

Local Analysis of Structural Limitations of Network Congestion Control

Krister Jacobsson and Håkan Hjalmarsson[†]

Abstract—Recently there have been a number of interesting contributions to the stability analysis of network congestion control based on fluid models. Here, we further this emerging analysis by studying the structural limitations that so called primal/dual congestion control algorithms impose. Such algorithms rely on aggregated information from a network path, e.g. TCP-Vegas use the aggregated queuing delay. We show through local analysis that this imposes certain limitations of feedback control. Viewed from the source side, the complementary sensitivity and the sensitivity functions are severely restricted when many sources share the same bottleneck. This impose that source control must be small enough to achieve suitable noise rejection.

In addition, a specialized congestion control paradigm where all sources share a common time-base is analyzed. For this scenario the analysis facilitates significantly and robustness limitations towards configuration changes is observed.

I. INTRODUCTION

The tremendous complexity of the Internet makes it extremely difficult to model and analyze, and it has been questioned if mathematical theory can offer any major improvements in this area. However, recently significant progress in the theoretical understanding of network congestion control has been made following seminal work by Kelly and coworkers [1], [2] (see also the surveys [3], [4] and the book [5]). The key to success is to work at the correct level of aggregation—which is, fluid flow models with validity at longer time-scales than the round-trip time (RTT). By explicitly model the congestion measure signal fed back to sources, and posing the network flow control as an optimization problem—where the objective is to maximize the total source utility—it is shown that the rate control problem can be solved in a completely decentralized manner [1], [6], under the constraint that each source has a concave utility function of its own rate.

This optimization perspective of the rate control problem has been taken in a number of contributions. It also allows for dynamical laws, and the developed algorithms can be classified as: (1) primal, when the control at the source is dynamic but the link uses a static law; (2) dual, when the link uses a dynamic law but the source control is static; and (3) primal/dual, when dynamic controls are used both at the source and the links, see [7], [8], [9], for nice overviews.

By appropriate choice of utility function—even protocols not based on optimization, such as TCP Reno—can be interpreted as distributed algorithms trying to maximize the total utility [9], [10]. Delay based protocols such as TCP Vegas [11] or TCP FAST [12] can be classified as a primal/dual algorithm with queuing delay as dynamic link price (which is estimated at the source).

To ensure that the system will reach and maintain a favorable equilibrium, it is important to assess the dynamical properties—such as stability and convergence—of the schemes. Instability means that the protocol is unable to sustain the equilibrium; and manifests itself as severe oscillations in aggregate traffic quantities, such as queue lengths.

In the community, the main focus has been on stability; with numerous contributions concentrating on proving stability for more or less general configurations and scenarios as a result. Stability of the basic schemes was established already in [1], [6] but under very idealized settings. Similar work can be found in e.g. [13], [14], [15], [16]; however all results mentioned above have ignored the effect of network delay, which is critical for stability. Local stability of Reno/RED with feedback delays has been studied in [17], [18]. The stability analysis reveals that these protocols tend to become unstable when the delay increases and, more surprisingly, when the capacity increases. This has spurred an intensive research in protocols that maintain local stability also for networks with high bandwidth-delay product, see e.g., [19], [20], [21]. Other examples of work proving local stability when taking delay into consideration are [22], [23], [24], [25]. To be able to achieve global results but still not ignoring delay, the authors behind [26], [27], [28], uses so called Lyapunov-Krasovskii and/or Lyapunov-Razumikhin functionals to establish global convergence. An alternative approach is taken in [29] where global stability of a TCP/AQM setting is analyzed via integral-quadratic constraints (IQC). A further discussion on different methods for network analysis (focusing on stability issues) can be found in [30].

This paper has a slightly different approach than the above mentioned (and to authors knowledge, existing) contributions—here, the emphasis is on the dynamical properties such as system adaption and disturbance rejection; rather than on stability alone. We argue that it is of great interest to expose underlying system limitations, considering both system insight and future protocol design. Objective evaluation of link- and source controller performance without knowledge about system restrictions is in vain.

The paper is organized as follows. In Section II a well-known model of a network is presented and linearized. This model is used for local analysis in Section III, which includes a problem formulation and results regarding inherent performance limitations in general. To achieve deeper insight about the system, a case study is performed in Section IV. Conclusions and suggestions of future work directions are listed in Section V.

II. NETWORK MODEL

In this part we present the mathematical abstraction of the network. The model and notation follows the spirit of previous work found in e.g. [7], [31], [12].

[†]K. Jacobsson (krister.jacobsson@s3.kth.se) and H. Hjalmarsson (hakan.hjalmarsson@s3.kth.se) are with the School of Electrical Engineering, KTH, SE-100 44 Stockholm, Sweden.

This work was supported by European Commission through the project EuroNGI, by Swedish Research Council and SCINT.

A. Fluid-flow model

A generic communication network can be modeled as an indexed set of L resources (or links) each with an associated finite capacity c_l in packets per seconds, where $l \in \{1, \dots, L\}$. The network is shared by a set of N flows, competing about the offered capacity and indexed with $n \in \{1, \dots, N\}$. Every flow is uniquely identified by its source-destination pair.

To represent a certain network configuration (i.e. to associate flows with links they utilize) a *routing matrix* $R \in \mathbb{R}^{L \times N}$ is introduced. It is assumed to remain fixed and it is defined by: $R_{ln} = 1$ if link l is used by source n , and 0 otherwise. Furthermore it is implicitly assumed that R is properly posed—the configuration it represents is realizable.

Let a packet that is sent by flow n at time t appear at link l at time $t + \tau_{ln}^f$, this *forward delay* τ_{ln}^f models the amount of time it takes to travel from source n to link l , and it accounts for total latency and queuing delays. Define the *backward delay* τ_{ln}^b in the same manner: the time it takes between a packet pass link l and the corresponding acknowledgment is received at source n . The *round-trip time* associated with source n , is in this context naturally defined as $\tau_n := \tau_{ln}^f + \tau_{ln}^b$. In a buffering network round-trip time is generally time-varying since queues normally are fluctuating. This is *not* accounted for in the model and since our analysis is linear τ_n , τ_{ln}^f and τ_{ln}^b are all instead assumed to be represented by their corresponding equilibrium values in the sequel. The validity of the model is therefore *only* in time-scales coarser than the round-trip time (see e.g. [7] for a more thoroughly discussion).

The (continuous) transmission rate $x_n(t)$ in packets per second of source n at time t is related to the source congestion window $w_n(t)$ and the round-trip time as $x_n(t) := \frac{w_n(t)}{\tau_n}$ and is accurate only for the longer time scales considered. The aggregate flow $y_l(t)$ at link l is straightforwardly determined by the equation

$$y_l(t) = \sum_{n=1}^N R_{ln} x_n(t - \tau_{ln}^f) =: r_f(x_n, \tau_{ln}^f) \quad (1)$$

where $y_l(t)$ must not exceed the associated capacity c_l in equilibrium.

Motivated by seminal work by Kelly and coworkers [1], [2] and Low and Lapsley [6] the congestion measure signal fed back to sources is modeled explicitly: each link has an associated congestion signal referred to as *price* $p_l(t)$, further it is assumed that every individual source has access to the *aggregate price* $q_n(t)$ of all links along its path. This is mathematically expressed as

$$q_n(t) = \sum_{l=1}^L R_{ln} p_l(t - \tau_{ln}^b) =: r_b(p_l, \tau_{ln}^b). \quad (2)$$

Obviously it is crucial that the actual price information at links is distributed to- or estimated at sources.

B. Primal/dual control

As presented in [1], [6], the introduction of the concepts of aggregate rate $y_l(t)$ and aggregate price $q_n(t)$, allow specifications of dynamical source- and link laws that solves the resource allocation problem in a distributed manner—that is, *without* the need of explicit communication between

sources. In, so called, primal/dual control [1] each source is assumed to adjust its transmission rate *dynamically* based on the aggregate price $q_n(t)$ observed,

$$\dot{x}_n = k_n(x_n, q_n, \tau_n), \quad (3)$$

and similarly each link updates its price *dynamically* based on the amount of traffic it carries according to

$$\dot{p}_l = f_l(p_l, y_l, c_l). \quad (4)$$

In a physical network the source control (3) is realized by a dynamical window algorithm in an appropriate window based transmission control protocol such as e.g. TCP NewReno [32] or TCP Vegas [11], while the link control (4) is managed by an AQM scheme such as RED [33] or REM [34].

At this stage, the only *apriori* assumptions we make about the control is that: (3) and (4) stabilize the closed-loop system presented in Fig. 1; and furthermore, that the




Fig. 1. Schematic representation of the delayed closed loop system of interconnected sources/resources.

sources have a so called *demand curve* $x_n^* = \tilde{k}_n(q_n^*)$, that is a monotonically decreasing function of price. The demand curve is obtained from the source dynamics when the system has settled, i.e. solving for x_n^* in $k_n(x_n^*, q_n^*, \tau_n) = 0$. The stability assumption is necessary to guarantee that the system strives towards the optimal point (x^*, p^*) , and the constraints on the demand function is essential to achieve a strictly concave increasing *utility function*—and more important: a concave optimization problem—which fits into the optimization framework presented in [1], and hence solves the resource allocation problem.

Note that we avoid to explicitly define the entity of the price signal to keep the discussion as general as possible.

C. Linearized system

Since the focus of this paper is on *local* properties of the interconnected source/resource system, it is suitable to study small perturbations $x = x^* + \delta x$, $y = y^* + \delta y$, $p = p^* + \delta p$, $q = q^* + \delta q$ around the equilibrium. Fundamental assumptions for the linear analysis to hold are [31]:

- i) that the link control (4) is such that there exist *no* queuing delay close to the equilibrium point—for example, a *virtual* queue with target capacity slightly below the physical capacity—with consequence that τ_{ln}^f and τ_{ln}^b are constant and only dependent on the latency;
- ii) equilibrium queues are sufficiently far away from link saturations (such as buffer limits), this yields that changes in a source's rate δx_n is observed by every bottleneck along its path;
- iii) non-bottleneck links have zero price, hence *only* bottleneck links are taken into account in the analysis and

elements corresponding to non-bottlenecks are removed (but still, implicitly included in the model through the latency they contribute with).

Before proceeding, remark the slight abuse of notation that follows: the variables (x, y, p, q) are from now all representing *perturbations*; and the elements in y and p corresponding to non-bottlenecks, have been removed.

Assumption i) implies that (1) and (2) are time invariant, the Laplace transform is applicable which yields that the linearized aggregate quantities can be expressed as

$$y(s) = R_f(s)x(s) \quad (5)$$

$$q(s) = R_b^T(s)p(s) \quad (6)$$

in the frequency domain (superscript T denotes transpose). The forward delay matrix $R_f(s)$ is obtained, using the routing matrix R , by: first, eliminating rows in R corresponding to non-bottleneck links; then, where 1's appear—replacing those by the appropriate frequency domain forward delay $e^{-\tau_{ln}^f s}$. The backward delay matrix $R_b(s)$ is obtained similarly, using backward delays $e^{-\tau_{ln}^b s}$ instead. That only bottlenecks are accounted for implies that there are always at least as many sources as links, i.e. $N \geq L$ is true by definition. (Adjacent links with identical capacity carrying the same traffic are viewed as one link). Furthermore it is assumed that $R_f(0) = R_b^T(0)$ is of full rank.

Denote the linearized frequency domain source control $K_n(s)$ and link dynamics $F_l(s)$. Define $\mathcal{K}(s) := \text{diag}(K_n(s))$ and $\mathcal{F}(s) := \text{diag}(F_l(s))$ which yields that

$$x(s) = \mathcal{K}(s)q(s) \quad (7)$$

$$p(s) = \mathcal{F}(s)y(s) \quad (8)$$

in vector form. Now: (5), (6), (7) and (8)—together—defines the linear closed loop system, representing the interconnected source/resource system at equilibrium.

Only the reduced linearized system (and linear control) is considered in the sequel of the paper. It is also sometimes implicitly understood—by the argument—if it is the time domain or corresponding frequency domain expression that is considered.

III. LOCAL ANALYSIS

In this section it is discussed which signal mappings that are motivated to study. A structural performance limitation is revealed which questions the performance of *any* distributed control.

A. Design objectives

1) *Disturbance attenuation*: Only persistent sources—so called, “elephants”—are included in the fluid-flow model introduced in the previous section. However, the network is also utilized by short “mice” traffic not accounted for. This additional traffic will influence individual link prices $p_l(t)$, and subsequently propagate to the persistent sources via the aggregate price $q(t)$. A natural way to model this is to treat the “mice” as noise on the link rates, or similarly: noise (filtered through link dynamics) affecting the prices. Another issue is that bottlenecks only are included in the network model. Nevertheless, also non-bottleneck links along a source's path might be exposed to “mice” cross-traffic—contributing to $q_n(t)$ occasionally through the (not modeled)

price mechanism. Therefore, it is here claimed: that the effect of non-modeled traffic together with other disturbances is naturally modeled as additional noise on the aggregate price $q(t)$. As a consequence, it is concluded that the signal mapping between the aggregate price and the source rate

$$G_{xq} : q \mapsto x \quad (9)$$

should attenuate noise; which here is stated as a design objective. Furthermore, we also remark that the performance of the short lived flows are strongly connected to the behavior of the persistent sources through the price dynamics. This should be accounted for in the control design objective as well, especially considering the global control law not discussed in this paper. As an example, a delay based scheme should converge towards minimal queues.

2) *Adaptation*: The modeled network is utilized by N persistent sources, coexisting with uncontrolled “mice” traffic treated as noise. In a real setting, however, the number of sources N is not fixed—even persistent sources connects and disconnects occasionally. In a system perspective, it is of great interest that remaining sources adapt to the new conditions suitable fast and smooth. Unfortunately, such variations in the configuration—which results in changes in the routing matrix R —are not handled by the model since R is *not* allowed to change dynamically, as stated in Section II. The phenomena must instead be dealt with in an alternative way.

When a persistent source disconnects it means that the links along the source's path will experience a sudden negative change in the aggregate rate. This spare bandwidth should be effectively used by the remaining sources. On the other hand, when a new persistent source connects it means a (more or less sudden) temporary increase in the aggregate rate in the links it sends over. This increase should be observed by the other affected sources, that subsequently decrease their rates, with the result that the system converges towards a new equilibrium. These phenomena deriving from the configuration changes, are modeled by applying suitable signal changes in the aggregate traffic $y_l(t)$, meanwhile keeping the network configuration fixed. To be able to parry such changes in the configuration, it is crucial for the system that the mapping from the aggregate rate to the source rate

$$G_{xy} : y \mapsto x \quad (10)$$

tracks sufficient high frequencies (i.e. have suitable high 'bandwidth' in control nomenclature). Remark that—since the analysis is linear—only configuration changes resulting in small enough changes in $y_l(t)$ is considered. However, since the equilibrium system obeys the linear dynamics initially, the discussion is also valid in the initial phase of more major changes.

3) *Summary of control objective*: The MIMO transfer function from the aggregate rate $y(t)$ to the source rate $x(t)$ is given by:

$$G_{xy}(s) = (I + \mathcal{K}(s)R_b^T(s)\mathcal{F}(s)R_f(s))^{-1} \mathcal{K}(s)R_b^T(s)\mathcal{F}(s), \quad (11)$$

and similarly between the aggregate price $q(t)$ and $x(t)$:

$$G_{xq}(s) = (I + \mathcal{K}(s)R_b^T(s)\mathcal{F}(s)R_f(s))^{-1} \mathcal{K}(s), \quad (12)$$

where I denotes the identity matrix. Based on the previous discussion we conclude that the design objective is twofold:

- a) $\sigma_n(G_{xq}(j\omega))$ should be small for all frequencies, and
- b) $\sigma_l(G_{xy}(j\omega)) \approx 1$ for all frequencies,

where $\sigma_i(A)$ denotes the i th singular value of a matrix A ; this is unfortunately contradictive, as shown in the next subsection.

B. Structural limitation

1) *Preliminaries:* Before we proceed with the discussion on the properties of the singular values of G_{xq} and G_{xy} , we present the following preliminaries.

Lemma 3.1 (Fan's theorem): Let $\mathbb{C}^{l \times n}$ denote the space of complex $l \times n$ matrices, and $\sigma_i(A)$ the i th singular value of A ordered decreasingly. The double sided inequality

$$\sigma_i(A) - \sigma_1(B) \leq \sigma_i(A + B) \leq \sigma_i(A) + \sigma_1(B) \quad (13)$$

holds for all $A, B \in \mathbb{C}^{l \times n}$ and $i = 1, \dots, \min(l, n)$.

Proof: See [35]. ■

This result is used in the theorem below.

Theorem 3.1: Assume $R_f, R_b^T \in \mathbb{C}^{L \times N}$ and $\mathcal{K}, \mathcal{F} \in \mathbb{C}^{N \times N}$ are all of full rank, then if $N > L$ the maximum singular value of the sensitivity function $S := (I + \mathcal{K}R_b^T \mathcal{F}R_f)^{-1}$ fulfills $\sigma_1(S) \geq 1$.

Proof: Define the complementary sensitivity function as $T := I - S = (I + \mathcal{K}R_b^T \mathcal{F}R_f)^{-1} \mathcal{K}R_b^T \mathcal{F}R_f$. By definition $T \in \mathbb{C}^{N \times N}$. However, $\text{rank}(T) = L$ as $\text{rank}(R_b^T \mathcal{F}R_f) = L$ and S and \mathcal{K} both have full rank N . It follows that $\sigma_i(T) = 0$ for $i = N - L + 1, \dots, N$. From Lemma 3.1 we conclude that $\sigma_i(T) = \sigma_i(I - S) \geq \sigma_i(I) - \sigma_1(-S) = 1 - \sigma_1(S)$ and subsequently $\sigma_1(S) \geq 1 - \sigma_i(T) \geq 1 - \sigma_N(T) = 1 - 0 = 1$. ■

Remark 3.1: Note that $\sigma_1(S(j\omega)) \geq 1$ is true for all frequencies $\omega \in [0, \infty)$. For every frequency there is thus a direction in which the sensitivity function amplifies signals.

Remark 3.2: By inspection of $S = I - T$ it is realized that all input signals $x(t)$ with Fourier transform $X(j\omega)$ in the null-space of T also pass through S unaffected, or formally $S(j\omega)X(j\omega) = (I - T(j\omega))X(j\omega) = X(j\omega)$ if $X(j\omega) \in \ker(T(j\omega)) = \ker(R_b^T \mathcal{F}R_f)$. From this it can be concluded that there, by construction, exist at least $N - L$ different input directions where S lacks noise attenuation. If N is much larger than L , this limits the control performance for such a system.

2) *Limitation:* Rewriting (11) and (12) in terms of S yields

$$G_{xy}(s) = S(s)\mathcal{K}(s)R_b^T(s)\mathcal{F}(s), \quad (14)$$

$$G_{xq}(s) = S(s)\mathcal{K}(s), \quad (15)$$

and we observe that the source control $\mathcal{K}(s)$ must be small at the frequencies where disturbances affects the aggregate price, this to guarantee that G_{xq} attenuates noise when sources outnumber bottlenecks, $N > L$. Since control is completely decentralized, and in addition should be applicable in numerous scenarios spanning a wide range of network configurations, there is no motivation—from a performance point of view—for control design focusing on certain signal directions (of course, MIMO stability must be fulfilled). Making the singular values of $\mathcal{K}(s)$ small; or explicitly, each $K_n(s)$ small, is hence the best that can be achieved in this particular case.

Assume $K_n(s)$ is forced to be small to achieve noise suppression in $G_{xq}(s)$; then, to compensate for that and to get $G_{xy} \approx 1$, the link law $F_l(s)$ must have high gain. Since both source and link control are real systems and hence must be causal, however, $K_n(s)$ and $F_l(s)$ must be proper—with upper bounded high frequency gain as a consequence. This implies that there—due to the structure of the system—exists an unavoidable tradeoff between tracking abilities in $G_{xy}(s)$ and damping in $G_{xq}(s)$, whenever the number of sources are greater than the number of bottlenecks. This obviously constrains the achievable performance since the source/resource system is so vulnerable to certain signal directions.

IV. CASE STUDY: HOMOGENEOUS NETWORK

In this section we further investigate the results from the previous sections by studying a less general setting: a homogeneous network employing universal linear source- and link laws. This dramatically simplifies the calculations but still yields insight of some properties of the system.

A. The network

A class of source/resource systems obeying a certain structure are explored. The assumptions about the system are the following:

- 1) the price mechanism is universal on a local basis; that is, the link transfer function matrix can be written as $\mathcal{F}(s) = F(s)I$, where $F(s)$ is a scalar. This includes, of course, the single bottleneck case using any permissible algorithm; but also, for example, in a delay based setting a network with homogeneous capacity, employing virtual queues with identical target rates as price control law.
- 2) All (persistent) sources share a common round-trip time τ . This could be achieved by introducing an universal delay, by virtually holding the acknowledgment—when received at source—an appropriate amount of time. Note that this *does not* automatically imply decreased performance [36] for the MIMO source/resource system, and remember that it is control at equilibrium that is considered (it is not necessary to impose delays in the transient phase). (Of course, it could also be due to a favorable network configuration as well.)
- 3) It is assumed that all sources employ identical rate controllers around equilibrium; that is, the source transfer function matrix can be expressed as $\mathcal{K}(s) = K(s)I$, where $K(s)$ is a scalar. This is partly motivated by the previous statement about a common round-trip time, but also, for example, constrains us to only consider source laws that are proportional fair.

Furthermore it is also assumed that both link- and source laws are such that the system is open loop stable.

It could be argued that a system obeying the simplification above are too specialized to be of interest. However, such an example provides insight in what can not be attained in the general case—at least. It also, e.g., includes the case of multiple homogeneous protocols sending over a single bottleneck which should be representative for many applications and hence quite generic.

Before proceeding with the dynamical properties of the presented system, conditions for local stability are studied.

B. Local stability constraint

1) *Loop gain and a structural property:* Define the open loop transfer function

$$L(s) := \mathcal{K}(s)R_b^T(s)\mathcal{F}(s)R_f(s), \quad (16)$$

which can be rewritten as

$$L(s) = K(s)F(s)e^{-\tau s}R_f^*(s)R_f(s), \quad (17)$$

(superscript * denotes conjugate transpose) by introducing the scalar laws, and using that the backward delay matrix can be expressed in terms of the forward delay matrix and the round-trip time as $R_b(s) = R_f(-s)\text{diag}(e^{-\tau n s}) = e^{-\tau s}R_f(-s)$.

Introduce $\bar{R}_f(s) := R_f^*(s)R_f(s)$ which has n real (only) eigenvalues λ_n (ordered decreasingly), since it is Hermitian by construction. Furthermore it is easily verified that λ_n is always positive. Using that the spectral radius of a matrix is bounded by any induced matrix norm yields

$$\begin{aligned} \lambda_1 &\leq \|\bar{R}_f(s)\| = \|R_f^*(s)R_f(s)\|_\infty \\ &\leq \|R_f^*(s)\|_\infty \|R_f(s)\|_\infty = L \cdot N, \end{aligned} \quad (18)$$

where $\|\cdot\|_\infty$ represent the max-row-sum norm. Hence $0 \leq \lambda_n \leq L \cdot N$ is true for all n .

2) *Stability using the generalized Nyquist criterion:* To show local stability for a MIMO system in the presence of time delays, the generalized Nyquist criterion [37] is suitable; it implies that if a system is open loop stable, the feedback system is locally stable if each eigenloci of the open loop transfer function does not encircle -1 . Interpolating this result to the source/resouce system yields that if $L_n(s) := K(s)F(s)e^{-\tau s}\lambda_n$ does not encircle -1 for all n , local stability is guaranteed.

Define the phase-crossover frequency ω_p by the identity

$$\begin{aligned} \arg L_n(j\omega_p) &= \arg K(j\omega_p)F(j\omega_p)e^{-j\omega_p\tau}\lambda_n \\ &= \arg K(j\omega_p)F(j\omega_p) - \omega_p\tau = -\pi. \end{aligned} \quad (19)$$

Obviously ω_p is independent of the forward- and backward-delay matrices respectively. Examine the magnitude of the eigenloci,

$$\begin{aligned} |L_n(j\omega_p)| &= |K(j\omega_p)F(j\omega_p)e^{-j\omega_p\tau}\lambda_n| \\ &\leq |K(j\omega_p)| |F(j\omega_p)| \cdot L \cdot N \end{aligned} \quad (20)$$

which gives that local stability is achieved as long as

$$|K(j\omega_p)| |F(j\omega_p)| < \frac{1}{L \cdot N} \quad (21)$$

is fulfilled. This is in line with the findings about gain constraints for local stability found in [38], [25], [31]. However, due to the special structure of the system, the constraint is in this case relaxed to a single frequency only (which is determined by source control $K(s)$, link control $F(s)$ and the universal round-trip delay τ).

C. Dynamical properties

The task of designing a source control $K(s)$ and a link control $F(s)$ to obtain a damped $G_{xq}(s)$ (12) and a high 'bandwidth' $G_{xy}(s)$ (11) is further discussed below.

Define $\mathcal{C} = K(s)F(s)e^{-s\tau}$ for a more compact notation. Represent the forward delay matrix in a singular value decomposition, i.e. $R_f(s) = U\Sigma V^*$ where the singular

values in Σ is ordered decreasingly. Now, the sensitivity function S can be written as

$$\begin{aligned} S &= (I + \mathcal{C}R_f^*R_f)^{-1} = (I + \mathcal{C}V\Sigma^T\Sigma V^*)^{-1} \\ &= V(I + \mathcal{C}\Sigma^T\Sigma)^{-1}V^* = VS_\sigma V^* \end{aligned} \quad (22)$$

where

$$S_\sigma(s) := \begin{pmatrix} \text{diag}\left(\frac{1}{1+K(s)F(s)e^{-s\tau}\sigma_l^2(R_f(s))}\right) & \mathbf{0} \\ \mathbf{0} & I \end{pmatrix} \quad (23)$$

whose diagonal elements are also the eigenvalues of S . Subsequently, the singular values of $G_{xq}(s)$ are obtained by combining (15), (22), (23), and thus L of them is given by

$$\sigma_n(G_{xq}(j\omega)) = \left| \frac{K(j\omega)}{1 + K(j\omega)F(j\omega)e^{-j\omega\tau}\sigma_l^2(R_f(j\omega))} \right|, \quad (24)$$

where $l = 1, \dots, L$, and the remaining $N - L$ ones by

$$\sigma_n(G_{xq}(j\omega)) = K(j\omega). \quad (25)$$

Observe that the ordering of the singular values is frequency dependent. From (24) and (25) it is explicit that the system is sensitive to disturbance inputs with Fourier transforms in the null-space of $R_f(s)$, and that $K(s)$ must be small to guarantee noise attenuation.

For G_{xy} , note that

$$G_{xy} = \mathcal{C}SR_f^* = \mathcal{C}VS_\sigma\Sigma U^*. \quad (26)$$

Now straightforward calculations yields

$$\sigma_l(G_{xy}(j\omega)) = \left| \frac{K(j\omega)F(j\omega)\sigma_l(R_f(j\omega))}{1 + K(j\omega)F(j\omega)e^{-j\omega\tau}\sigma_l^2(R_f(j\omega))} \right|. \quad (27)$$

The issue of choosing $K(s)$ and $F(s)$ such that (24), (25) and (27) fulfills desirable properties is similar to the SISO control problem presented in Fig.2, with the objective to design a small $K(s)$ and an $F(s)$ such that the mapping $r \mapsto u$ is small and $r \mapsto y$ is approximately one. Using




Fig. 2. Block diagram of the SISO representative of (24), i.e. $r \mapsto u$, and (27), i.e. $r \mapsto y$.

high loop-gain via $F(s)$ (subject to stability constraints) as custom in SISO design does not solve the problem in this particular case. This is realized by observing that when $|K(j\omega)F(j\omega)\sigma_l^2(R_f(j\omega))| \gg 1$ we have

$$\begin{aligned} \sigma_l(G_{xy}(j\omega)) &\approx \left| \frac{K(j\omega)F(j\omega)\sigma_l(R_f(j\omega))}{K(j\omega)F(j\omega)e^{-j\omega\tau}\sigma_l^2(R_f(j\omega))} \right| \\ &= \frac{1}{\sigma_l(R_f(j\omega))}. \end{aligned} \quad (28)$$

This means that the network configuration (routing and forward/backward delay distribution) crucially affects system performance and can not be compensated for with decentralized control. The conclusion is that it is impossible to

make control robust with respect to configuration from a performance point of view using high loop gain. It also suggests that having an integrating term in $F(s)$ (as, e.g., in the case of a FIFO queue policy) when not accounted for with derivative terms in $K(s)$, may not be a suitable choice since static properties become heavily dependent on the configuration of the routing matrix R . However, alternatives have not yet been studied.

V. CONCLUSIONS AND FUTURE WORK

In this paper we have taken a slightly different approach than customary in literature on network congestion control with a control theoretic perspective. Rather than focusing on stability alone, we have studied dynamical properties of a source/resource system such as adaptation and disturbance attenuation on a local basis. By identifying and analyzing relevant signal mappings the problem is formulated, and it is revealed that, viewed from the source side, the complementary sensitivity and the sensitivity functions are severely restricted when many sources share the same bottleneck. It is also concluded that source control must be small at frequencies where disturbances appear to avoid noise amplification, this since certain signal directions are impossible to influence with link control. By studying a less general setting, a homogeneous set of sources with universal round-trip time sending over a network employing identical link control, it is furthermore shown that from a performance point of view the source/resource system can not be made robust with respect to changes in configuration.

A more explicit formulation of a suitable low-gain rate control in conjunction with a high-gain price mechanism obeying stability constraints is an issue for the future. The design should take the full non-linear control into consideration. We also point to the need for rigorous discrete event simulations (using, e.g., NS-2 or similar) to validate results, or rather justify the fluid flow model that is used. This is currently under investigation and will be reported elsewhere.

REFERENCES

- [1] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [2] F. P. Kelly, "Mathematical modelling of the Internet," in *Fourth International Congress on Industrial and Applied Mathematics, Edinburgh, Scotland*, July 1999.
- [3] F. Kelly, "Fairness and stability of end-to-end congestion control," *European Journal of Control*, vol. 9, pp. 159–176, 2003.
- [4] S. H. Low and R. Srikant, "A mathematical framework for designing a low-loss, low-delay Internet," *Networks and Spatial Economics*, vol. 4, pp. 75–101, 2004.
- [5] R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.
- [6] S. H. Low and D. E. Lapsley, "Optimization flow control – I: Basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, 1999.
- [7] S. H. Low, F. Paganini, and J. C. Doyle, "Internet congestion control," *Control Systems Magazine*, vol. 22, no. 1, pp. 28–43, February 2002.
- [8] S. Liu, T. Basar, and R. Srikant, "Controlling the Internet: A survey and some new results," in *Proceedings of the 42nd IEEE Conference on Decision and Control*, Maui, Hawaii USA, 2003, pp. 3048–3057.
- [9] S. H. Low and R. Srikant, "A mathematical framework for designing a low-loss, low-delay Internet," *Networks and Spatial Economics*, January-February 2003, special issue on "Crossovers between Transportation Planning and Telecommunications".
- [10] S. H. Low, "A duality model of TCP and queue management algorithms," in *Proceedings of ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, 2000.
- [11] L. S. Brakmo and L. L. Peterson, "TCP Vegas: end-to-end congestion avoidance on a global Internet," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 8, pp. 1465–1480, 1995.
- [12] C. Jin, D. X. Wei, and S. H. Low, "FAST TCP: motivation, architecture, algorithms, performance," in *Proceedings of IEEE Infocom*. IEEE, March 2004.
- [13] S. Kunniyur and R. Srikant, "A time-scale decomposition approach to adaptive ECN marking," *IEEE Transactions on Automatic Control*, vol. 47, no. 6, pp. 884–894, 2002.
- [14] E. Altman, T. Basar, and R. Srikant, "Robust rate control for ABR sources," in *IEEE Infocom*, San Francisco, CA, 1998, pp. 166–173.
- [15] C. V. Hollot and Y. Chait, "Nonlinear stability analysis for a class of TCP/AQM networks," in *IEEE CDC*, Orlando, FL, December 2001, pp. 2309–2314.
- [16] J. T. Wen and M. Arcak, "A unifying passivity framework for network flow control," *IEEE Transactions on Automatic Control*, vol. 49, no. 2, pp. 162–174, 2004.
- [17] C. Hollot, V. Mishra, D. Towsley, and W. B. Gong, "A control theoretic analysis of RED," in *Proceedings of IEEE Infocom 2001*, 2001.
- [18] S. H. Low, F. Paganini, J. Wang, S. A. Adlakha, and J. C. Doyle, "Dynamics of TCP/RED and a scalable control," in *IEEE Infocom*, vol. 1, New York, June 2002, pp. 239–248.
- [19] F. Paganini, J. C. Doyle, and S. H. Low, "Scalable laws for stable network control," in *IEEE CDC*, vol. 1, Orlando, FL, December 2001, pp. 185–190.
- [20] S. Floyd, "Highspeed TCP for large congestion windows," February 2003, RFC 3649.
- [21] T. Kelly, "Scalable TCP: improving performance in highspeed wide area networks," in *First International Workshop of Protocols for Fast Long-Distance Networks*, 2003.
- [22] R. Johari and D. Tan, "End-to-end congestion control for the Internet: delays and instability," *IEEE/ACM Transactions on Networking*, vol. 6, no. 9, pp. 818–832, 2001.
- [23] L. Massoulié, "Stability of distributed congestion control with heterogeneous feedback delays," Microsoft Research, Cambridge, UK, Technical Report, 2000.
- [24] G. Vinnicombe, "On the stability of end-to-end congestion control for the Internet," Cambridge University, Cambridge, UK, Technical Report CUED/F-INFENG/TR.398, 2000.
- [25] —, "On the stability of networks operating TCP-like congestion control," in *IFAC World Congress*, Barcelona, Spain, July 2002.
- [26] S. Deb and R. Srikant, "Global stability of congestion controllers for the internet," *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 1055–1060, June 2003.
- [27] L. Ying, G. E. Dullerud, and R. Srikant, "Global stability of internet congestion controllers with heterogeneous delays," in *Proceedings of the American Control Conference 2004*, July 2004.
- [28] F. Mazenc and S. I. Niculescu, "Remarks on the stability of a class of TCP-like congestion control models," in *IEEE Proceedings of Conference on Decision and Control*. IEEE, December 2003, pp. 5591–5594.
- [29] M. Peet and S. Lall, "On global stability of Internet congestion control," in *IEEE Proceedings of Conference on Decision and Control*. IEEE, December 2004, pp. 1035–1041.
- [30] A. Papachristodoulou, L. Li, and J. C. Doyle, "Methodological frameworks for large-scale network analysis and design," *ACM SIGCOMM Computer Communications Review*, vol. 34, no. 3, pp. 7–20, July 2004.
- [31] F. Paganini, Z. Wang, J. C. Doyle, and S. H. Low, "Congestion control for high performance, stability, and fairness in general networks," *IEEE/ACM Trans. Netw.*, vol. 13, no. 1, pp. 43–56, 2005.
- [32] S. Floyd and T. Henderson, "The NewReno modification to TCP's fast recovery algorithm," RFC 2582, April 1999.
- [33] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397–413, 1993.
- [34] S. Athuraliya, V. H. Li, S. H. Low, and Q. Yin, "REM: active queue management," *IEEE Network*, vol. 15, no. 3, pp. 48–53, 2001.
- [35] R. Horn and C. R. Johnson, *Topics in Matrix Analysis*. Cambridge, UK: Cambridge University Press.
- [36] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control*. Wiley, 1996.
- [37] C. A. Desoer and Y.-T. Wang, "On the generalized nyquist stability criterion," *IEEE Transactions on Automatic Control*, vol. 25, no. 2, pp. 187–196, April 1980.
- [38] F. Paganini, J. Doyle, and S. H. Low, "A control theoretical look at internet congestion control," in *Multidisciplinary Research in Control: The Mohammed Dahleh Symposium 2002*, ser. Lecture Notes in Control and Information Sciences, L. Giarré and B. Bamieh, Eds. Springer-Verlag, 2003, vol. 289.