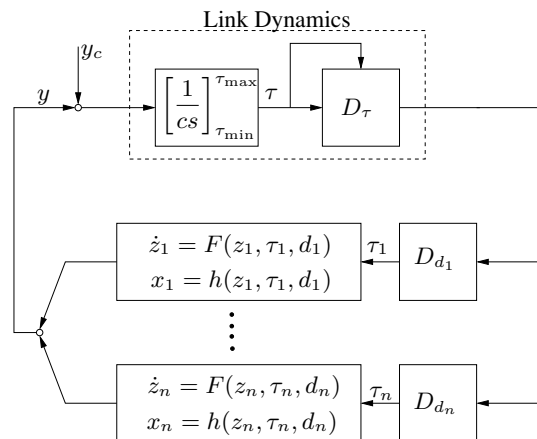


Fig. 2. A flow model for the single-link-multiple-user case.

with a time-varying delay operator

$$(D_{\tau}v)(t) = v(t - \tau(t)). \quad (1)$$

The buffer length τ represents a natural measure of the congestion in the link but this information is not available at the user controller until one round-trip time, $T_k = \tau + d_k$, after the corresponding packets were sent. This delay will have a negative effect of the achievable control performance. Note $0 \leq \tau_{\min} < \tau_{\max}$ models the size of the buffer.

We will investigate how different signaling information affects the fairness, efficiency, and the control performance of the system. Each user controller has the general form

$$\begin{aligned} \dot{z}_k &= F(z_k, \tau_k, d_k), \\ x_k &= h(z_k, \tau_k, d_k), \end{aligned}$$

i.e. there is a direct term between the rate and the delays. Throughout the paper, we make the following assumption on the controllers.

Assumption 1: F, h are continuously differentiable, and $F_z := \frac{\partial F}{\partial z}(z, \tau, d)$ is a Hurwitz matrix for all admissible z, τ, d .

The assumption means we only consider smooth and locally stable controllers. Furthermore, we consider the following three special cases of control architecture:

RTTC: Only the round-trip times $T_k = \tau + d_k$ are used for control, i.e.

$$\dot{z}_k = F(z_k, T_k), \quad x_k = h(z_k, T_k). \quad (2)$$

This structure is easy to implement since the round trip time can be measured with high accuracy. However, we will see in the next section that this signaling information is insufficient and either leads to a lack of fairness or to poor control performance.

QC: Only the queuing delay τ_k is used for control, i.e.

$$\dot{z}_k = F(z_k, \tau_k), \quad x_k = h(z_k, \tau_k). \quad (3)$$

In this case it is possible to achieve fairness and full utilization of the link capacity, but we show in Section IV

that the control performance is highly dependent on the propagation delays.

QPC: Both the queuing delay τ and the propagation delays are used for control, i.e.

$$\dot{z}_k = F(z_k, \tau_k, d_k), \quad x_k = h(z_k, \tau_k, d_k). \quad (4)$$

Using this structure, it is possible to achieve fairness and full utilization of the link capacity, and it is also possible to maintain robustness of stability under changing round-trip times, as discussed in Section IV.

III. FEASIBILITY OF FAIRNESS AND EFFICIENCY

In this section we investigate the equilibrium properties of the flow control system. Let us first focus on the case with one control input, i.e. the cases RTTC and QC. At stationarity it holds

$$\begin{aligned} 0 &= F(z_k^o, u_k^o), \quad k = 1, \dots, n, \\ x_k^o &= h(z_k^o, u_k^o), \quad k = 1, \dots, n, \end{aligned}$$

where z_k^o, u_k^o, x_k^o , denotes the equilibrium point ($u_k = T_k$ in RTTC and $u_k = \tau_k$ in QC). Since F_z is invertible by Assumption 1, it follows by the implicit function theorem that there (locally) is a smooth function $z_k^o = z_k^o(u_k^o)$, and thus a smooth function g defined by

$$x_k^o = h(z_k^o(u_k^o), u_k^o) =: g(u_k^o). \quad (5)$$

In the following, we assume that the function g is globally uniquely defined for all admissible control inputs. We will show in this and the next section that performance requirements on fairness and stability imposes further restrictions on g .

We use the following definition of fairness.

Definition 1: The single-link-multiple-users flow control model is said to be *fair* if for any set of propagation delays and any constant cross traffic y_c^o , the equilibrium rates x_k^o are equal. It is *efficient* if $\sum_{k=1}^n x_k^o = c - y_c^o$, for all $y_c^o \in [0, c]$ and $n \geq 1$, i.e. all users get an equal share of the available bandwidth.

1) *Control Based on Round-Trip Time (RTTC):* We first consider the case when the round-trip times are used for control, i.e. when $T_k = \tau + d_k^o$ and the control is defined by F and h in (2). Then we have the following negative result.

Observation 1: It is impossible to obtain both fairness and efficiency if round-trip times are used for control.

To justify the claim, let us assume the system is fair and efficient for a particular amount of cross traffic. By equation (5) with $u_k^o = T_k^o$ we have $x^o = g(T_k^o)$. Since this equality must true for arbitrary T_k^o , it follows that g is a constant function. If the desirable efficient equilibrium x^o now changes to $x^o + \delta x^o$, due to varying cross traffic for example, there is no solution to the equations $x^o + \delta x^o = g(T_k^o)$. This contradicts that the system can be both fair and efficient.

2) *Control Based on Queuing Delay (QC):* We will assume that the queuing delay is perfectly known and consider the rate-based control law in Figure 2. This in particular implies that F and h are given as in (3). Since the same control law is used by all users, the equilibrium rates are identical and a function of the equilibrium delay.

Observation 2: If queuing delay is used for control then the flow control model is fair and efficient if the function $g(\cdot)$ defined in (5) is surjective and such that $g(\tau_{\min}) = c$ and $g(\tau_{\max}) = 0$.

To justify this claim, note first that fairness follows from

$$x^o := x_k^o = g(\tau^o). \quad (6)$$

We have the following three cases at the equilibrium

$$\begin{aligned} x^o &= \frac{c - y_c^o}{n}, & \tau^o &\in (\tau_{\min}, \tau_{\max}), \\ x^o &\geq \frac{c - y_c^o}{n}, & \tau^o &= \tau_{\max}, \\ x^o &\leq \frac{c - y_c^o}{n}, & \tau^o &= \tau_{\min}, \end{aligned}$$

In the first case $y_c^o \in [0, c]$ and $n \geq 1$, which implies $\text{cl } g((\tau_{\min}, \tau_{\max})) = [0, c]$, where cl denotes the closure of a set. The last two cases corresponds to an over-utilization of the link capacity when the queue is full and an under-utilization of the link when the queue is empty. In order to prevent over-utilization (independent of cross-traffic) when $\tau = \tau_{\max}$ we need $g(\tau_{\max}) = 0$ and to prevent under-utilization (independent of cross-traffic) when $\tau = \tau_{\min}$ we need $g(\tau_{\min}) \geq c/n$. Since $n \geq 1$ this implies that $g(\tau_{\min}) = c$ suffices.

3) *Control Based on Queuing and Propagation Delay (QPC):* Since the control (4) is more general than the control (3), it follows that efficiency and fairness can be obtained if the conditions in Observation 2 are satisfied.

4) *Active Queue Management:* In the next section, we show that in order for the closed loop system to be stable, the function g in (5) must be monotonically decreasing. The queue length at equilibrium in an efficient system is then uniquely given as

$$\tau^o = g^{-1} \left(\frac{c - y_c^o}{n} \right),$$

i.e. it depends on the cross traffic and the number of users. This is sometimes undesirable and motivates the use of active queue management. We may assume without loss of generality that $\tau_{\min} = 0$ (otherwise let $\tau := \tau - \tau_{\min}$).

Consider the link dynamics (adopted from [2]) (here $\tilde{c} = c - y_c^o$ denotes the available capacity)

$$\begin{aligned} \dot{\tau}(t) &= \begin{cases} \frac{1}{c} \min\{y(t) - \tilde{c}, 0\}, & \tau = \tau_{\max}, \\ \frac{1}{c} (y(t) - \tilde{c}), & \tau \in (0, \tau_{\max}), \\ \frac{1}{c} \max\{y(t) - \tilde{c}, 0\}, & \tau = 0, \end{cases} \\ \dot{p}(t) &= \begin{cases} \alpha(\tau(t) - \tau_r) + y(t) - \tilde{c}, & p > 0, \\ \max\{\alpha(\tau(t) - \tau_r) + y(t) - \tilde{c}, 0\}, & p = 0, \end{cases} \end{aligned}$$

where τ denotes the buffer length and p is the congestion measure used for control, i.e. $u_k^o = p_k^o$ in (5). The reference

$\tau_r \in (0, \tau_{\max})$ is the desired buffer length and α is a positive design parameter. In the next result we provide conditions under which the buffer length converges to this reference value.

Observation 3: The flow control system with active queue management is fair and efficient if the function $g(\cdot)$ in (5) is surjective, and such that $g(0) = c$. Moreover, under these conditions the equilibrium queue length will be $\tau^o = \tau_r$ unless $n = 1$ and $y_c^o = 0$ in which case we only can ensure $\tau^o \in [0, \tau_r]$ (which is realistic since a single user have access to all the link capacity).

To justify this claim, note that since $x^o := x_k^o = g(p^o)$, $k = 1, \dots, n$ it follows that the flow distribution is fair at equilibrium. Next we prove efficiency.

Clearly we need $\text{cl } g([0, \infty)) = [0, c]$ for the efficiency. To justify the end point constraints, we first assume the network resources are under-utilized at equilibrium, i.e. $x^o = g(p^o) < \tilde{c}/n$. This would imply that $\tau^o = p^o = 0$ and thus $g(p^o) = g(0) = c$; a contradiction. If the network resources are over-utilized at equilibrium, i.e. $x^o = g(p^o) > \tilde{c}/n$, then $\tau^o = \tau_{\max}$ and hence $\alpha(\tau^o - \tau_r) + nx^o - \tilde{c} > 0$, which contradicts that the congestion signal p is at equilibrium. Thus, efficiency is proven.

Finally, if $\tau^o \neq \tau_r$ then either $\dot{p} \neq 0$ which contradicts that the solution is at an efficient equilibrium or $\tau^o < \tau_r$ and $p^o = 0$, which implies $x^o = g(0) = c$, i.e. there is only one user exploiting all the link capacity.

IV. CONTROLLER LIMITATIONS

In Section III, it was shown that to obtain fairness and efficiency, at least the information structure QC in (3) was needed. In this section, we show what limitations there are in the structure QC to ensure stability and robustness of the closed-loop system. This is done using local analysis. This means the closed-loop system is linearized, and that the time-varying nature of the delay D_τ is neglected. A justification of the later simplification is given in the appendix.

A. Some Basic Properties for Rate Based Control

Our first observation is far from sensational but still worth stating (see (6)):

Observation 4: When all users use the same feedback signal, the rate equilibria will be identical. Moreover, different rate equilibria have to correspond to different equilibrium delays.

We will now further examine the required properties of the function g which relate the delay and rate equilibria. For this purpose, we consider the linearized system

$$\begin{aligned} \dot{\Delta z} &= F_z \Delta z + F_\tau \Delta \tau, \\ x &= h_z \Delta z + h_\tau \Delta \tau, \end{aligned}$$

where $F_z = \frac{\partial F}{\partial z}(z^o(\tau^o), \tau^o)$ etc. This corresponds to the controller

$$C(s, \tau^o) = h_z(sI - F_z)^{-1} F_\tau + h_\tau, \quad (7)$$

where s denotes the Laplace variable. The loop gain of the linearized system can be written

$$L(s) = \frac{1}{cs} e^{-s\tau^o} \sum_k e^{-sd_k} C(s, \tau^o).$$

To ensure local stability, the controller has to ensure negative feedback is obtained. In fact when there are no time-delays and the controller is stable, a necessary requirement is that the static gain of the controller is negative.

Proposition 1: Consider the loop gain

$$L(s) = \frac{1}{s} C(s),$$

where $C(s) = B(s)/A(s)$ is a proper rational stable controller. If the feedback system is internally stable then $C(0) < 0$.

Proof: Let $B(s) = b_0 s^n + \dots + b_n$ and $A(s) = s^n + \dots + a_n$. Then $a_n > 0$ since the controller is assumed stable and the characteristic equation for the closed loop system is $sA(s) - B(s)$ which has b_n as its constant coefficient. Thus closed loop stability implies $b_n < 0$. Thus $C(0) = b_n/a_n < 0$. ■

In view of the preceding observation, the control law in Figure 2 should be such that the static gain of the linearization satisfies

$$C(0, \tau^o) = -h_z F_z^{-1} F_\tau + h_\tau < 0. \quad (8)$$

Now let us return to g defined in (5). Using the implicit function theorem gives

$$g'(\tau^o) = -h_z F_z^{-1} F_\tau + h_\tau = C(0, \tau^o).$$

In light of the inequality in (8) we can make our next observation:

Observation 5: For a locally stabilizing control law that is in-itself stable, the function (5) relating equilibrium delay to equilibrium rate is strictly monotonically decreasing.

Notice that strict monotonicity makes the function g invertible and it is thus possible to infer the equilibrium delay from the equilibrium rate. This is sometimes desirable from a control perspective and thus another reason for ensuring that g is strictly monotonic.

B. Properties under Full Utilization

We will here revisit the equilibrium properties discussed in Section III in the new context where the assumption on finiteness of the queue has been removed.

The desired conditions at stationarity are

$$\sum_k x_k^o = c - y_c^o, \quad x_k^o = x^o, \quad k = 1, \dots, n$$

resulting in

$$x^o = \frac{c - y_c^o}{n}. \quad (9)$$

Now we return to the properties of $x^o = g(\tau^o)$. We already have established that g is monotonically decreasing. Since $x^o \geq 0$, g is bounded from below and thus the limit $\lim_{\tau \rightarrow \infty} g(\tau)$ exists. If (9) is to be satisfied for any c , y_c^o

and n the control law must be such that this limit is zero. For the same reason $\lim_{\tau \rightarrow 0^+} g(\tau) = +\infty$ must hold. We summarize these requirements:

Proposition 2: In order to ensure that full utilization is ensured regardless of capacity c , cross-traffic intensity y_c^o and number of users n , the control law must be such that

$$\lim_{\tau \rightarrow 0^+} g(\tau) = +\infty, \quad (10)$$

$$\lim_{\tau \rightarrow \infty} g(\tau) = 0. \quad (11)$$

The loop gain of the linearized system can be written

$$\begin{aligned} L(s) &= \frac{1}{cs} C(s, \tau^o) e^{-s\tau^o} \sum_{k=1}^n e^{-sd_k} \\ &= \frac{1}{(nx^o + y_c^o)s} C(s, \tau^o) e^{-s\tau^o} \sum_{k=1}^n e^{-sd_k} \\ &= \frac{1}{(x^o + \frac{y_c^o}{n})s} C(s, \tau^o) e^{-s\tau^o} \frac{1}{n} \sum_{k=1}^n e^{-sd_k}. \end{aligned} \quad (12)$$

The next question is whether it is possible to design a $C(s, \tau^o)$ such that the linearized system is stable and has good performance regardless of c , y_c^o , n and propagation delays d_1, \dots, d_n . We make the following two observations.

Observation 6: The decentralized controllers RTTC and QC with the control architectures (2)–(3) can always be destabilized for large enough propagation delays.

Observation 7: The decentralized controllers RTTC, QC, and QPC with the control architectures (2)–(4) can all suffer from poor performance (low bandwidth) under varying cross-traffic.

We can justify the first observations as follows. Suppose that for given c , y_c^o , n and propagation delays d_1, \dots, d_n we have designed L to have stability. Notice that due to the integrator there must be frequencies for which $|L| > 1$. Now replace d_k with $d_o + d_k$, $k = 1, \dots, n$, and let $d_o \rightarrow \infty$. Then the magnitude curve of L remains the same but there is an additional delay e^{-sd_o} which eventually will render the system unstable.

We can justify the second observation as follows. Suppose that for given c , y_c^o , n and propagation delays d_1, \dots, d_n we have designed L to have stability and good performance. Now keeping d_1, \dots, d_n , n and x^o constant (the latter condition implies that the controller remains the same), increase y_c^o (and consequently also c). This will cause gain of L to decrease at all frequencies and by making y_c^o sufficiently large the gain of L can be made arbitrarily small.

If we assume the controllers have the structure QPC in (4), it is possible to avoid destabilization due to varying propagation delays, since the controllers are in the form $C(s, \tau^o, d_k)$. However, the problem with cross-traffic does not disappear with additional information about propagation delays. This will be illustrated in the case of FAST TCP in Section VI. Cross-traffic therefore remains a large obstacle for all the proposed information structures.

Hence, it seems the buffer must signal the users how much cross-traffic is present in order for the system to

have robustness of performance. Limitations on achievable bandwidth is further studied in the next section.

V. LIMITATIONS ON ATTAINABLE BANDWIDTH

It is well known that time delays give serious limitations on the attainable bandwidth of a closed-loop control system, see for example [1]. The bandwidth is here characterized by the smallest frequency ω_{gc} such that the loop gain satisfies $|L(j\omega_{gc})| = 1$. Disturbances of frequency lower than ω_{gc} are attenuated in the closed-loop system. Next, we discuss how various distributions of the propagation delay and the buffer delay changes the performance of QC controllers. We do not discuss RTTC controllers here because we showed earlier they are not fair and efficient. QPC controllers are briefly discussed at the end of this section.

A. Delay-Induced Limitations

To get some further insight we first consider the simplest case when $d_k = d$, $k = 1, \dots, n$. This gives the loop gain

$$L(s) = \frac{n}{cs} C(s, \tau^o) e^{-s(\tau^o+d)}.$$

We make the following observation.

Observation 8: For a QC-controlled system with a homogeneous propagation delay d , the bandwidth ω_{gc} cannot be much larger than $\pi/(2(\tau^o + d))$.

Thus it is clearly desirable to keep down both propagation delays and buffer delay to improve performance.

We can justify the observation using the Nyquist stability criterion that says that

$$\arg L(j\omega_{gc}) > 0 \quad (13)$$

to obtain stability of the closed-loop system. Note that positive feedback is here assumed. Assume furthermore that we want

$$\left. \frac{d \log |L(j\omega)|}{d \log \omega} \right|_{\omega=\omega_{gc}} \approx -1. \quad (14)$$

The condition (14) ensures that the bandwidth does not vary too much under small changes in the loop gain, and gives a (large) phase margin of $\pi/2$ radians when there are no time delays and the controller is minimum phase. This follows from Bode's relation, see for example [1]. If the time delay $e^{-s(\tau^o+d)}$ is added in the loop, the stability condition (13) becomes $\pi/2 - \omega_{gc}\tau^o - \omega_{gc}d > 0$, which gives the bound $\omega_{gc} < \frac{\pi}{2(\tau^o+d)}$ in the observation.

A slightly more complicated case is when the links have different, but small, propagation delays. We make the following observation.

Observation 9: For a QC-controlled system with small heterogeneous propagation delays d_k , the bandwidth ω_{gc} cannot be much larger than $\pi/(2(\tau^o + \bar{d}))$, where \bar{d} is the mean propagation delay.

The justification is as follows. Define the sum-of-propagation-delays factor as

$$d(s) := \frac{1}{n} \sum_{k=1}^n e^{-sd_k}, \quad (15)$$

with a slight abuse of notation. When the delays d_k are heterogeneous, $d(s)$ can exhibit very complicated behavior for large $|s|$ as we shall see later. For small complex frequencies $|s|$ or small d_k , the following approximation is valid, however,

$$\begin{aligned} d(s) &= (n - sd_1 - sd_2 - \dots - sd_n + O(|s|^2))/n \\ &= e^{-s\bar{d}} + O(|s|^2), \quad |s| \rightarrow 0, \end{aligned} \quad (16)$$

where \bar{d} is the mean propagation delay, $\bar{d} := \frac{1}{n} \sum_{k=1}^n d_k$. Hence, if the propagation delays are sufficiently small, we can take $d(s)$ into account as an additional time delay \bar{d} . The loop gain of the system is

$$L(s) = \frac{n}{cs} C(s, \tau^o) e^{-s\tau^o} d(s).$$

Equations (13) and (14) give that we must have $\pi/2 - \omega_{gc}\tau^o - \omega_{gc}\bar{d} > 0$, which yields the bandwidth bound

$$\omega_{gc} < \frac{\pi}{2(\tau^o + \bar{d})}. \quad (17)$$

The relation (17) only holds when the approximation (16) is valid. When there are n_1 sources of delay d_1 , and n_2 sources of delay d_2 , the analysis can be made more exact. We then have the following observation.

Observation 10: For a QC-controlled system with n_1 propagation delays d_1 and $n_2 \leq n_1$ propagation delays d_2 , the bandwidth ω_{gc} cannot be much larger than

$$\frac{\pi/2 - \arcsin n_2/n_1}{\tau^o + d_1},$$

if d_2 is variable.

Hence, if $n_2 \rightarrow n_1$ it is impossible to stabilize the closed loop for arbitrary delays d_2 under the condition (14). One then need to accept a less steep slope of the loop gain at ω_{gc} , and the bandwidth becomes more sensitive to small gain variations.

The justification is as follows. Let $n = n_1 + n_2$ and define $\Delta d := |d_2 - d_1|$. We then have,

$$d(j\omega) = \frac{e^{-j\omega d_1}}{n} (n_1 + n_2 e^{\pm j\omega \Delta d}).$$

The magnitude $|d(j\omega)|$ oscillates between the extremum values

$$\begin{aligned} |d(j\omega_{\max})| &= 1, & \omega_{\max} &= \frac{2k\pi}{\Delta d}, \\ |d(j\omega_{\min})| &= \frac{n_1 - n_2}{n}, & \omega_{\min} &= \frac{\pi}{\Delta d}(1 + 2k), \quad k \in \mathbb{Z}. \end{aligned}$$

Furthermore, the phase of $d(j\omega)$ becomes

$$\arg d(j\omega) = -\omega d_1 - \phi_{\Delta d}(\omega),$$

where $\phi_{\Delta d}(\omega) \in [-\arcsin \frac{n_2}{n_1}, \arcsin \frac{n_2}{n_1}]$, and $\phi_{\Delta d}(\omega)$ oscillates with the same frequency period as $|d(j\omega)|$. Hence if $n_1 \gg n_2$, there will not be very large variations in the magnitude of $d(j\omega)$. In these cases it is enough to take $d(s)$ into account as an extra time delay d_1 to the loop gain $L(s)$.

On the other hand, if $n_1 \approx n_2$, there will be very large variations in the magnitude and phase of $d(j\omega)$, with a

frequency period inversely proportional to the difference in time delay Δd . The phase correction $\phi_{\Delta d}$ attains values close to $\pm\pi/2$. If there are two groups of sources of roughly equal size and with very different propagation delay (Δd large), the loop gain will thus have zeros starting to appear already at the very low frequency $\pi/\Delta d$. The effect of this is further discussed in Section V-B. Equations (13) and (14) give that

$$\pi/2 - \omega_{gc}\tau^o - \omega_{gc}d_1 - \phi_{\Delta d}(\omega_{gc}) > 0$$

to ensure stability. Using that $\phi_{\Delta d}(\omega_{gc})$ is in the interval $[-\arcsin \frac{n_2}{n_1}, \arcsin \frac{n_2}{n_1}]$, we obtain the performance bound

$$\omega_{gc} < \frac{\pi/2 - \arcsin n_2/n_1}{\tau^o + d_1},$$

if we assume that the controller must be robust to arbitrary time delays d_2 .

As a last case to consider, let us assume that the delays d_k in $d(s)$ come from the same probability distribution, and are independent. For large n , the law of large numbers can be used to conclude that

$$d(s) = \frac{1}{n} \sum_{k=1}^n e^{-sd_k} \approx \mathbf{E}[e^{-sd}],$$

i.e. this term can be replaced by the moment generating function of the assumed distributions for the propagation delays. For example, assuming the propagation delays to be $\chi^2(m)$, gives

$$\frac{1}{n} \sum_{k=1}^n e^{-sd_k} \approx \mathbf{E}[e^{-sd}] = \frac{1}{(1 + 2s)^{m/2}}.$$

This is certainly a very simple expression, and is even a minimum phase transfer function, that gives the bound $\omega_{gc} < \pi/2\tau^o$. Hence, it seems like if the propagation delays are many, and enough spread out, their sum can be quite well behaved. The consequences of this relation is to be further investigated.

B. Zeros on Imaginary Axis

When there are two sources (or two groups of equal size) with propagation delays $d_1 \neq d_2$, the sum-of-propagation-delay factor $d(s)$ will have zeros on the imaginary axis,

$$d(j\omega_{k,z}) = 0, \quad \omega_{k,z} = \frac{\pi}{\Delta d}(1 + 2k), \quad \Delta d = |d_2 - d_1|,$$

where $k \in \mathbb{Z}$. Potentially these zeros can be canceled by adding poles in the controller $C(s)$, but this requires a very detailed knowledge of the difference Δd and would violate the assumption of a stable controller. Hence, the loop gain $L(s)$ inherits these zeros, and

$$L(j\omega_{k,z}) = 0.$$

This has at least two consequences for the performance of the closed-loop control system.

The first consequence is of importance if the bandwidth ω_{gc} is larger than $\omega_{1,z} = \pi/\Delta d$. Because of the zeros of $L(j\omega)$, the sensitivity function $S(s) = 1/(1 - L(s))$ is necessarily equal to one at $s = j\omega_{k,z}$. Now, for reasonable

control designs, $S(j\omega) \approx 0$ for (most) frequencies up to ω_{gc} . This means that there will be peaks in $|S(j\omega)|$ at all $\omega = \omega_{k,z} < \omega_{gc}$. These peaks result in an oscillation of frequency $\frac{\pi}{\Delta d}$ that is present in the response to disturbances acting on the system. The only way to avoid these oscillations is to lower the bandwidth below $\frac{\pi}{\Delta d}$.

The second consequence (which follows from the first) concerns the response to periodic oscillations in the cross-traffic flow-rate y_c^o . If its oscillation frequency is equal to $\omega_{k,z}$, the oscillation magnitude amplification of the flow-rate of source k is given by

$$|x_k(j\omega_{k,z})|/|y_c(j\omega_{k,z})| = \frac{\Delta d}{(1+2k)\pi c} |C(j\omega_{k,z})|.$$

If Δd is large, this can result in very large oscillations in x_k which clearly is not desirable for source k .

C. Attainable Bandwidth Using QPC

The previous limitations apply when the control structure QC is applied. It is clear that if QPC is used, more can be done since the controller takes the form $C(s, \tau^o, d_k)$. For the sake of simplicity, assume such a controller can be factorized into

$$C(s, \tau^o, d_k) = C_1(s, \tau^o) C_2(s, d_k).$$

The loop gain then becomes

$$L(s) = \frac{n}{cs} C_1(s, \tau^o) e^{-s\tau^o} \frac{1}{n} \sum_k C_2(s, d_k) e^{-sd_k}.$$

An optimal predictor $C_2(s, d_k) = e^{sd_k}$ would be ideal here, since the sum-of-propagation-delays factor would disappear. In this case the performance bound $\omega_{gc} < \pi/2\tau^o$ holds. However, this is impossible to implement. A simple approximation that cancels the sum-of-propagation-delays factor for sufficiently small propagation delays is $C_2(s, d_k) = 1 + sd_k$.

An alternative would be to consider a control scheme that has scale invariance of the variable s with respect to delays [11]. For example, we will in the next section show that FAST TCP has a loop gain on the form

$$L(s) = \frac{1}{c} \sum_{k=1}^n \frac{e^{-sT_k^o}}{sT_k^o} C(sT_k^o, \tau^o, d_k).$$

The first factor has a Nyquist curve that is independent of the round trip time $T_k^o = \tau^o + d_k$. This implies that the closed loop system can be made robust to variations in the delay provided that the second factor is such that $C(sT_k^o, \tau^o, d_k)$ does not vary much with the parameters τ^o and d_k . This can be understood since then the Nyquist curves of the terms in the expressions for the loop gain are close and the location of their linear combination easy to predict.

VI. APPLICATION TO FAST TCP

Consider first-order QPC user controllers in the form

$$\begin{aligned} \dot{z}_k &= F(z_k, \tau_k, d_k) = k(\tau_k, d_k)(z_k - a(\tau_k, d_k)), \\ x_k &= h(z_k, \tau_k, d_k) = b(\tau_k, d_k)z_k + e(\tau_k, d_k). \end{aligned} \quad (18)$$

In equilibrium, the flow of these controllers is given by

$$x_k^o = b(\tau^o, d_k)a(\tau^o, d_k) + e(\tau^o, d_k) =: g(\tau^o, d_k).$$

The FAST TCP protocol [15] is a protocol designed for high communication speed systems with large propagation delays. In [14], [4], a continuous-time flow model is derived for FAST TCP. It can be written in the form (18),

$$\begin{aligned} \dot{z}_k &= F(z_k, \tau_k, d_k) \\ &= \frac{\log(1 - \gamma\tau_k/(\tau_k + d_k))}{\tau_k + d_k} \left(z_k - \frac{\alpha(\tau_k + d_k)}{\tau_k} \right), \\ x_k &= h(z_k, \tau_k, d_k) = z_k/(\tau_k + d_k), \end{aligned}$$

where γ, α are real tuning parameters. We note that FAST TCP is a window based scheme and compared to [4] we have used a less accurate model for the relationship between rate x_k and the window size z_k . The equilibrium flow function g becomes

$$x^o = x_k^o = g(\tau^o, d_k) = \frac{\alpha}{\tau^o},$$

independently of k . The FAST TCP protocol is thus fair and efficient, in the sense of our previous definitions.

Next, let us compute the loop gain for FAST TCP. First we then need to linearize the system in equilibrium. Let us define

$$\zeta_k := \frac{\tau^o + d_k}{\log(1 - \gamma\tau^o/(\tau^o + d_k))}.$$

The transfer function from τ_k to x_k of the controller becomes

$$C(s, \tau^o, d_k) = \frac{\alpha}{\tau^o(\tau^o + d_k)} \left(\frac{d_k}{\tau^o(s\zeta_k - 1)} - 1 \right),$$

resulting in a loop gain

$$L(s) = \frac{\alpha}{c\tau^o s} e^{-s\tau^o} \sum_k \frac{e^{-sd_k}}{\tau^o + d_k} \left(\frac{d_k}{\tau^o(s\zeta_k - 1)} - 1 \right).$$

In equilibrium, efficiency gives that $nx^o = \tilde{c} = c - y_c$, and that $\tau^o = n\alpha/\tilde{c}$, resulting in the loop gain

$$L(s) = \frac{\tilde{c}}{c} \frac{1}{n} \sum_{k=1}^n \frac{e^{-sT_k^o}}{sT_k^o} \left(\frac{d_k}{\tau^o(s\zeta_k - 1)} - 1 \right),$$

where $T_k^o = d_k + \tau^o$ is the round-trip time at steady state. In order to analyze the loop gain we define

$$\psi_k := -\log(1 - \gamma\tau^o/(\tau^o + d_k)).$$

Then we get the equivalent expression

$$L(s) = -\frac{\tilde{c}}{c} \frac{1}{n} \sum_{k=1}^n \frac{e^{-sT_k^o}}{sT_k^o} \frac{sT_k^o + \psi_k T_k^o / \tau^o}{sT_k^o + \psi_k},$$

which can be simplified further provided that the gain γ is sufficiently small such that the following approximation is valid

$$\psi_k \approx \frac{\gamma\tau^o}{\tau^o + d_k} = \frac{\gamma\tau^o}{T_k^o}.$$

Then the loop gain simplifies to

$$\begin{aligned} L(s) &= -\frac{\tilde{c}}{c} \frac{1}{n} \sum_{k=1}^n \frac{e^{-sT_k^o}}{sT_k^o} \frac{sT_k^o + \gamma}{sT_k^o + \frac{\gamma T_k^o}{T_k^o}} \\ &= -\frac{\tilde{c}}{c} \frac{1}{n} \sum_{k=1}^n \frac{e^{-sT_k^o}}{sT_k^o} \frac{sT_k^o + \gamma}{sT_k^o + \gamma/(1 + d_k \tilde{c}/(n\alpha))}. \end{aligned}$$

From this expression we conclude the following properties of TCP-FAST

- 1) $0 < \tilde{c}/c \leq 1$ so there is a limit on how much the capacity influences the magnitude $|L(s)|$. In particular when $y_c = c$, the loop gain is zero.
- 2) The dynamics of FAST TCP is to large extent scale-invariant with respect to the round trip time, see the previous section for a discussion.
- 3) The variation in the gain and the time constant due to variations in the propagation delay scales as $d_k \tilde{c}/(n\alpha)$. Hence, the negative effects of heterogeneous propagation delays dissappear as the number of users increases.

The first property is as we have argued above inevitable unless information about the cross traffic is signaled to the users.

REFERENCES

- [1] K. J. Åström and R. M. Murray. *Feedback Systems – An Introduction for Scientists and Engineers*. Princeton University Press, 2008.
- [2] S. Athuraliya, S.H. Low, V.H. Li, and Y. Qinghe. REM: active queue management. *IEEE Network*, 15(3):48–53, 2001.
- [3] B. Briscoe. A fairer, faster Internet protocol. *IEEE Spectrum*, December 2008.
- [4] K. Jacobsson. *Dynamical modeling of Internet congestion control*. PhD thesis, Royal Institute of Technology (KTH), 2008.
- [5] J.M. Jaffe. Flow control power is nondecentralizable. *IEEE Transactions on Communications*, 29(9):1301–1306, September 1981.
- [6] C.-Y. Kao and B. Lincoln. Simple stability criteria for systems with time-varying delays. *Automatica*, 2004.
- [7] F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [8] S. Low, F. Paganini, and J. Doyle. Internet congestion control. *IEEE Control Systems Magazine*, 2(1):28–43, 2002.
- [9] Matthew Mathis. Reflections on the TCP macroscopic model. *ACM SIGCOMM Computer Communication Review*, 39(1):47–49, 2009.
- [10] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5), October 2000.
- [11] F. Paganini, Z. Wang, J. C. Doyle, and S. H. Low. Congestion control for high performance, stability, and fairness in general networks. *IEEE/ACM Transactions on Networking*, 13(1):43–56, 2005.
- [12] S. Shenker. A theoretical analysis of feedback flow control. In *Proceedings of ACM SIGCOM '90 Symposium*, Philadelphia, PA, September 1990.
- [13] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhäuser, 2004.
- [14] A. Tang, L. L. H. Andrew, K. Jacobsson, K. H. Johansson, S. H. Low, and H. Hjalmarsson. Window flow control: macroscopic properties from microscopic factors. In *IEEE Infocom*, Phoenix, AZ, USA, 2008.
- [15] D. X. Wei, C. Jin, S. H. Low, and S. Hegde. FAST TCP: motivation, architecture, algorithms, performance. *IEEE/ACM Transactions on Networking*, 14(6):1246–1259, 14(6):1246–1259, 2006.

APPENDIX: TIME-VARYING DELAY STABILITY

In our analysis above we used linearized models where the delay imposed on the packet flow is approximated by the equilibrium value of the queue length. It is well known that time-variations in the delay may cause stability problems

even if the system is linear. The following simple example illustrates the potential problem.

Example 1: Consider the simple linearized model

$$\begin{aligned} \Delta \dot{\tau}(t) &= \Delta x(t) + \Delta y_c(t) \\ \Delta x(t) &= -\Delta \tau(t - \tau(t)) \end{aligned} \quad (19)$$

where $\tau(t) = \tau^o + \Delta \tau(t)$. If the time-varying delay in (19) is replaced by the steady state approximation $e^{-s\tau^o}$, then it follows from the Nyquist criterion that the corresponding simplified system is stable if $\tau^o < \pi/2$. Let us instead assume (19) to be initially at equilibrium, i.e. $\Delta \tau(t) = \Delta x(t) = \Delta y_c(t) = 0$, $t < 0$, and let the cross traffic make a unit step at $t = 1$. Then the solution will be $\Delta \tau(t) = t$ and $\Delta x(t) = -\Delta \tau(t - \tau(t)) = -\Delta \tau(t - \tau^o - t) = -\Delta \tau(-\tau^o) = 0$. Hence, the queue grows without bound.

Despite, the negative result above it is possible to justify that equilibrium analysis above is locally correct. The linearized dynamics considered in the paper can be formulated as

$$\begin{aligned} v &= D_\tau \circ \frac{1}{s}(w + r), \\ w &= G(s)v, \end{aligned}$$

where $w = \Delta x$, $r = \Delta y_c$, $v(t) = \Delta \tau(t - \tau(t))$, D_τ is the time-varying delay operator in (1) and

$$G(s) = nC(s, \tau^o)d(s),$$

where $C(s, \tau^o)$ was defined in (7) and $d(s)$ in (15). After a loop transformation we obtain the equivalent system

$$\begin{aligned} \tilde{v} &= S_\tau(\tilde{w} + \tilde{r}), \\ \tilde{w} &= H(s)\tilde{v}, \end{aligned}$$

where $S_\tau = (D_\tau - D_{\tau^o}) \circ \frac{1}{s}$ and

$$\begin{aligned} H(s) &= \frac{sG(s)}{s - e^{-s\tau^o}G(s)}, \\ \tilde{r} &= \left(1 + \frac{e^{-s\tau^o}G(s)}{s - e^{-s\tau^o}G(s)}\right)r, \end{aligned}$$

where $H(s)$ is assumed stable, i.e. the system is stable if the queuing delay is replaced by its average value. The delay operator S_τ was studied in [6] where it was shown that its gain is proportional to $\|\tau - \tau^o\|_\infty$. In our system τ is a part of the state and its size is not known beforehand. To circumvent this problem we use the norm $\|v\| = \max(\|v\|_\infty, \|v\|_2)$. We have the system gains

$$\begin{aligned} \|H\| &= \sup_{\tilde{v} \neq 0} \frac{\|H\tilde{v}\|}{\|\tilde{v}\|} = \|H\|_1, \\ \|S_\tau\| &= \sup_{\tilde{w} \neq 0} \frac{\|S_\tau\tilde{w}\|}{\|\tilde{w}\|} = \|\tau - \tau^o\|_\infty, \end{aligned}$$

where the last property is proven analogously as in [6].

It is a consequence of the small gain theorem that the system is locally stable around the equilibrium if for some $\gamma < 1/\|H\|_1$, provided the input is bounded by

$$\|\tilde{r}\| \leq 1 - \|H\|_1 \gamma.$$

This proves our claim.