

Methodological Frameworks for Large-scale Network Analysis and Design

Antonis
Papachristodoulou
California Institute
of Technology
Pasadena CA 91125
antonis@cds.caltech.edu

Lun Li
California Institute
of Technology
Pasadena CA 91125
lun@cds.caltech.edu

John C Doyle
California Institute
of Technology
Pasadena CA 91125
doyle@cds.caltech.edu

ABSTRACT

This paper emphasizes the need for methodological frameworks for analysis and design of large scale networks which are independent of specific design innovations and their advocacy, with the aim of making networking a more systematic engineering discipline. Networking problems have largely confounded existing theory, and innovation based on intuition has dominated design. This paper will illustrate potential pitfalls of this practice. The general aim is to illustrate universal aspects of theoretical and methodological research that can be applied to network design and verification. The issues focused on will include the choice of models, including the relationship between flow and packet level descriptions, the need to account for uncertainty generated by modelling abstractions, and the challenges of dealing with network scale. The rigorous comparison of proposed schemes will be illustrated using various abstractions. While standard tools from robust control theory have been applied in this area, we will also illustrate how network-specific challenges can drive the development of new mathematics that expand their range of applicability, and how many enormous challenges remain.

1. INTRODUCTION

The recent technological advances and the ubiquity of communication and large-scale networks have presented researchers with new design and analysis challenges. The integration of heterogeneous modules and their complex interactions make the problem seem intractable. The added requirements of reliability and evolvability call for an integration of new mathematical tools and more methodological frameworks to facilitate and also formulate the design and analysis processes. There are many systems around us that have evolved to have these properties, such as social and ecological systems and biological networks etc. Analysis of their functionality reveals the principles that either nature

or society followed to evolve. The question that arises then is how to design and analyze large-scale networks for robustness, reliability, and evolvability.

The multi-objective requirements for analysis and design of such systems should be based on rigorous and repeatable design methodologies and systematic evaluation frameworks. A simple example is the Internet where the first step towards the construction of protocols was based largely on intuition. Despite its brilliant innovation and enormous success in this particular case this far, this approach is definitely not adequate in the long run for large-scale network designs, as underestimating the importance of certain system features is usually only revealed in surprising failures after implementation. Such frameworks should not only provide comparability of various proposed designs but also facilitate the analysis of their functionality and bring to the surface pitfalls and propose cures. It is important both to provide analysis methods that go behind mere simulation of point scenarios, as well as methods that transcend the advocacy of a single design solution.

In this paper we consider as a case study network congestion control for the Internet, where the introduction of a unified framework enabled comparability and provided a better understanding of the shortcomings of certain TCP/AQM schemes, and allowing the construction of others. Indeed, the advances in network technology have revealed both deficiencies in the ad-hoc schemes as well as rigorous explanations for their scalability, which is astonishing, even in retrospect. A well-known result is that the performance of some TCP/AQM schemes deteriorates on increasing bandwidth-delay product networks [12]. While TCP/AQM will be the focus of this paper, as it has had a recent surge in progress, our aim is to use it to illustrate broader methodological issues that we claim are universally relevant to all complex network research.

Developing robust models for the modules comprising the system is the first, very important step both for analysis and design. For the various modules in congestion control, models of varying complexity can be constructed, ranging from stochastic multi-scale ones to simple deterministic [18]. Through a series of simple assumptions very fine scale models can be simplified into ones that current analysis tools can handle. This simplification process should account for all the approximations made and the analysis methodology should

take them into consideration, for otherwise the wrong conclusion for the original system properties may be obtained. Robust control theory offers a unique framework for robust model building by providing tools that can account for all the simplifications in the reduction process in a systematic way, through the introduction of uncertainty resulting in a parameterized family of models [2]. Analysis then follows by considering this *set* of models rather than a single nominal one. The uncertainty can be static or dynamic, in which case parts of the dynamics are ‘lumped’ together by considering worst-case scenario of input-output properties. Even when designing the system, static and dynamic uncertainty in the various system parameters can be incorporated in a similar manner.

In particular for network congestion control, this process produces uncertain deterministic nonlinear delay-differential equation models, for which in general no scalable analysis tools are available at the nonlinear level. Ignoring the effect of delays can be misleading in both analysis and design. Analysis by linearization allows scalability, however it can only provide a *local* picture which also can be misleading. It is common practice to simulate using **ns-2** more complex model descriptions: this can be used to investigate specific system behaviors, but can not guarantee functionality under all possible parameters and initial conditions, and cannot be used to study large networks. Thus what is really required is a scalable nonlinear analysis procedure that includes the effect of delays. This would allow scalable verifiability, one of the most important issues in design, both for TCP/AQM and hopefully ultimately for complex systems in general. In this paper we will show two alternative ways of investigating these uncertain nonlinear, delay-differential equation models: under fixed system dynamic structures, we will provide scalable conditions for the functionality of the network for arbitrary network sizes. And for various system descriptions we will provide a framework for establishing the functionality of specific network topologies with a complexity that varies in a polynomial manner with the size of the system description [25].

Both procedures have their own merits, and are equally important. For simple network topologies and arbitrary dynamics the algorithmic procedure can provide simple proofs of the functionality with tight conditions on the parameters in the system. This allows the verifiability of model descriptions not possible in the past, and expands the level of modelling complexity for which tools exist, from linear to nonlinear with delays. On the other hand, the procedure leading to scalable conditions on the functionality of arbitrary sized networks can provide conservative (relaxed) conditions for functionality at the price of being scalable.

The paper is organized as follows. In section 2 we present the unified framework that was used for network congestion control, and the optimization/duality formulation which decomposed the centralized problem into two decoupled ones. In section 3 we will present how robust models can be constructed for the various modules that can enable design, and in section 4 we will show how incomplete modelling or failure to take into account important features of the problem can lead to disastrous designs. In section 5 we present methodologies for evaluating the properties of networks that scale

with the network size, and present recent tools that allow analysis of even more complicated system descriptions.

2. UNIFIED FRAMEWORKS

Technological advances in the past years have called for the design of even more complex large-scale networks that offer evolvability and robustness. This can only be achieved by decentralizing the complex multi-objective optimization problem into smaller ones that can be solved efficiently, and at the same time solving the original one. At this top-level view of the problem, one is not concerned with the dynamics that the various modules will be endowed with; but rather how the whole network will synergistically operate so as to reach the desired equilibrium in a decentralized way. The next step, which we will discuss in the sequel, is the choice of appropriate dynamics for the modules so as to asymptotically stabilize them around this equilibrium.

The attempt to set up this mathematical framework for network congestion control began in the early 90’s. In [27], an individual feedback scheme with fairly shared gateways had been proposed to achieve a time-scale invariant, fair, stable and robust performance design. Jain [7] brings up the issue of proper congestion control mechanisms under heavy load. Kevshav [9] develops a continuous-time model to unify the treatment of various control laws. However, these models are mainly based on queuing systems which cannot capture the static and dynamic properties of the transport layer protocol.

Following the methodology defined in [8, 14, 13], consider a network of L communication links shared by S sources shown in Figure 1. The routing matrix R is given by:

$$R_{li} = \begin{cases} 1 & \text{if source } i \text{ uses link } l \\ 0 & \text{otherwise} \end{cases}$$

Each source i has an associated transmission rate x_i . All sources whose flow passes through link l contribute to the *aggregate rate* y_l for link l , the rates being added with some forward time delay $\tau_{i,l}^f$. Hence we have:

$$y_l(t) = \sum_{i=1}^S R_{li} x_i(t - \tau_{i,l}^f) \triangleq r_f(x_i, \tau_{i,l}^f) \quad (1)$$

The link l reacts to the aggregate rate y_l by setting congestion information p_l , the price at link l . This is the Active Queue Management (AQM) part of the picture that is to be designed. The prices of all the links that source i uses are aggregated to form q_i , the *aggregate price* for source i , again through a delay $\tau_{i,l}^b$:

$$q_i(t) = \sum_{l=1}^L R_{li} p_l(t - \tau_{i,l}^b) \triangleq r_b(p_l, \tau_{i,l}^b) \quad (2)$$

The prices q_i can then be used to set the rate of source i , x_i , which completes the picture. The whole interconnection is shown in Figure 1. The forward and backward delays can be combined to yield the Round Trip Time (RTT) for each source i :

$$\tau_i = \tau_{i,l}^f + \tau_{i,l}^b$$

The capacity of the link is assumed to be c_l . The functions f and g shown in Figure 1 are the source law and the link law




Figure 1: The Internet as an interconnection of sources and links through delays.

respectively. This setting is *universal*, and the only missing blocks are the two control laws that describe how the i th source reacts to the price signal q_i that it receives

$$\dot{x}_i = f_i(x_i, q_i, \tau_i), \quad (3)$$

and how the l th link reacts to the aggregate rate y_l it observes

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

Here f_i models TCP algorithms (e.g. Reno or Vegas) and g_l models AQM algorithms (e.g. RED, REM).

In order to understand the meaning of the variables p_l, q_i , let us consider the properties of the equilibrium of this system, which we assume is given by the vector quantities x^*, y^*, p^*, q^* . The framework we use will be based on convex optimization. Firstly, we have the following relationships from (1) and (2) at equilibrium:

$$y^* = R x^*, \quad q^* = R^T p^*.$$

We now assume that as the aggregate price signal at equilibrium q_i^* increases, then the demanded transmission rate at the source should decrease, i.e. the two are related by

$$x_i^* = F_i(q_i^*),$$

where F_i is a positive, strictly monotone decreasing function, which is the solution of $f_i(x_i^*, q_i^*) = 0$ where f_i is given by Equation (3). Alternatively, one can think that the sources have a certain *utility* if allowed a certain transmission rate, the utility satisfying:

$$U_i'(x_i) = F_i^{-1}(x_i).$$

This relationship implies that $U_i(x_i)$ is a monotonically increasing strictly concave function. Under these assumptions, the equilibrium rate x_i for each source solves

$$\max_{x_i} U_i(x_i) - x_i q_i^* \quad (5)$$

which means that the sources are trying to maximize their profit: maximize their utility, but at the same time minimize the cost of having high rates; q_i^* can be thought of as the *price* per unit flow that the sources have to pay. Different protocols correspond to different utility functions U_i , and to different dynamic laws (3–4) that attempt, in a decentralized way, to reach the appropriate equilibrium. For this reason this framework allows comparability of different designs.

The role of prices is to coordinate the actions of individual sources so as to align individual optimality with social optimality, i.e., to ensure that the solution of (5) also solves the network resource allocation problem

$$\begin{aligned} \max_{x_i \geq 0} \quad & \sum_{i=1}^S U_i(x_i) \\ \text{s.t.} \quad & \sum_{i=1}^S R_{li} x_i \leq c_l, \quad \forall l = 1, \dots, L, \end{aligned} \quad (6)$$

where the inequality constraint is the natural limitation that the sum of all transmission rates through link l has to be less than or equal to its capacity. The uniqueness of the optimal solution to the above problem is guaranteed since the U_i are strictly concave functions and the program is convex. The above optimization problem cannot be solved in a decentralized way, as the source rates are coupled in the shared links through the inequality constraints and solving for x^* would require cooperation among possibly all sources. To solve it in a distributed manner over a large network we can decompose it into a primal problem that the sources are trying to solve and a dual that the links are trying to solve, regarding the sources x_i as primal variables and the prices set by the links p_l as the dual variables. Specifically, consider the Lagrangian for program (6):

$$\begin{aligned} L(x, p) &= \sum_i U_i(x_i) - \sum_l p_l (y_l - c_l) \\ &= \sum_i \left(U_i(x_i) - x_i \sum_l R_{li} p_l \right) + \sum_l p_l c_l, \end{aligned}$$

where $p_l \geq 0$. The dual problem can then be written as

$$\min_{p_l \geq 0} \max_{x_i \geq 0} \sum_i \left(U_i(x_i) - x_i \sum_l R_{li} p_l \right) + \sum_l p_l c_l.$$

By the properties of the dual solution and the fact that the programs are convex, at the optimum p^* the x_i that maximizes the individual profit (5) is the same as the unique solution to the network problem (6). If the equilibrium prices p^* are made to align with the Lagrange multipliers, the individual optima — computed in a decentralized fashion by the sources — will align with the global optima of (6).

The dynamical system defined by (3–4) with delays ignored aims to drive the system close to or exactly at the optimal point (x^*, p^*) , using well-known gradient algorithms. After [8], it is customary to call ‘dual’ the congestion control scheme with dynamics at the links but a static source law; and ‘primal’ one with dynamics at the sources and a static link law. If both source and link laws have dynamics, the scheme is termed ‘primal-dual’.

Let us consider the dual case. The algorithm that will guarantee that these equilibrium prices are Lagrange multipliers, is based on a gradient method:

$$\begin{aligned} \dot{p}_l(t) &= \begin{cases} \frac{y_l - c_l}{c_l} & \text{if } p_l(t) > 0; \\ \max\{0, \frac{y_l - c_l}{c_l}\} & \text{if } p_l(t) = 0. \end{cases} \\ &\triangleq g_l(y_l, c_l) \end{aligned} \quad (7)$$

The choice of these dynamics is key to the proof of scalable functionality as we will see in section 5, and is a result of

the fact that the gradient of the Lagrangian only depends on aggregate rates y_l . If the above equations are at equilibrium, we have $y_l^* = c_l$ for $p_l \neq 0$ and so equilibrium prices are indeed Lagrange multipliers. A similar approach can be used to construct a primal algorithm, as explained in detail in [8].

The routing matrix R is assumed fixed and full row rank. This means that there are no algebraic constraints between link flows, i.e. they can vary independently by choice of source flows x_i . As a consequence, equilibrium prices are uniquely determined.

3. ROBUST MODEL BUILDING

Having identified the framework in which TCP/AQM will be analyzed, one has to develop models for the two modules: the sources and the links. Both consist of software and hardware; inside hardware are uncertain devices connected so as to provide the functionality hardware is meant to have. Hardware response to software commands is therefore assumed instantaneous and is not modelled. On the other hand software should be modelled in some detail, as it defines the response of the routers and the sources in the network. Fine scale, stochastic models for TCP and AQM can be developed, which are usually simplified into deterministic delay differential models that are easier to work with. For example, the first fluid model of Reno/RED was derived from a stochastic model of throughput as a function of loss rate and round-trip delay [16, 20]. From this, a set of deterministic delay differential equations were derived [18] and linearized [6]. Let us outline the steps followed to achieve this.

First, model the change in the window size $dw_i(t)$ according to a Poisson process with a time varying rate $\lambda_i(t)$ modelling the packet loss $N_i(t)$:

$$\begin{aligned} dw_i(t) &= \frac{dt}{\tau_i(t)} - \frac{w_i(t)}{2} dN_i(t), \\ \lambda_i(t) &= \frac{p_i(t - \tau_i)w_i(t - \tau_i)}{\tau_i(t - \tau_i)} \end{aligned}$$

where p_i is the packet marking/dropping probability and w_i/τ_i is the packet sending rate. This takes account of the fact that the ACK received or the loss detected corresponds to a packet which was sent one round trip time (RTT) earlier, and comes delayed by some backward time delay. In going from the stochastic to the deterministic model in [18], the approximation $E[w_i(t)dN_i(t)] \approx E[w_i(t)]E[dN_i(t)]$ was used which was supported both experimentally and theoretically. It is important to stress that no further simplification should be made on the delays as this usually leads to the wrong conclusions, as we will see later. The final nonlinear delay differential equation of TCP becomes

$$\dot{w}_i(t) = \frac{1}{\tau_i(t - \tau_i)} - \frac{w_i(t)w_i(t - \tau_i)}{2\tau_i(t - \tau_i)} q_i(t - \tau_i).$$

At the AQM side, the average queue length $r(t)$ change and instantaneous queue length $b(t)$ for RED are modelled as:

$$\dot{r}(t) = \alpha c(r(t) - b(t)), \quad \dot{b}(t) = \sum_i \frac{w_i(t)}{\tau_i(t)} - c, \quad 0 < \alpha < 1,$$

where c is the link capacity and α is a constant. The aggregate window size at the link is the sum of the window sizes at the sources after a forward trip time. Linearization around equilibrium for N sources with homogeneous delays gives:

$$\dot{w}(t) = -\frac{N}{\tau^*{}^2 c} (w(t) + w(t - \tau^*)) - \frac{\tau^* c^2}{2N^2} p(t - \tau^*) \quad (8)$$

$$\dot{b}(t) = \frac{N}{\tau^*} w(t) - \frac{1}{\tau^*} b(t) \quad (9)$$

$$\dot{p}(t) = -\alpha c(p(t) - \rho b(t)), \quad (10)$$

where τ^* is the round-trip time at equilibrium. ρ is an RED parameter and we assume that RED is operating at the early congestion phase.

Another model for TCP/RED, developed in [15] considers a general multi-source multi-link model, where link prices are considered as marking probabilities at each link and the window size is updated directly according to aggregate link prices at the source. If $w_i(t)$ is the window size of source i at time t , in the congestion avoidance phase of Reno the window update follows

$$\dot{w}_i(t) = x_i(t - \tau_i(t))(1 - q_i(t)) \frac{1}{w_i(t)} - x_i(t - \tau_i(t))q_i(t) \frac{w_i(t)}{2}. \quad (11)$$

The first term on the right corresponds to the fact that the window size $w_i(t)$ is increased by $1/w_i(t)$ by a fraction of $1 - q_i(t)$ when positive ACK is received, delayed appropriately by $\tau_i(t)$. Similarly, the second term models the multiplicative decrease effect of the window size due to packet drop. The model of RED at the AQM side includes modelling the instantaneous queue dynamics $b_l(t)$ of link l at time t , the average queue length $r_l(t)$ and the marking probability $p_l(t)$:

$$\dot{b}_l(t) = y_l(t) - c_l \quad (12)$$

$$\dot{r}_l(t) = -\alpha_l c_l(r_l(t) - b_l(t)) \quad (13)$$

$$p_l(t) = \rho_l(r_l(t) - \underline{b}_l) \quad (14)$$

where $0 < \alpha_l < 1$ is constant, ρ_l, \underline{b}_l are RED parameters. Here the RED marking mechanism is only described at the bottleneck links where marking probability is strictly positive and RED works at its early random dropping phase. One can show that if forward delay is ignored and a single bottleneck used by homogeneous sources is considered, the model in [15] can be simplified to the one in [6]. Equations (8), (11) correspond to the source law while Equations (9–10) and (12–14) represent the link law, shown in Figure 1.

Control theory has developed a methodology for robust model building: all assumptions made during the modelling process are accounted for as some form of uncertainty in the model description. Different types of uncertainty, either of a static form (parametric) or of a dynamic form can be encapsulated in this framework. With this methodology, the ‘details’ are removed from the model yet their significance is accounted for in an uncertain description that is easier to work with. This is done by embedding the model in a *set* of models parameterized in some way by the uncertain parameters. These parameters are used to encapsulate approximations and simplifications made during the modelling process, with the price of some potential conservativeness. Failure to do so

however may lead to the optimistically wrong conclusions, the simplest being the problem of *instability*, although at first it might not seem relevant to congestion control. Instability has a lot of effects. Jittering in the source rate and delays can cause problems in short-lived flows, ‘mice’, that are delay and loss sensitive. *Predictability* of the network behavior in the response at a traffic demand is lost; indeed a more natural response would be a stable convergence to a window size without going into a limit cycle. Also, the average *throughput* may be reduced as a result of the shape of the limit cycle.

EXAMPLE 1. Parametric uncertainty. As a motivating example, consider homogeneous sources sharing a single bottleneck react to the aggregate price q and their current rate x as follows:

$$\dot{x} = \ell - 2q - \frac{3}{2}x + qx \quad (15)$$

where ℓ is a parameter that results from modelling software that is written at the sources end. Routers integrate excess rate, i.e.

$$\dot{p} = \frac{y(t)}{c} - 1 \Rightarrow \dot{q} = \frac{x(t - \tau)}{c} - 1. \quad (16)$$

The value of ℓ is not known exactly, but is believed to be around 2.5, so we fix its value to 2.5. We require the system to be stable for $\tau \in [0, \tau^*)$. Linearization of the above system about the equilibrium $(x^*, q^*) = (c, \frac{1.5c - \ell}{c - 2})$ gives:

$$\begin{aligned} & \frac{d}{dt} \begin{bmatrix} \delta q \\ \delta x \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 \\ c - 2 & \frac{3 - \ell}{c - 2} \end{bmatrix} \begin{bmatrix} \delta q \\ \delta x \end{bmatrix} + \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \delta q(t - \tau) \\ \delta x(t - \tau) \end{bmatrix} \end{aligned}$$

This linearization is stable for $\tau = 0$ for $\ell < 3$ and $c - 2 < 0$ under the additional constraint that the equilibrium is in the positive orthant which requires that $\ell > 1.5c$. The nominal value of $\ell = 2.5$ implies stability for $\tau = 0$, but we have already said that this value was uncertain. Indeed, one can see from the analysis that the system is unstable for $\ell > 3$.

Stability for $\tau = 0$ is therefore a function of c and ℓ . The fact that we fixed $\ell = 2.5$ while we know that ℓ is uncertain might cause problems, as if it were allowed to vary in $2 \leq \ell \leq 3.25$ then the system could go unstable. A way to avoid it is to quantify the uncertainty in ℓ as the model is constructed from the software level. If $2 \leq \ell \leq 2.75$ then stability of the equilibrium can be proven for all possible values of ℓ i.e. the system is *robustly stable*. In some other cases the uncertainty is of a more ‘dynamic’ nature. In this case a different approach has to be followed.

EXAMPLE 2. Dynamic Uncertainty. Suppose that we have the same dynamics for the source as before, but the queue dynamics are different:

$$\begin{aligned} \dot{b} &= x(t - \tau) - c \\ \dot{q} &= -\alpha(qc - b) \\ \dot{x} &= \ell - 2q - \frac{3}{2}x + qx, \end{aligned}$$

where α is an unknown positive parameter. The map from b to q has worst-gain $1/c$ for all times, i.e. we can write:

$$\int_0^T \left\{ q^2(t) - \frac{1}{c^2} b^2(t) \right\} dt \leq 0, \quad \forall T. \quad (17)$$

Equation (17) is what is known as an *Integral Quadratic Constraint (IQC)*, and can be used to encapsulate the dynamics of q into the system description, giving:

$$\begin{aligned} \dot{b} &= x(t - \tau) - c \\ \dot{x} &= \ell - 2q - \frac{3}{2}x + qx \\ 0 &\geq \int_0^T \left\{ q^2(t) - \frac{1}{c^2} b^2(t) \right\} dt. \end{aligned}$$

Tools for investigating the stability in the aforementioned setting have been developed in [17], but can only be applied to the linearized system. The same setting can be used to cover linear time-invariant and time-varying uncertainties, high order dynamics, periodic uncertainties, stiff dynamics, time delays, nonlinearities, random noise, etc.

After a robust model has been constructed, design (or redesign) of some of the parameters can be done in such a way so as to make the system properties better. For example, in the case of TCP/RED various models have been developed to simplify current TCP/AQM protocols [6, 22, 15] all cast under the distributed unified framework developed earlier, shown in Figure 1. These were constructed based on fluid model descriptions that provide a full understanding of equilibrium properties of general large-scale networks under end-to-end control but also facilitate investigation of the dynamical properties of the equilibrium, especially in the presence of feedback delay [22, 23, 30, 11, 29, 1, 3]. Analysis of Equations (8–10) or (11–14) can lead to an appreciation of the deficiencies of TCP/RED as the link capacity and the RTT is increased. It also provides a framework for determining values for the various RED parameters that increase the stability margins of the system [6]. More detailed TCP/RED model analysis reveals that stability is an intrinsic problem which cannot be completely fixed by tuning parameters, and there is a trade-off between stability and response speed. Various controllers, such as P and PI [6] have been proposed. This is a very nice example of how a mathematical analysis of a system allows precise statements to be made about its properties. One might argue that the model that was used is very crude; we have seen how IQCs and parametric uncertainty can help keep track of unmodelled dynamics. Until recently, analysis of uncertain systems was only possible when they are linearized; a framework for nonlinear robust analysis has been developed, which will be presented in section 5.

4. ANALYSIS PROCEDURES AND PITFALLS

After a robust model describing the modules comprising a complex system is built, a systematic analysis procedure is required to evaluate its properties. It has been argued that the simplest model for network congestion control is in terms of a deterministic flow model in the form of a delay-differential equation. Any further simplification to an ordinary differential equation description might be detrimental

in the analysis procedure, and care has to be taken to avoid this. In this section we will identify such and other pitfalls.

Analyzing the stability properties of systems in general is a rather difficult task, in particular when considering nonlinear models. On the other hand, if one considers the linearization of these models about some operating point, then conclusions can be drawn more easily, but they are restricted in some region around the equilibrium.

EXAMPLE 3. Linear versus Nonlinear. Let us consider the system described by Equation (17) whose linearization about the equilibrium was analyzed in Example 1. Figure 2(a) shows what is called the phase-plane for $\ell = 2.5, \tau = 0, c = 1$ for the linearization. We see that under all initial conditions all trajectories tend to the equilibrium. Figure 2(b) shows the full nonlinear system. The thicker line denotes a separatrix, i.e. a trajectory that separates the vector field into two disjoint parts; trajectories on the right of this separatrix can never reach the stable equilibrium. One could be misled into concluding that the equilibrium is globally stable by looking at the linearization, (Figure 2(a)) where all nonlinear effects have been removed and all trajectories tend to the equilibrium, making the local picture global. At this point we should note that for most TCP/AQM algorithms the nonlinear equilibrium point is globally stable, especially if they can be cast in the duality framework of [8]. The above scenario is possible if care is not taken in the design.

In the case of the nonlinear system description, most analysis methods concentrate in the construction of what is called a *Lyapunov function*. Lyapunov functions are nothing but energy-like functions for the system: Their minimum is at the equilibrium, they are positive everywhere else, and their time derivative along the system's trajectories is non-increasing. The task is to construct a Lyapunov function as shown in Figure 2(b), that will provide an *exact* stability proof far away from the equilibrium that linearization can never provide. The concept behind Lyapunov functions is very important. Level curves of Lyapunov functions, shown in Figure 2(b) provide 'trapping regions' that the flow can never get out. Whereas simulation tries to follow the evolution of the system when released from a particular initial condition, the Lyapunov function constructs these nested regions which restrict the flow of the system to within their boundaries. ns-2 simulations can never present the true properties of the system as they can never exhaust all initial conditions and scenarios. The total energy for the system was always a very good candidate for a Lyapunov function, but many systems come from simplifications of higher order systems, and there is no intuition on what the Lyapunov function should look like. We will see in the next section tools for algorithmic analysis of such systems with and without delays.

Delay-differential equations (i.e. models with $\tau \neq 0$) can also be analyzed through linearization and the construction of what is called a Nyquist plot, as in the next example.

EXAMPLE 4. Linear analysis with delay. Suppose $\ell = 2.5$ and $c = 1$ in the model given by Equation (17), and take

Laplace transforms:

$$\frac{d}{dt}\delta q(t) = \delta x(t - \tau) \Rightarrow \delta q(s) = \frac{e^{-s\tau}}{s}\delta x(s) \quad (18)$$

$$\frac{d}{dt}\delta x(t) = -\delta q(t) - \frac{1}{2}\delta x(t) \Rightarrow \delta x(s) = -\frac{1}{(s + 0.5)}\delta q(s). \quad (19)$$

We can consider the frequency domain description as the interconnection of two systems, shown in Figure 2(c), and use a Nyquist argument to test for stability as τ increases [28]. See Figure 2(d) for more details. What we should also point out is that appropriate scaling of the feedback gains by τ allows stability that is independent of delay size, a feature that was used in the development of scalable TCP/AQM algorithms [22].

To analyze network congestion control for arbitrary networks with multiple delays one can use a generalization of Nyquist's criterion [29]. This graphical test allows scalability, but only for the *linearized* versions of the full nonlinear time-delay system. For nonlinear systems with delays it is very difficult to draw conclusions on the stability properties. A Lyapunov argument can also be used but the construction of the relevant Lyapunov functionals is far more difficult than the un-delayed case. Yet designing a congestion control scheme ignoring the effect of delays can be problematic. Let us give an example of the hazard that this can cause.

EXAMPLE 5. The pitfall of ignoring delays. We presented in an earlier section a model for TCP/RED in which the relationship between q , the queue size and p , the marking probability is given by

$$q(s) = \frac{K_m e^{-\tau s}}{(T_1 s + 1)(T_2 s + 1)} p(s)$$

where s is a frequency domain variable, $K_m = \frac{(c\tau)^3}{4N^2}$, $T_1 = \frac{c\tau^2}{2N}$ and $T_2 = \tau$ are TCP parameters, and N the number of active TCP sessions. The factor $e^{-\tau s}$ accounts for the RTT.

A similar simple model can be constructed for RED as

$$p(s) = \frac{1 + k_1 s}{k_2 s} [q_{dem}(s) - q(s)] - K_h q(s)$$

where $q_{dem}(s)$ is the demanded queue size. Choosing the parameters k_1 , k_2 and K_h one can hope to design schemes that tend to stabilize the system and provide better performance. One such attempt was the one followed in [5], which however proceeded in the design by assuming that the term $e^{-s\tau}$ is absent. A simulation for $N = 60$, $\tau = 0.22s$, $c = 1250 \text{ pkts/s}$, $K_h = 0.0014$, $k_1 = 0.4$ and $k_2 = 200$, values that the authors suggest, reveals a major pitfall: the design is unstable when the term $e^{-s\tau}$ is included, see Figure 3. The purpose of redesigning the RED parameters was to achieve a faster response by increasing the gains, but the only reason that this was thought to be possible was the absence of the delay.

We now consider yet another pitfall that may cause problems: It relates to homogeneous versus heterogeneous delays and adding 'exotic' nonlinearities to the system.

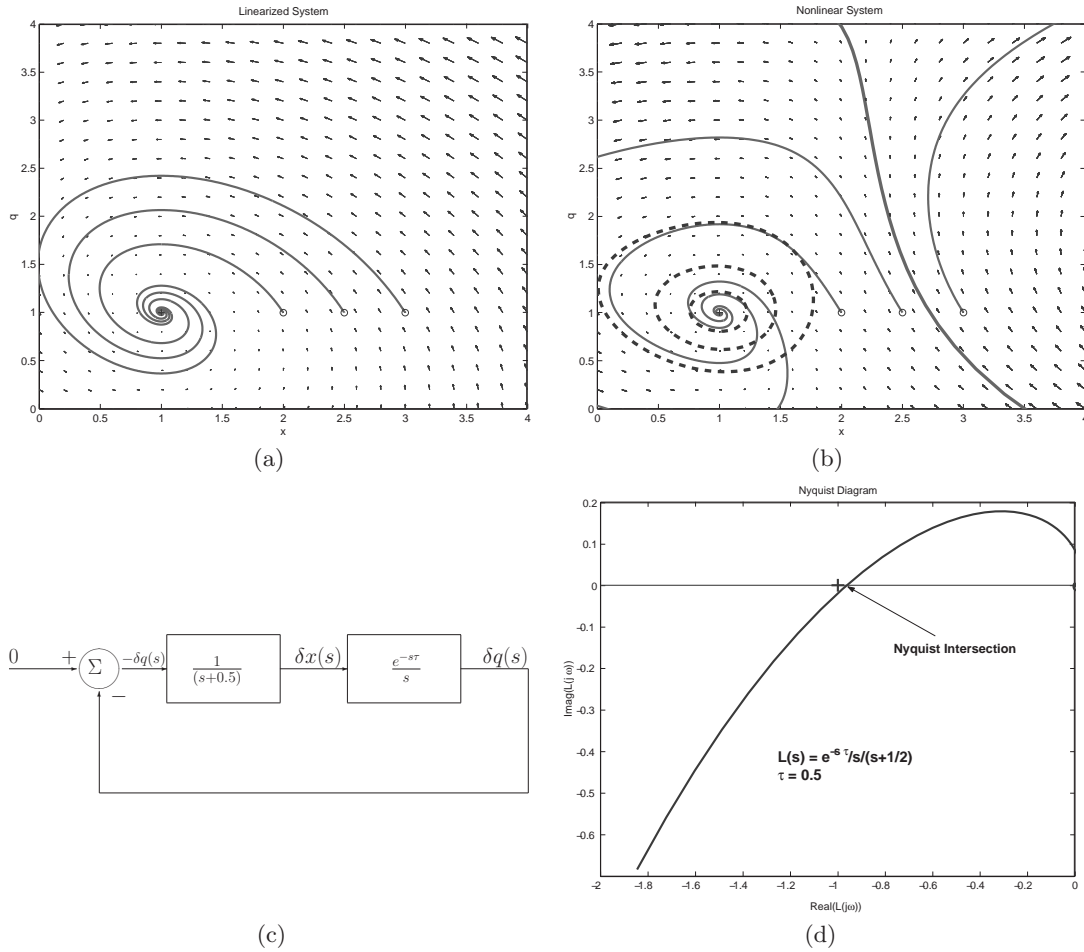


Figure 2: Figures (a,b): Solid lines are trajectories originating from ‘o’, arrows denote vector field, equilibrium is denoted by ‘+’. Figure (a) shows the linearization of the system given by Equation (17) about the equilibrium for $\tau = 0$. Figure (b) shows the phase plane of the full nonlinear system. The level curves of a Lyapunov function that is used to prove stability of the equilibrium are shown dashed in Figure (b). Figures (c,d): Figure (c) shows a block diagram for the system described by Equations (18–19). Figure (d) shows a Nyquist plot, i.e. Real vs Imaginary part of the Loop Transfer Function (the product of the blocks shown in Figure (c)) evaluated at $j\omega$ as ω varies. If this plot encircles the -1 point, the closed loop system is unstable. As τ is increased, this occurs at $\tau^* = 0.5204$. The Nyquist plot is shown for $\tau = 0.5$.

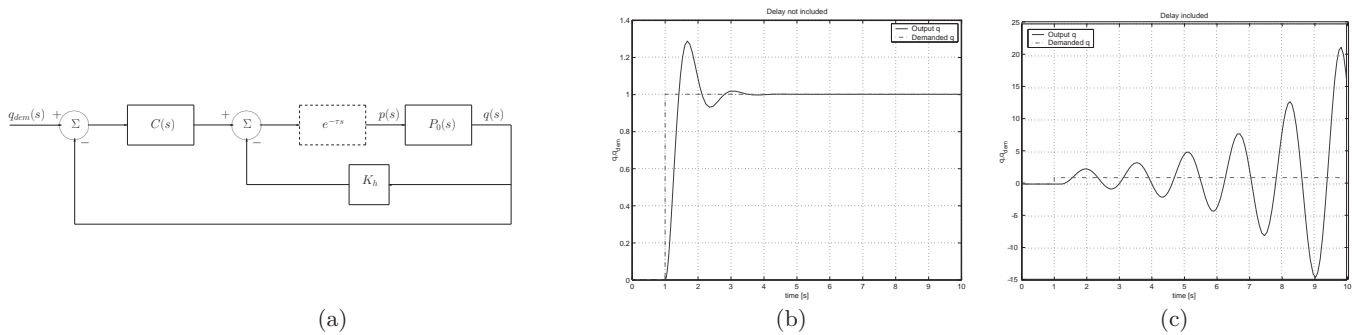


Figure 3: How the omission of delays can lead to a disastrous design. Figure (a) shows a block diagram for the system as it was explained in Example 5. The dashed block corresponds to a delay of τ seconds. If the delay is ignored, the system is stabilized by the control law, as a simulation shows in Figure (b). In Figure (c), the same system is simulated with the delay block introduced, showing the deficiency in the design.

EXAMPLE 6. Homogeneous versus Heterogeneous sources. Consider two sources sharing the same link using FAST [23] as their congestion control scheme. The equations are given by:

$$\begin{aligned}\dot{p} &= \frac{x_{\max,1}}{c} e^{-\frac{\alpha_1}{\tau_1} p(t-\tau_1)} + \frac{x_{\max,2}}{c} e^{-\frac{\alpha_2}{\tau_2} p(t-\tau_2)} - 1 \\ x_1(t) &= x_{\max,1} e^{-\frac{\alpha_1}{\tau_1} p(t-\frac{\tau_1}{2})} \\ x_2(t) &= x_{\max,2} e^{-\frac{\alpha_2}{\tau_2} p(t-\frac{\tau_2}{2})}\end{aligned}$$

where $x_{\max,1}$ and $x_{\max,2}$ are parameters that are chosen so that the equilibrium values $x_1^* = x_2^* = 0.5$ for $p^* = 1$. The linearization of the above system about the equilibrium is

$$\dot{p} = -\frac{\alpha_1}{2\tau_1} p(t-\tau_1) - \frac{\alpha_2}{2\tau_2} p(t-\tau_2)$$

for which the stability condition is $\alpha_i < \frac{\pi}{2}$. We choose $\alpha_1 = \alpha_2 = 0.3$; when the rates are far from their equilibrium values this choice of α_i will cause the system to react very slowly, as shown in Figure 4(a) so we propose to increase the value of α_i according to the distance from the equilibrium, as shown in Figure 4(b). When the two delays are the same, $\tau_1 = \tau_2 = 1$, the new scheme behaves better than the old one. Nonetheless, when $\tau_1 = 0.4$, $\tau_2 = 1$, the system engages in chattering between being ‘too fast’ and ‘too slow’, as shown in Figure 4(c).

This shows that the switching behavior of systems may not be predictable by linearization, and that designing for homogeneous sources may lead to problems for heterogeneous ones.

In the preceding examples, we have seen that the behavior of nonlinear systems far away from the equilibrium cannot be predicted by linearization. In the penultimate section of this paper, we will present recent advances in control theory that can help analyze nonlinear systems directly.

5. ANALYSIS TOOLS FOR NONLINEAR SYSTEMS

The above examples have demonstrated that an important feature of network congestion control is that the feedback mechanism involves delays that cannot be ignored. This important modelling aspect comes with a price, as scalable nonlinear analysis of such systems is quite difficult.

It was mentioned in the previous section that for nonlinear analysis, most methodologies are based on Lyapunov functions, but these are in general difficult to construct. In specific cases, for which the dynamics of the modules comprising the network are judiciously chosen, it is possible to generate scalable proofs for the functionality of the network, i.e. construct Lyapunov functions whose properties scale with the size of the network. However in general such an endeavor would be intractable, in particular for systems which possess dynamics that are not chosen in a way that would enable this.

We now formulate more formally this concept and then proceed with a network example. Consider the equilibrium of interest of the system $\frac{dx}{dt} = f(x)$ to be at the origin, i.e.

$f(0) = 0$. Then a Lyapunov function is given by the following theorem:

THEOREM 7. (Lyapunov) [10] Let $x = 0$ be an equilibrium point of the system $\frac{dx}{dt} = \dot{x} = f(x)$, $x \in \mathbb{R}^n$ (i.e. $f(0) = 0$), and let $D \subset \mathbb{R}^n$ be a domain containing $x = 0$. Let V be a continuously differentiable function defined on D taking values in \mathbb{R} such that:

$$V(0) = 0 \text{ and } V(x) > 0 \text{ in } D - \{0\} \quad (20)$$

$$\frac{dV}{dt} = \dot{V} = \frac{\partial V}{\partial x} f(x) \leq 0 \text{ in } D \quad (21)$$

Then $x = 0$ is stable and V is termed a Lyapunov Function. Moreover, if

$$\dot{V} < 0 \text{ in } D$$

then $x = 0$ is asymptotically stable.

Condition (20) is the positive definiteness condition on $V(x)$ and condition (21) is the negative semidefiniteness of its time derivative.

EXAMPLE 8. Scalable properties using Lyapunov certificates In dual congestion control algorithms the links have the dynamics given by (7). At the sources, we have the following static law:

$$x_i = U_i'^{-1}(q_i). \quad (22)$$

Combining (1,2) and (7, 22) the system has the following closed loop dynamics:

$$\dot{p}_l(t) = \sum_{i=1}^S \frac{R_{li}}{c_l} U_i'^{-1} \left(\sum_{m=1}^L R_{mi} p_m(t - \tau_{i,l}^f - \tau_{i,m}^b) \right) - 1 \quad (23)$$

for $p_l > 0$, and \dot{p}_l is equal to the positive projection of the right hand side of (23) if $p_l = 0$. The undelayed version of this closed loop system is given by

$$\dot{p}_l(t) = \sum_{i=1}^S \frac{R_{li}}{c_l} U_i'^{-1} \left(\sum_{m=1}^L R_{mi} p_m(t) \right) - 1 \quad (24)$$

for $p_l > 0$, and \dot{p}_l is equal to the positive projection of the right hand side of (24) if $p_l = 0$.

THEOREM 9. For fixed full rank R , the (unique) equilibrium of (24) is asymptotically stable for all non-negative initial conditions.

PROOF. Consider the function:

$$V(p) = \sum_{i=1}^L (c_i - y_i^*) p_i + \sum_{i=1}^S \int_{q_i^*}^{q_i} (x_i^* - U_i'^{-1}(Q)) dQ.$$

where $*$ denote equilibrium (optimal) values. This function is positive definite, as argued in [28], and has its minimum at the equilibrium, which is unique under the rank assumption

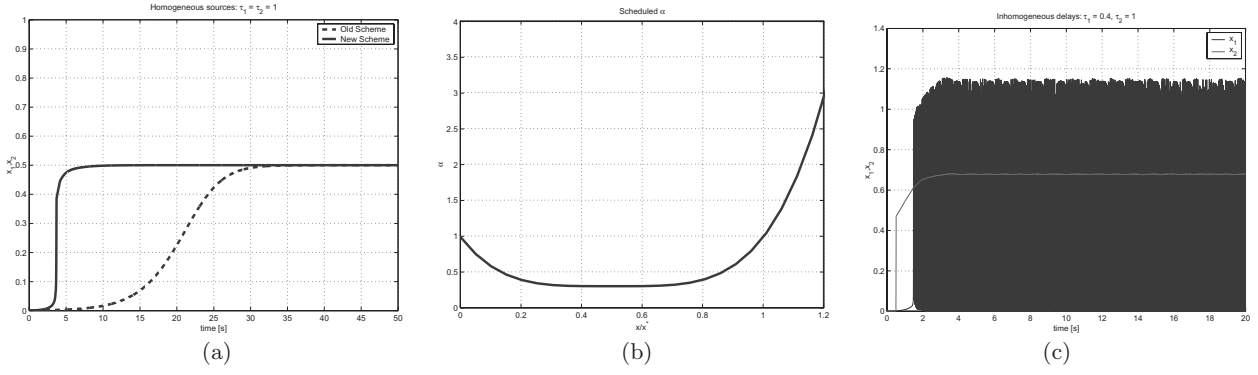


Figure 4: The new proposed design and the effects it has. Figure (a) shows the faster response of the proposed system for two homogeneous sources sharing the same link. Figure (b) shows the scheduling scheme for α_i based on the distance of the rate from the equilibrium one. Figure (c) shows the same law under heterogeneous sources — the system engages in chattering and oscillation.

on R . The time derivative of this Lyapunov function is:

$$\begin{aligned}
 \frac{dV}{dt} &= \sum_{l=1}^L (c_l - y_l^*) \dot{p}_l + \sum_{i=1}^S (x_i^* - U_i'^{-1}(q_i)) \dot{q}_i \\
 &= \sum_{l=1}^L (c_l - y_l^*) \dot{p}_l + \sum_{i=1}^S \sum_{l=1}^L (x_i^* - U_i'^{-1}(q_i)) R_{li} \dot{p}_l \\
 &= \sum_{l=1}^L \left\{ (c_l - \sum_{i=1}^S R_{li} U_i'^{-1} \left(\sum_{m=1}^L R_{mi} p_m \right)) \right\} \dot{p}_l \\
 &= - \sum_{l=1}^L c_l \dot{p}_l^2
 \end{aligned}$$

Now $\dot{V} = 0$ only when each link satisfies $y_l = c_l$ or $y_l < c_l$ and $p_l = 0$. Therefore the optimal solution to the general optimization problem (6) is globally stable if the delays are ignored, under the proposed dynamics at the links and sources. \square

The above theorem presented a scalable proof methodology for establishing the functionality of arbitrary networks, but ignored an important feature that as we saw in the earlier section may be disastrous - the effect of delays. For the linearization of the system with heterogeneous delays included, the use of frequency domain methodologies leads to the following result:

THEOREM 10. [21] Let $M_i = \sum_{l=1}^L R_{li}$. The linearization of the system given by (23) about the (unique) equilibrium is asymptotically stable if

$$\frac{1}{c_l} \sum_{i=1}^S \frac{R_{li} M_i \tau_i}{|U_i''(x_i^*)|} < \frac{\pi}{2}$$

In order to ensure stability for arbitrary topologies for the nonlinear congestion control scheme with delays, we have to use a time-domain methodology, i.e. a Lyapunov based argument. The extension of Lyapunov's theorem for time delayed systems (described by Functional Differential Equations) requires the use of *Lyapunov Functionals*, the so-called Lyapunov-Krasovskii functionals. In this paper we

will construct a Lyapunov-Krasovskii functional that scales with the network size. The system to be analyzed is described by Equation (23).

Recall that the utility function is a continuously differentiable, non-decreasing, strictly concave function. Therefore $U_i''(x_i) < 0$ everywhere. Let γ_i be the lower bound for $|U_i''(x_i)|$, so that

$$|U_i''(x_i)| \geq \gamma_i > 0, \quad \forall i$$

We have the following result:

THEOREM 11. The equilibrium of the system described by (23) is asymptotically stable for arbitrary delays, provided that

$$\frac{1}{c_l} \sum_{i=1}^S \frac{R_{li} M_i \tau_i}{\gamma_i} < \frac{2}{3}$$

and the matrix R is an arbitrary full rank, fixed routing matrix.

PROOF. We will prove that the following system is stable:

$$\dot{p}_l(t) = \sum_{i=1}^S \frac{R_{li}}{c_l} U_i'^{-1} \left(\sum_{m=1}^L R_{mi} p_m(t - \bar{\tau}_i) \right) - 1 \quad (25)$$

where $\bar{\tau}_i = \max_{l,m} (\tau_{i,l}^f + \tau_{i,m}^b)$. Consider V_1 given by

$$V_1(p) = \sum_{l=1}^L (c_l - y_l^*) p_l + \sum_{i=1}^S \int_{\sum_{l=1}^L R_{li} p_l^*}^{\sum_{l=1}^L R_{li} p_l} (x_i^* - U_i'^{-1}(Q)) dQ$$

$V_1 > 0$ apart at the equilibrium, and is radially unbounded [8, 28], as we saw in the proof of the previous theorem. Now

$$\dot{V}_1(p) = - \sum_{l=1}^L c_l \dot{p}_l \dot{p}_{l,u} = - \sum_{l=1}^L c_l \dot{p}_l^2 - \sum_{l=1}^L c_l \dot{p}_l (\dot{p}_{l,u} - \dot{p}_l),$$

where $\dot{p}_{l,u}$ corresponds to the undelayed version of (25),

Equation (24). The second term of \dot{V}_1 is equal to:

$$\begin{aligned}
& - \sum_{l=1}^L c_l \dot{p}_l (\dot{p}_{l,u} - \dot{p}_l) \\
& = - \sum_{l=1}^L \sum_{i=1}^S R_{li} \dot{p}_l \int_{-\bar{\tau}_i}^0 \frac{d}{dt} U_i'^{-1} \left(\sum_{n=1}^L R_{ni} p_n(t+\theta) \right) d\theta \\
& = - \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \int_{-\bar{\tau}_i}^0 \frac{R_{li} R_{mi} \dot{p}_m(t+\theta)}{U_i'' \left(U_i^{-1} \left(\sum_{n=1}^L R_{ni} p_n(t+\theta) \right) \right)} d\theta \\
& \leq \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{li} R_{mi}}{\gamma_i} \int_{-\bar{\tau}_i}^0 |\dot{p}_l| |\dot{p}_m(t+\theta)| d\theta
\end{aligned}$$

Putting everything together we have:

$$\begin{aligned}
\dot{V}_1(p) & \leq - \sum_{l=1}^L c_l \dot{p}_l^2 \\
& + \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{li} R_{mi}}{\gamma_i} \int_{-\bar{\tau}_i}^0 |\dot{p}_l| |\dot{p}_m(t+\theta)| d\theta.
\end{aligned}$$

Now let us concentrate on the second term.

$$\begin{aligned}
& \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{li} R_{mi}}{\gamma_i} \int_{-\bar{\tau}_i}^0 |\dot{p}_l| |\dot{p}_m(t+\theta)| d\theta \\
& \leq \frac{1}{2} \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{mi} R_{li} \bar{\tau}_i}{\gamma_i} \dot{p}_l^2 \\
& + \frac{1}{2} \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{mi} R_{li}}{\gamma_i} \int_{-\bar{\tau}_i}^0 \dot{p}_m^2(t+\theta) d\theta,
\end{aligned}$$

where the formula $2ab \leq a^2 + b^2$ was used. This gives as an estimate for \dot{V}_1 :

$$\begin{aligned}
\dot{V}_1(p) & \leq - \sum_{l=1}^L c_l \dot{p}_l^2 \\
& + \frac{1}{2} \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{mi} R_{li} \bar{\tau}_i}{\gamma_i} \dot{p}_l^2 \\
& + \frac{1}{2} \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{mi} R_{li}}{\gamma_i} \int_{-\bar{\tau}_i}^0 \dot{p}_m^2(t+\theta) d\theta.
\end{aligned}$$

We now introduce another term in the Lyapunov functional,

$$V_2 = \frac{1}{2} \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{mi} R_{li}}{\gamma_i} \int_{-\bar{\tau}_i}^t \int_{t+\theta}^t \dot{p}_m^2(\zeta) d\zeta d\theta.$$

Then

$$\begin{aligned}
\dot{V}_2 & = \frac{1}{2} \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{mi} R_{li} \bar{\tau}_i}{\gamma_i} \dot{p}_l^2 \\
& - \frac{1}{2} \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{mi} R_{li}}{\gamma_i} \int_{-\bar{\tau}_i}^0 \dot{p}_m^2(t+\theta) d\theta
\end{aligned}$$

Let $V = V_1 + V_2$. Then we get

$$\dot{V} \leq - \sum_{l=1}^L c_l \dot{p}_l^2 + \sum_{l=1}^L \sum_{i=1}^S \sum_{m=1}^L \frac{R_{mi} R_{li} \bar{\tau}_i}{\gamma_i} \dot{p}_l^2.$$

The coefficient of each term \dot{p}_l^2 is

$$- c_l + \sum_{i=1}^S \frac{R_{li} M_i \bar{\tau}_i}{\gamma_i}$$

If

$$\frac{1}{c_l} \sum_{i=1}^S \frac{R_{li} M_i \bar{\tau}_i}{\gamma_i} < 1 \quad (26)$$

for all l then $\dot{V} \leq 0$, and the Lyapunov-Krasovskii conditions of stability [4] are satisfied. A simple LaSalle type argument [4] gives asymptotic stability. It is easy to see that $\bar{\tau}_i < \frac{3}{2} \tau_i$, so the above condition becomes:

$$\frac{1}{c_l} \sum_{i=1}^S \frac{R_{li} M_i \tau_i}{\gamma_i} < \frac{2}{3} \quad (27)$$

This completes the proof. \square

A general class of utility functions is [19]

$$U_i(x_i) = \begin{cases} w_i \frac{x_i^{1-\alpha_i}}{1-\alpha_i}, & \alpha_i > 0, \quad \alpha_i \neq 1 \\ w_i \log x_i, & \alpha_i = 1. \end{cases}$$

If $x_i < x_{\max,i}$, then the value of γ is

$$\gamma = \frac{\alpha_i w_i}{x_{\max,i}^{1+\alpha_i}} \quad (28)$$

$U_i''(x_i)$ is non-decreasing, so the value of γ is strictly smaller than the value of $|U_i''(x_i^*)|$, which is part of the conservativeness between the nonlinear result and the linearization.

The above result provides a *scalable* condition for the stability of the whole network. The fact that such a condition could be obtained in the first place is a result of the decentralization scheme and the simple interactions of the TCP and AQM parts of the algorithm, but is also a result of the dynamics that were chosen. In dynamical systems terms, the undelayed closed loop system is a *weighted potential* system, i.e. there is a potential function V so that $\dot{p}_l = -\alpha_l \frac{\partial V}{\partial p_l}$.

In the analysis of general networks and interactions, it may be difficult to construct a scalable result of the form outlined above, as the construction of Lyapunov functions and functionals is not intuitive. Moreover the dynamics of the system that we wish to analyze may come from simplifications and modelling procedures that can be captured by uncertainty as described in the earlier section. The need for tools for systems analysis is evident, and we will now introduce a way to construct these Lyapunov functions and functionals algorithmically given a system model.

Let $u \in \mathbb{R}^m$ denote the parameters, either static or dynamic that appear in the system given by

$$\dot{x} = f_x(x, u),$$

In this case we can construct *Parameterized Lyapunov functions* to answer the more interesting systems analysis questions such as robust functionality, i.e. to test that a property holds for a range of parameters. Their construction is much harder than in the case in which the system does not have any parameters, but if one such Lyapunov function

was available, it would provide a short proof that the system is stable for a range of parameters, something that simulation cannot answer exactly. Both parametric uncertainty and Integral quadratic constraints can be tested using an extension of Lyapunov's theorem. Usually the description of such parameterized systems requires certain constraints to be imposed. These can be of 3 types: Equality, Inequality and Integral Quadratic Constraints. Let us denote them by

$$\begin{aligned} a_{i_1}(x, u) &\leq 0, \quad \text{for } i_1 = 1, \dots, N_1, \\ b_{i_2}(x, u) &= 0, \quad \text{for } i_2 = 1, \dots, N_2, \\ \int_0^T c_{i_3}(x, u) dt &\leq 0, \text{ for } i_3 = 1, \dots, N_3, \text{ and } \forall T \geq 0. \end{aligned}$$

Here $x \in \mathbb{R}^n$ is the state of the system. We assume that the a_{i_1} 's, b_{i_2} 's, and c_{i_3} 's are polynomial functions in (x, u) , and $f_x(x, u)$ is a vector of polynomial functions in (x, u) with no singularity in \mathcal{D} , where $\mathcal{D} \subset \mathbb{R}^{m+n}$ is defined as

$$\mathcal{D} = \{(x, u) \in \mathbb{R}^{m+n} \mid a_{i_1}(x, u) \leq 0, b_{i_2}(x, u) = 0, \forall i_1, i_2\}.$$

Without loss of generality, it is assumed that $f_x(x, u) = 0$ for $x = 0$ and $u \in \mathcal{D}_u^0$, where

$$\mathcal{D}_u^0 = \{u \in \mathbb{R}^m \mid (0, u) \in \mathcal{D}\}.$$

we have the following theorem as an extension of Lyapunov's stability theorem, and which can be used to prove that the origin is a stable equilibrium of the above constrained system.

THEOREM 12. [25] *Suppose that for the above system there exist polynomial functions¹ $V(x)$, $w(x, u)$, $p_{i_1}(x, u)$, $q_{i_2}(x, u)$, and constants $r_{i_3} \geq 0$ such that*

- $V(x)$ is positive definite² in a neighborhood of the origin.
- $w(x, u) > 0$ and $p_{i_1}(x, u) \geq 0$ in \mathcal{D} .

Then

$$\begin{aligned} &-\frac{\partial V}{\partial x} f_x(x, u) + \sum p_{i_1}(x, u) a_{i_1}(x, u) \\ &+ \sum q_{i_2}(x, u) b_{i_2}(x, u) + \sum r_{i_3} c_{i_3}(x, u) \geq 0 \end{aligned}$$

will guarantee that the origin of the state space is a stable equilibrium of the system.

The problem with constructing Lyapunov functions amounts to checking the non-negativity conditions that appear in Theorem 12 efficiently. This is an absolutely crucial element in using Lyapunov methods for nonlinear dynamical systems, but this task is known to be computationally hard, when the order of the polynomial is greater than or equal to 4. In fact, there is no algorithm that will answer the question 'Does this polynomial take only non-negative values when evaluated at every point in its domain?' efficiently. Altering the question to 'Can this polynomial be expressed as a sum of other polynomials squared?', i.e. trying to construct a *Sum of Squares* (SOS) decomposition for it *can* be solved

¹Although not written explicitly here, we assume that we keep track of the indices.

²Strictly speaking, it is enough to require V to have a local minimum at the origin.

efficiently. If a Sum of Squares decomposition is found, this implies that the polynomial is non-negative. The converse is not true: there are polynomials that are non-negative but for which there is no Sum of Squares decomposition.

So how about using the Sum of Squares (SOS) decomposition to check the two Lyapunov conditions? This idea is indeed the step that opened up the way to an algorithmic analysis of nonlinear systems. The framework of Theorem 12 allows parameterized Lyapunov functions to be constructed in the same unified manner for systems that have parametric uncertainty or whose description involves IQCs, the use of which is important for robust modelling as was discussed in section 3. Hybrid and switching systems can also be dealt with directly using the same framework. We can now obtain information about the properties of the system further away from the equilibrium, that no linearization procedure could provide us. Constructing the Lyapunov function as a Sum of Squares polynomial can be done using SOSTOOLS [26], a software written for this purpose. The Lyapunov function shown in Figure 2(b) was constructed algorithmically using SOSTOOLS.

It has been argued that nonlinear deterministic time delay models are the simplest models that one can have for network congestion control schemes. To study the stability of equilibria of time-delay systems, one can use Lyapunov-Krasovskii (L-K) functionals, the natural extension of Lyapunov functions as a tool for stability analysis for ODEs. Can the SOS methodology that was used for nonlinear systems be extended to nonlinear time-delay systems? Indeed this is possible. The functionals that we use have kernels that are *polynomials*. The same methodology can be used to analyze robust stability of nonlinear time delay systems under parametric uncertainty [24].

Consider a time-delay system with a parameter p :

$$\dot{z}(t) = f(z_t, p),$$

where $z_t(\theta) = z(t + \theta)$, $\theta \in [-\tau, 0]$, f completely continuous and $p \in P$ given by

$$P = \{p \in \mathbb{R}^m \mid q_i(p) \geq 0, i = 1, \dots, N\}, \quad (29)$$

i.e. the uncertainty set is captured by a set of inequalities. The condition $f(z_0, p) = 0$ is an equality constraint describing how the equilibrium z_0 moves as p is allowed to change in P .

Let $x(t) = z(t) - z_0$, and transform the system into one with an equilibrium at the origin:

$$\dot{x}(t) = f(x_t + z_0, p) \quad (30)$$

$$0 = f(z_0, p) \quad (31)$$

The stability of this system can be handled directly using the above tools by constructing a *Parameter Dependent* Lyapunov functional. For example we can use:

$$\begin{aligned} V(x_t, p) &= a_0(x(t), p) + \int_{-\tau}^0 a_1(\theta, x(t), x(t + \theta), p) d\theta \\ &+ \int_{-\tau}^0 \int_{t+\theta}^t b_1(x(\zeta), p) d\zeta d\theta. \end{aligned}$$

Then we have the following conditions for stability:

PROPOSITION 13. Consider the system given by Equation (30), where $p \in P$ as defined by Equation (29). Suppose that there exist polynomials $a_0(x(t), p)$, $a_1(\theta, x(t), x(t+\theta), p)$ and $b_1(x(\zeta), p)$ and a positive definite function $\varphi(x(t))$ such that the following conditions hold:

1. $a_0(x(t), p) - \varphi(x(t)) \geq 0, \forall p \in P$,
2. $a_1(\theta, x(t), x(t+\theta), p) \geq 0 \forall \theta \in [-\tau, 0], p \in P$,
3. $b_1(x(\zeta), p) \geq 0 \forall p \in P$,
4. $a_1(0, x(t), x(t), p) - a_1(-\tau, x(t), x(t-\tau), p) + \frac{da_0}{dx(t)} f + \tau b_1(x(t), p) - \tau b_1(x(t+\theta), p) + \tau \frac{\partial a_1}{\partial x(t)} f - \tau \frac{\partial a_1}{\partial \theta} \leq 0, \forall \theta \in [-\tau, 0]$ and $p \in P$ and when Equation (31) is satisfied.

Then the equilibrium 0 of the system given by Equations (30–31) is robustly globally uniformly stable for all $p \in P$ and for all delays in the interval $[0, \tau]$.

To use this proposition, first construct the polynomials a_0 , a_1 and b_1 in SOSTOOLS. The function φ is constructed so as to be positive definite. To impose the conditions $\theta \in [-\tau, 0]$, we use a process similar to the S-procedure and adjoin the condition $\theta(\theta + \tau) \leq 0$ to the positivity and negativity conditions. The equality constraints given by Equation (31) that may arise during the transformation process can be adjoined using appropriate polynomial multipliers [25]. More details can be found in [24].

Then the four Sum of Squares conditions corresponding to the four nonnegativity conditions in Proposition 13 will be four constraints in a relevant Sum of Squares programme which can be solved using SOSTOOLS [26].

EXAMPLE 14. **Nonlinear analysis with delay.** Consider the system that was developed in Example 1, given by Equations (15–16), for $c = 1$ and $\ell = 2.5$. We can construct a Lyapunov functional of the form (32) for $|q-1| < 0.6$ and $|x_i - 1| < 0.6$ for $\tau = 0.48$. Recall that the linearization showed stability for $\tau = 0.5204$.

We will now analyze the stability properties of a congestion control scheme that is described in [23]. Although the vector field is not polynomial, a change of variables is suggested that renders it polynomial. Consider the network given by Figure 5. The routing matrix R in this case is equal to:

$$R = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$$

The network dynamics are given by

$$\dot{p}_i(t) = \sum_{i=1}^S \frac{R_{li}}{C_l} x_{\max, i} e^{\frac{-\alpha_i \sum_{m=1}^L R_{mi} p_m(t - \tau_{i,l}^f - \tau_{i,m}^b)}{M_i \tau_i}} - 1, \quad (32)$$

for $p_l > 0$, and \dot{p}_i is equal to the positive projection of the right hand side of (32) if $p_i = 0$. Here M_i is an upper bound on the number of bottleneck links that source i sees in its path, α_i are source gains, and $x_{\max, i}$ are source constants.




Figure 5: A simple network.

We set $\alpha_1 = \alpha_2 = \alpha_3 = \alpha$, $M_1 = M_2 = 1$, $M_3 = 2$. To impose fairness, we enforce $x_i^* = 1$ so that $x_{\max, 1}^{\tau_1} x_{\max, 2}^{\tau_2} = x_{\max, 3}^{2\tau_3}$ and $y_{01} = y_{02} = c_1 = c_2 = 2$. The system equations then become:

$$\begin{aligned} \dot{p}_1 &= \frac{1}{2} (x_{\max, 1} e^{\frac{-\alpha p_1(t - \tau_1)}{\tau_1}} \\ &\quad + x_{\max, 3} e^{\frac{-\alpha(p_1(t - \bar{\tau}_3) + p_2(t - \bar{\tau}_3))}{2\tau_3}} - 2) \\ \dot{p}_2 &= \frac{1}{2} (x_{\max, 2} e^{\frac{-\alpha p_2(t - \tau_2)}{\tau_2}} \\ &\quad + x_{\max, 3} e^{\frac{-\alpha(p_1(t - \bar{\tau}_3) + p_2(t - \bar{\tau}_3))}{2\tau_3}} - 2) \end{aligned}$$

where $\bar{\tau}_3 = \max_{l,m} \tau_{3,l}^f + \tau_{3,m}^b$ for $l, m = 1, 2$. We now perform the following change of variables:

$$z_1 = x_{\max, 1}^{\frac{\tau_1}{2\tau_3}} e^{\frac{-\alpha p_1(t)}{2\tau_3}} - 1, \quad z_2 = x_{\max, 2}^{\frac{\tau_2}{2\tau_3}} e^{\frac{-\alpha p_2(t)}{2\tau_3}} - 1$$

to get

$$\begin{aligned} \dot{z}_1 &= -\frac{\alpha}{4\tau_3} [z_1(t) + 1] [(z_1(t - \tau_1) + 1)^{\frac{2\tau_3}{\tau_1}} \\ &\quad + (z_1(t - \bar{\tau}_3) + 1)(z_2(t - \bar{\tau}_3) + 1) - 2] \\ \dot{z}_2 &= -\frac{\alpha}{4\tau_3} [z_2(t) + 1] [(z_2(t - \tau_1) + 1)^{\frac{2\tau_3}{\tau_2}} \\ &\quad + (z_1(t - \bar{\tau}_3) + 1)(z_2(t - \bar{\tau}_3) + 1) - 2] \end{aligned}$$

Consider in the case in which $2\tau_1 = 2\tau_2 = \tau_3 = 2\tau$. For $\alpha = 1$ we can construct a Lyapunov functional that proves stability of the equilibrium.

We note that Equation (32) is a special case of (23) with

$$U_i(x_i) = \frac{\tau_i M_i}{\alpha_i} x_i \left(1 - \log \frac{x_i}{x_{\max, i}} \right)$$

We can now compare the result obtained in Theorem 11 with the SOS based result. We see that the condition in Theorem 11 is more conservative; it gives the following two conditions:

$$\begin{aligned} \alpha_1 x_{\max, 1} + \alpha_3 x_{\max, 3} &< \frac{4}{3} \\ \alpha_2 x_{\max, 2} + \alpha_3 x_{\max, 3} &< \frac{4}{3} \end{aligned}$$

Therefore the best α_i for which stability can be proven satisfies $\alpha_i < \frac{2}{3}$ (when $x_i^* = x_{\max, i} = 1$). The Lyapunov functional we constructed using the Sum of Squares decomposition holds locally, for fixed ratios of inhomogeneous time-delays, but for $\alpha_i = 1$. Therefore we see the advantages and disadvantages of the two methodologies. Although conservative, the result of Theorem 11 is what ensures the func-

tionality of the system for arbitrary network sizes; the conservativeness can be reduced if the topology of the network is known exactly.

6. CONCLUSIONS

In this paper, we made the thesis that optimization based decompositions of complex systems into interacting modules facilitates analysis, comparability and verifiability of the desired system properties. The modularity that such decompositions offer, and which at first endows the systems with an apparent complexity should be taken advantage of. Aiming for such decompositions is beneficial both for analysis and design.

We also stressed that in any analysis procedure it is important to construct *robust* models for the modules, as this will capture the uncertainty in modelling and component parameters so that it be taken account in the design process. The new tools that we developed in the previous section allow us to analyze such systems even at the nonlinear level, and expand the applicability of this methodology.

Complex systems and large scale networks will dominate the future societies as technology advances. Designing such systems is more than art based on intuition. It is widely appreciated that network congestion control for the Internet is probably the only complex system for which we have a good understanding of the interaction of the various modules at the TCP/AQM level. The system can be designed by resorting to a solid methodological framework that provides the desired functionality at equilibrium, based on an optimization scheme; and the correct dynamics can be chosen for the various modules to drive the system to the equilibrium — the right choice of dynamics are key to the scalability of the verification result.

The success in designing network congestion control schemes for the Internet through a mathematical formulation which enables understanding of its functionality and the limitations that features such as delays pose, allows us to believe that similar hierarchical structures can enable understanding and design of other complex systems in the future. We envision that solid methodological frameworks can be used to formulate and solve the design problem this way and the resulting system's functionality can be proven in a structured way. Apart from the specific analysis results that one can produce by hand, the algorithmic procedure we propose can be used to analyze more complicated system descriptions therefore increasing the set of model building blocks that can be used in the construction of future mathematical frameworks for complex system analysis.

7. REFERENCES

- [1] H. Choe and S. H. Low. Stabilized Vegas. In *Proc. of IEEE Infocom*, 2003.
- [2] G. E. Dullerud and F. Paganini. *A Course in Robust Control Theory: A Convex Approach*. Springer-Verlag, 2000.
- [3] X. Fan, M. Arcak, and J. Wen. Robustness of network flow control against disturbances and time-delay. To appear in *System and Control Letters*, 2004.
- [4] J. K. Hale and S. M. V. Lunel. *Introduction to Functional Differential Equations*. Applied Mathematical Sciences (99). Springer-Verlag, 1993.
- [5] Z. Heying, L. Baohong, and D. Wenhua. Design of a robust active queue management algorithm based on feedback compensation. In *ACM Sigcomm*, 2003.
- [6] C. V. Hollot, V. Misra, D. Towsley, and W. Gong. A control theoretic analysis of red. In *Proc. of IEEE Infocom*, 2001.
- [7] R. Jain. Congestion control in computer networks: Issues and trends. *IEEE Network Magazine*, pages 24–30, May 1990.
- [8] 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.
- [9] S. Keshav. A control-theoretic approach to flow control. In *ACM Sigcomm*, 1991.
- [10] H. K. Khalil. *Nonlinear Systems*. Prentice Hall, Inc., second edition, 1996.
- [11] S. Kunniyur and R. Srikant. A time-scale decomposition approach to adaptive ECN marking. *IEEE Transactions on Automatic Control*, 47(6):882–894, 2002.
- [12] S. H. Low. A duality model of TCP and queue management algorithms. *IEEE/ACM Transactions on Networking*, 11(4):525–536, 2003.
- [13] S. H. Low and D. E. Lapsley. Optimization flow control, i: basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 1999.
- [14] S. H. Low, F. Paganini, and J. C. Doyle. Internet congestion control. *IEEE Control Systems Magazine*, 2002.
- [15] S. H. Low, F. Paganini, J. Wang, and J. C. Doyle. Linear stability of TCP/RED and a scalable control. *Computer Networks Journal*, 43:633–647, 2003.
- [16] M. Mathis, J. Semske, J. Mahdavi, and T. Ott. The macroscopic behavior of the TCP congestion avoidance algorithm. *Computer Communication Review*, 27(3):67–82, 1997.
- [17] A. Megretski and A. Rantzer. System analysis via integral quadratic constraints. *IEEE Transactions on Automatic Control*, 42(6):819–830, 1997.
- [18] V. Misra, W.-B. Gong, and D. Towsley. Fluid-based analysis of a network of AQM routers supporting TCP flows with an application to RED. In *ACM Sigcomm*, 2000.
- [19] J. Mo and J. Warland. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, 2000.
- [20] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling TCP throughput: A simple model and its empirical validation. In *ACM Sigcomm*, 2001.

- [21] F. Paganini, J. Doyle, and S. Low. Scalable laws for stable network congestion control. In *Proceedings of the 40th IEEE Conf. on Decision and Control, Orlando, FL*, 2001.
- [22] F. Paganini, J. C. Doyle, and S. H. Low. Scalable laws for stable network congestion control. In *Proc. of the 40th Conference on Decision and Control*, 2001.
- [23] F. Paganini, Z. Wang, J. C. Doyle, and S. H. Low. Congestion control for high performance, stability and fairness in general networks. To appear in *IEEE/ACM Transactions on Networking*, 2004.
- [24] A. Papachristodoulou. Analysis of nonlinear time delay systems using the sum of squares decomposition. In *Proc. of the American Control Conference*, 2004.
- [25] A. Papachristodoulou and S. Prajna. On the construction of Lyapunov functions using the sum of squares decomposition. In *Proc. of the 41st IEEE Conference on Decision and Control*, 2002.
- [26] S. Prajna, A. Papachristodoulou, and P. Parrilo. Introducing SOSTOOLS: A general purpose sum of squares programming solver. In *Proc. of the 41st IEEE Conference on Decision and Control*, 2002.
- [27] S. Shenker. A theoretical analysis of feedback flow control. In *ACM Sigcomm*, 1990.
- [28] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhäuser, 2003.
- [29] G. Vinnicombe. Robust congestion control for the Internet. submitted for publication.
- [30] G. Vinnicombe. On the stability of networks operating TCP-like congestion control. In *Proc. of IFAC World Congress*, 2002.