Flow Control via Pricing: a Feedback Perspective

Fernando Paganini*
UCLA Electrical Engineering
Los Angeles, CA 90095-1594.
paganini@ee.ucla.edu

Abstract

This paper discusses network flow control from the perspective of multivariable feedback theory. We build on recent research on flow control based on pricing signals, which has cast these problems in terms amenable to mathematical analysis, and more specifically we focus on the control algorithm of [8]. After reviewing our recent result regarding the stability of this method, we move on to issues of performance and delay robustness. These are studied with tools of linear loop-shaping, which exhibit the main tradeoffs involved in parameter selection for these algorithms. An illustrative example is provided.

1 Introduction

Feedback mechanisms are of widespread use for controlling flows in networks such as the Internet; indeed some form of feedback is inevitable in order for resources to be efficiently used in the presence of high variability and uncertainty. However, feedback theory is rarely invoked for such designs: at first sight, practical congestion control mechanisms as those in TCP and its variants [3, 2] appear far removed from mathematically grounded control.

Recent research (see [1, 4, 5, 7, 8] and references therein) in flow control based on propagation of price signals has, however, strikingly closed this gap and opened the door for a mathematical theory of network flow control. These models can be used both to interpret the mechanism implicit in existing protocols [4, 9], and to propose alternatives based on more explicit price signaling. In particular, Kelly *et al.* [4, 5] have developed a framework with suggestive economic theory interpretations that leads to two dual families of flow control algorithms. Employing continuous time models, analytical results are obtained regarding equilibrium fairness, stability, and convergence rates for these systems, as well as studies of stochastic disturbances and time delays.

A related approach has been developed by Low and co-workers [7, 8], based on discrete-time models. In it, prices are interpreted as Lagrange multipliers for the optimization of total source *utility*, subject to link capacity constraints. This interpretation leads to convergence proofs for a first-order version of the price update law [7]. A second-order version was later proposed in [8] to address some drawbacks of the first order

*Research supported by NSF CAREER Award ECS-9875056, and the David and Lucille Packard Foundation. The author thanks Steven Low for his guidance in this topic.

algorithm. This alternative has satisfactory performance in simulations, but analytical results have been more difficult to obtain. Very recently [11], employing a continuous-time model and a Lyapunov argument, we have been able to prove stability for this case. An outline of this result is provided in Section 3 of this paper.

The main thrust of the present paper is, however, to move beyond stability and apply to this problem the perspective of the multivariable feedback designer. Here, as in other applications, the main concern is the fundamental tradeoff between dynamic *performance* (transient response, tracking capability, etc.) and *robustness* to uncertain or under-modeled effects (e.g. time delays). In this regard, we show in Section 4, working with linearized models, that a loopshaping point of view can provide insight on the design tradeoffs involved in parameter choices for the algorithm of [8]. Section 5 contains a simulation example that illustrates these findings in the nonlinear setting.

2 Problem Formulation

We are concerned with a system of L communication links shared by a set of K sources. The routing matrix R, of dimensions $L \times K$, is defined by

$$R_{lk} = \begin{cases} 1 & \text{if source } k \text{ uses link } l \\ 0 & \text{otherwise} \end{cases} .$$

- **For each link l we have:**
  - the capacity c_l;
  - the aggregate rate y_l of all flows through the link;
  - the backlog b_l;
  - the price signal p_l.
- **For each source k we have:**
  - The source rate x_k;
  - The aggregate price q_k of all links used by source k.

We use vector notation to collect the above variables across all links or sources; thus we define $c, y, b, p \in \mathbb{R}^L$, and $x, q \in \mathbb{R}^K$. The following relationships are immediate[1]:

$$y = Rx, \qquad q = R^T p. \tag{1}$$

It is assumed that links are able to measure their own aggregate rate, and that sources are fed back their aggregate price; see [5, 8] for discussions on implementation of such mechanisms. For the moment we also assume that propagation of rates and prices is instantaneous; we later discuss the effect of delay.

With the information given so far, we have depicted in Figure 1 a generic feedback structure for flow control based on price propagation; what remains to be specified is:

- How the links fix their prices based on link utilization.
- How the sources fix their rates based on their aggregate price.

[1] R^T is the matrix transpose of R.

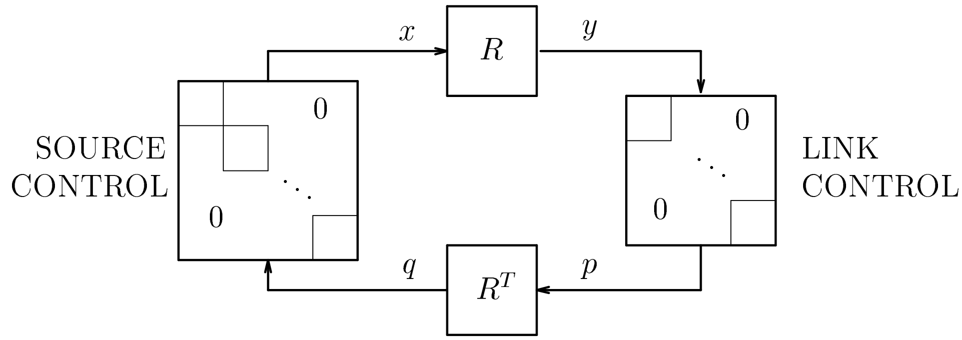

Figure 1: General flow control structure based on pricing signals.

These operations are up to the designer, but have a main restriction: both must be *decentralized*, as indicated in the figure by a block diagonal structure. For instance the source rate x_k can only depend on the corresponding aggregate price q_k. We are thus in the realm of decentralized feedback control. The objective of this feedback is to allow for flows to dynamically adapt to changing conditions in traffic demand, link capacity, routing, etc. Key design considerations are thus to ensure dynamic stability and regulation of these systems around equilibria that satisfy some desirable static properties (full resource utilization, fairness among sources).

In this context, Kelly *et al.* [5] have proposed two flow control methods:

(i) The "primal" algorithm uses a first-order continuous-time model for the source rate control, and a static law for the price update.

(ii) The "dual" method uses a first-order dynamics for the price update, and a static law for source rate selection.

We will not describe the specific laws here, but indicate that [5] describes fairness properties of the resulting equilibria, and proves that these are global attractors by means of a Lyapunov argument. These equilibria can also be interpreted in economic terms as balancing network supply with source demand, expressed in terms of a *utility* function.

In this paper we will focus on the approach of Low and coworkers [7, 8], which is similar to the dual method above, in particular it is based on a static law for source control, and a dynamic price control law at the links.

Source rate control. According to [7, 8], for a given total price q_k, the sources pick the rate that maximizes

$$\mathcal{U}_k(x_k) - x_k q_k$$

over x_k, where $\mathcal{U}_k(x_k)$ is the source utility function, assumed to be strictly concave[2]. Assuming $\mathcal{U}_k(\cdot)$ is differentiable, the maximum is achieved at $x_k = \mathcal{U}_k'^{-1}(q_k)$, where $\mathcal{U}_k'^{-1}$ is the inverse function of the derivative of $\mathcal{U}_k$. We denote henceforth

$$f_k(q_k) := \mathcal{U}_k'^{-1}(q_k).$$

Notice that $\mathcal{U}_k'$ is strictly decreasing in $x_k > 0$, hence f_k is a strictly monotone decreasing function of q_k. In vector notation, we summarize the above equations for source rates as

$$x = f(q). \tag{2}$$

[2][7] allows for the inclusion of maximum and minimum constraints for x_k, for simplicity we will not impose those here (as, for instance, in logarithmic utility functions $\mathcal{U}_k(x_k) = w_k \log(x_k)$).

Link price control: We adopt a continuous time version of the dynamics from [8]; for each l, backlogs and prices evolve according to

$$\frac{db_l}{dt} = \begin{cases} (y_l - c_l) & \text{if } b_l(t) > 0; \\ \max\{0, (y_l - c_l)\} & \text{if } b_l(t) = 0. \end{cases} \tag{3}$$

$$\frac{dp_l}{dt} = \begin{cases} \gamma(\alpha_l b_l + y_l - c_l) & \text{if } p_l(t) > 0; \\ \max\{0, \gamma(\alpha_l b_l + y_l - c_l)\} & \text{if } p_l(t) = 0. \end{cases} \tag{4}$$

The above switched differential equations enforce non-negativity constraints in all variables. Here $\gamma > 0$ and $\alpha_l > 0$ are constants. Setting $\alpha_l = 0$ corresponds to the first-order price update law in [7]; in this case prices are proportional to backlogs, which may lead to an undesirable equilibrium backlog, hence the use of $\alpha_l > 0$ in [8].

Let (b^*, p^*) be an equilibrium of the above system. We also use the notation $q^* = R^T p^*$ for the equilibrium source prices, $x^* = f(q^*)$ for the equilibrium source rates, and $y^* = Rx^*$ for the equilibrium link rates. It is not difficult to see that we must have $b^* = 0$. Indeed, if $b_l^* > 0$ then we would have $y_l^* = c_l$ so $\dot{p}_l > 0$, which contradicts equilibrium. Now p^* need not be zero, indeed its nonzero components correspond to links where $y_l \equiv c_l$, i.e. where the capacity constraint is active (bottleneck links).

The relation of this algorithm with optimization is explained in detail in [7], but we briefly outline it here. The key observation is that an equilibrium point p^*, x^* satisfying the equations (1-4) will be a saddle point of the optimization

$$\min_{p \geq 0} \max_{x} \left(\sum_{k=1}^{K} \mathcal{U}_k(x_k) + p^T(c - Rx) \right).$$

This is the Lagrangian dual of the convex program of maximizing the overall utility $\sum_{k=1}^{K} \mathcal{U}_k(x_k)$, subject to link capacity constraints $Rx \leq c$.. Furthermore, the first-order link price algorithm of [7], corresponding to $\alpha_l = 0$, can be interpreted as a gradient algorithm for the above maximization in p; this is used to establish convergence. No such interpretation is available in the second-order case of [8] we are considering.

It follows from duality theory that x^* must be the unique global optimum of the latter problem; therefore y^*, q^* are also unique. p^* need not be unique, because in general the capacity constraints might not be independent. To simplify the further development and obtain a unique equilibrium price, we assume from now on that the matrix R is of full row rank. Thus, given a vector q of aggregate source prices, there is a unique p satisfying $q = R^T p$. Some comments on generalizing this assumption are given in [11].

3 Stability

In this section we outline a proof of global stability for the price update law of (3-4); this supplements available results for the systems in [5, 7]. We state the following result:

Theorem 1. *Given the system (1-4), assume $f_k(q_k)$ is strictly decreasing in $q_k > 0$, and that R is of full row rank. Then the unique equilibrium point $b^* = 0$, p^* is globally asymptotically stable.*

A detailed proof is given in [11]. Here we will briefly outline the argument, which is based on the Lyapunov function

$$V(b, p) = \sum_{l=1}^{L} [\alpha_l \gamma \frac{b_l^2}{2} + (c_l - y_l^*) p_l] + \sum_{k=1}^{K} \phi_k(q_k),$$

$$\text{where} \quad \phi_k(q_k) := \int_{q_k^*}^{q_k} (x_k^* - f_k(\sigma))d\sigma.$$

The monotonicity of f_k makes ϕ_k non-negative, which implies $V(b,p)$ is non-negative (notice $y_l^* \leq c_l$). It is further argued in [11] that under the rank assumption on R, $V(b,p)$ only vanishes at equilibrium and is radially unbounded. The derivative of $V(b,p)$ along trajectories of the system is given by

$$\dot{V} = \sum_{l=1}^{L}[\alpha_l \gamma b_l \dot{b}_l + (c_l - y_l^*)\dot{p}_l] + \sum_{k=1}^{K}(x_k^* - f_k(q_k))\dot{q}_k.$$

Now observe that last term above can be written as

$$\sum_{k=1}^{K}(x_k^* - x_k)\dot{q}_k = (x^* - x)^T \dot{q} = (x^* - x)^T R^T \dot{p} = (y^* - y)^T \dot{p} = \sum_{l=1}^{L}(y_l^* - y_l)\dot{p}_l, \tag{5}$$

and thus

$$\dot{V} = \sum_{l=1}^{L}[\alpha_l \gamma b_l \dot{b}_l + (c_l - y_l^*)\dot{p}_l + (y_l^* - y_l)\dot{p}_l] = \sum_{l=1}^{L} \nu_l,$$

where we have denoted $\nu_l := \alpha_l \gamma b_l \dot{b}_l + (c_l - y_l)\dot{p}_l$.

Consider the case where $b_l > 0$, $p_l > 0$; then the dynamics (3-4) give

$$\nu_l = \alpha_l \gamma b_l (y_l - c_l) + (c_l - y_l)\gamma(\alpha_l b_l + y_l - c_l) = -\gamma(y_l - c_l)^2 \leq 0.$$

The limiting cases $b_l = 0$, $p_l = 0$ are more delicate, but a detailed study in [11] shows that indeed $\nu_l \leq 0$ at all times. Therefore our Lyapunov function is decreasing along trajectories of the system. Finalizing the proof of global asymptotic stability involves ensuring that $\dot{V} \equiv 0$ can only occur at equilibrium, and invoking Lasalle's invariance principle (see [6]) to show trajectories must converge to equilibrium over time. For full details on this last step see [11].

This theorem implies, under mild assumptions, that the flow control method proposed in [8] is sound and will, in the continuum limit, converge asymptotically to a desirable equilibrium where the overall utility is maximized, and backlogs are cleared.

4 Performance and Robustness

Global stability is a fundamental property, and it is fortunate to be able to establish it with such generality; however it does not suffice to guarantee a feedback system behaves in a satisfactory way. As in other feedback design problems, key questions are:

(i) What is the dynamic performance of this system? There are many ways to specify this, but mainly one wants to know that in addition to "eventually" settling down to equilibrium, the system is capable of quickly reacting to changing conditions. Of relevance here are the transient response from non-equilibrium initial conditions (these arise due to changes in network users, links, or routing), and the ability to track variations in link capacity c (e.g. in an Available Bit Rate scenario).

(ii) What is the robustness of this system to deviations from the idealized model? There are again many such deviations, many of which can be modeled as stochastic disturbances around equilibrium. From a stability point of view, however, the key concern is the effect of system *delay*, which arises both in propagating prices and rates and also is inherent in discrete-time price update implementations.

Some of these issues have been studied in [5] for the pricing algorithms considered there. Here we will explore some of these questions for the algorithm in [8], and cast them in the language of multivariable feedback control theory. We will work with a *linearized* model of the dynamics around the equilibrium. Let $\tilde{x}$, $\tilde{y}$, etc, denote the deviations of the system variables from their equilibrium values x^*, y^*, and so on. We also allow fluctuations $\tilde{c}$ in the link capacities. Then:

- Linearizing the source rate function $x = f(q)$ around the point (x^*, q^*), we write $\tilde{x} \approx -D\tilde{q}$, where
$$D = -\mathrm{diag}(f_k'(q_k^*)) = -\mathrm{diag}\left(\frac{1}{\mathcal{U}_k''(x_k^*)}\right).$$
Here and in the sequel, $\mathrm{diag}(\cdots)$ denotes a diagonal matrix with the enclosed diagonal elements. We have used above the fact that f_k is $(\mathcal{U}_k')^{-1}$. The sign convention is for convenience, making D a positive definite matrix (since $\mathcal{U}_k$ is concave).

- We consider the linear approximation
$$\frac{d\tilde{b}_l}{dt} = (\tilde{y}_l - \tilde{c}_l) \tag{6}$$
$$\frac{d\tilde{p}_l}{dt} = \gamma(\alpha_l \tilde{b}_l + \tilde{y}_l - \tilde{c}_l) \tag{7}$$

to the price/backlog dynamics. Clearly, (7) is a valid local model for bottleneck links, where the equilibrium $p_l^* > 0$. The argument is less clear for (6), since the equilibrium of (3) occurs at the point of discontinuity; one case in which it would apply is if the b_l are taken to measure the deviation from a nonzero backlog b_{l0}, as is also considered in [8]. Given these limitations, as well as the inherent locality of any linear model, our analysis to follow does not carry the same weight as the earlier stability result. Nevertheless we will see below that this first-cut linear model can have significant predictive power.

With this model, we can represent the linearized dynamics by the block diagram on the top of Figure 2; in it, s is the usual Laplace variable, and for simplicity we have assumed that the α_l have the same value α across all links.

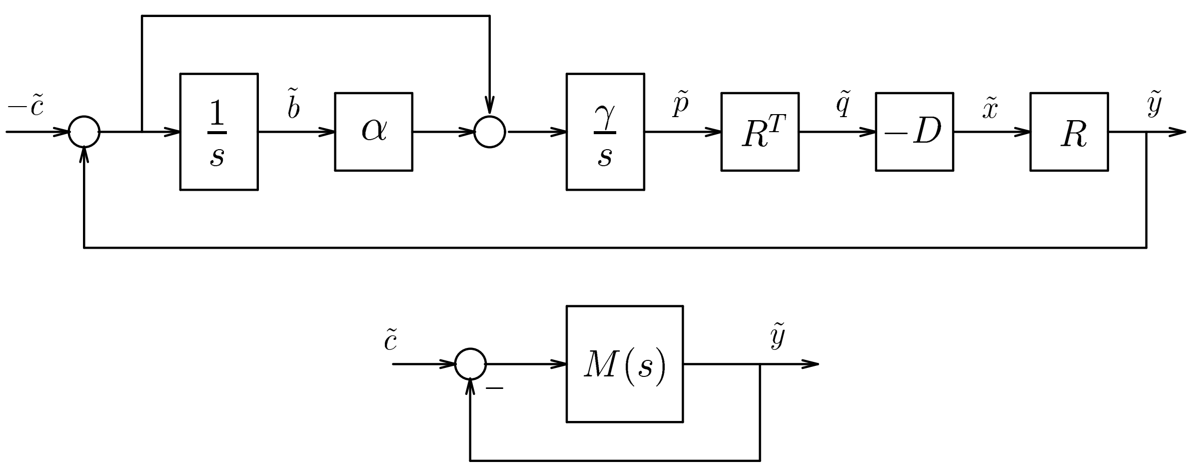

Figure 2: Linearized model and standard unity feedback form

The bottom figure is obtained from the above by absorbing the negative sign and defining the loop gain transfer function matrix
$$M(s) = \frac{\gamma}{s}\left(1 + \frac{\alpha}{s}\right) RDR^T.$$

In this way we are led to a standard unity feedback configuration; we can now bring in standard tools from multivariable control, to analyze:

1. The transient response. In particular, the closed loop modes are given by the roots of $\det(I + M(s)) = 0$.

2. The tracking performance: i.e. how well the link utilization tracks link capacity variations. Here we can look at the singular values of the sensitivity function
$$S(j\omega) = (I + M(j\omega))^{-1}$$
which maps $\tilde{c}$ to tracking error, and study the range of ω for which $S(j\omega)$ is small.

3. Stability margin to delay. One can insert delays of the form $e^{-\tau_i s}$ to loop channels and study bounds on delays consistent with stability.

These tasks are greatly simplified by the observation that the loop transfer function is easily diagonalized: indeed the matrix RDR^T is symmetric, positive definite and hence can be written as $U^T \Lambda U$, where U is a unitary matrix and Λ is the diagonal matrix of positive eigenvalues $\lambda_1, \ldots \lambda_L$. Given this fact, we have
$$M(s) = U^T \mathrm{diag}\Big(m_1(s), \ldots, m_L(s)\Big) U,$$
with $m_i(s) = \frac{\gamma\lambda_i}{s}\left(1 + \frac{\alpha}{s}\right)$. This means that the multivariable feedback can be easily studied by L single loops of loop-gain $m_i(s)$, leading to the following conclusions:

1. The closed loop modes are the roots of $s^2 + (\gamma\lambda_i)s + \gamma\lambda_i\alpha = 0$, for $i = 1 \ldots L$. It is easily verified that the dominant pole of the above equation can have negative real part no faster than $\gamma\lambda_i/2$; therefore the decay exponent of the overall transient dynamics is limited by $\gamma\lambda_{\min}/2$, where $\lambda_{\min}$ is the smallest eigenvalue of RDR^T.

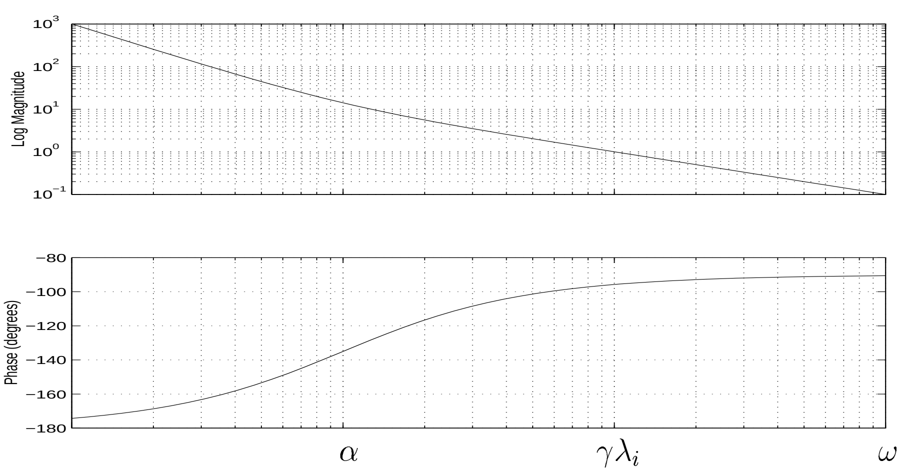

Figure 3: Bode plot of $m_i(j\omega)$

2. The sensitivity function diagonalizes as $S(j\omega) = U \mathrm{diag}\Big((1 + m_i(j\omega))^{-1}\Big) U^T$, so its singular values are determined by the "loopshape" functions $m_i(j\omega)$. In Figure 3 we sketch its Bode plot for the case $\alpha << \gamma\lambda_i$. The frequency range for which $|m_i(j\omega)| >> 1$ is associated with good tracking performance. Since the cross-over frequency is about $\gamma\lambda_i$ in this case, once again we find the figure of merit $\gamma\lambda_{\min}$ as a measure of the bandwidth over which our system exhibits good tracking.

3. Finally, we turn to the issue of robustness to delay. The analysis is simplest for a common delay $e^{-\tau s}$ in all loops; this directly translates to a modified loopshape function $m_i(j\omega)e^{-\tau j\omega}$; for this to yield stability, in classical terms we must have a *phase margin* greater than $\omega_c\tau$ in the original Bode plot, where ω_c is the crossover frequency.

In fact the best phase margin is about 90 degrees, achieved when $\alpha << \gamma\lambda_i$, which motivates our earlier choice. Then $\omega_c \approx \gamma\lambda_i$ and the stability condition is approximately

$$\tau < \pi/(2\gamma\lambda_i).$$

Here the condition is more stringent for the *largest* eigenvalue $\lambda_{\max}$ of RDR^T, and is *inversely* proportional to γ.

We summarize the conclusions of our linear analysis:

- The parameter γ directly affects dynamic performance, but inversely affects delay robustness; a judicious design must balance this tradeoff.
- A similar role is played by the matrix RDR^T, which represents the sensitivity of link rates to price increases. Note, however, that while its *minimum* eigenvalue dictates performance, its *maximum* eigenvalue impacts delay stability. Hence, for an ill-conditioned matrix RDR^T the design tradeoff will become more difficult; this observation is well known in multivariable feedback design problems (see, e.g. [10]).
- For the parameter α, robustness to delay indicates the use of $\alpha << \gamma\lambda_i$.

5 Example

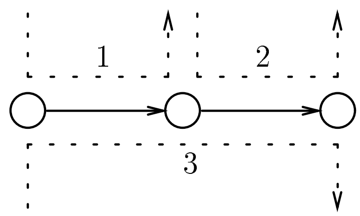

Figure 4: Example: network of 2 links, 3 sources

Consider the simple 2-link network of Figure 4; we assign capacity $c_1 = 0.9$ to the left link, and $c_2 = 2.4$ to the right link. Three sources, each with utility function $\mathcal{U}(x) = \log(x)$, share the network as indicated by dashed lines. Thus the routing matrix is

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

It is easily verified that the equilibrium point for the price-based control is given by

$$p^* = \begin{bmatrix} 2 \\ 0.5 \end{bmatrix}, \quad q^* = \begin{bmatrix} 2 \\ 0.5 \\ 2.5 \end{bmatrix}, \quad x^* = \begin{bmatrix} 0.5 \\ 2 \\ 0.4 \end{bmatrix}, \quad y^* = c = \begin{bmatrix} 0.9 \\ 2.4 \end{bmatrix}. \tag{8}$$

The linearized matrix

$$D = -\mathrm{diag}\left(\frac{1}{U''(x_1^*)}, \frac{1}{U''(x_2^*)}, \frac{1}{U''(x_3^*)}\right) = \begin{bmatrix} 0.25 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 0.16 \end{bmatrix}$$

leads to $RDR^T = \begin{bmatrix} 0.41 & 0.16 \\ 0.16 & 4.16 \end{bmatrix}$, with eigenvalues $\lambda_{\min} = 0.4$, $\lambda_{\max} = 4.16$.

We see that we are in an ill-conditioned case, with $\lambda_{\max}/\lambda_{\min} \approx 10$.

5.1 Study of Tracking Performance

We consider the price dynamics of (3-4) with parameters $\gamma = 0.1$, $\alpha = 0.01$. Note that $\gamma\lambda_{\min} = 0.04$, so we are in the situation of Figure 3, and in particular, the cross-over frequency of $m_1(j\omega)$ is $\omega_{c1} \approx 0.04$. Our linear analysis predicts that variations in capacity will be tracked well up to this bandwidth, but not beyond.

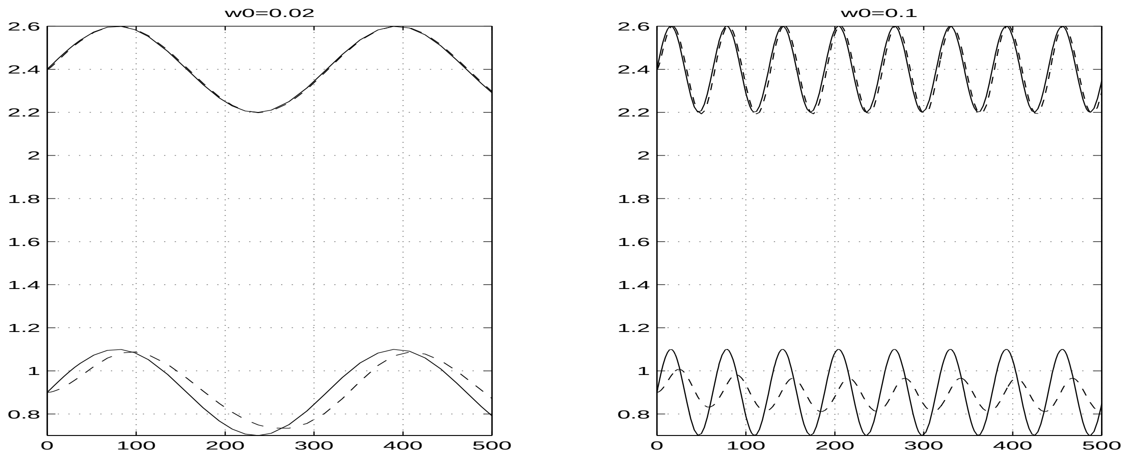

Figure 5: Link capacities (solid) and flows (dashed). Left: $\omega_0 = 0.02$; right: $\omega_0 = 0.1$

We simulated the nonlinear dynamics (1-4) using the Matlab stiff solver 'ode15s', and introducing the sinusoidal variations $\tilde{c}_l(t) = 0.2\sin(\omega_0 t)$, to the capacity values from (8), for several values of ω_0. Figure 5 presents the results for $\omega_0 = 0.02$ and $\omega_0 = 0.1$. We confirm the prediction that at least one of the link flows loses the ability to track the network capacity as frequency grows beyond the critical value of $\omega_0 = 0.04$. The fact that the other flow is able to track closely corresponds here to the matrix RDR^T being close to diagonal; in a general situation, many or all links could lose their tracking capability.

5.2 Study of Robustness to Delay

In a second study, we restored the capacity to the constant value in (8), but now inserted a delay in the feedback loop. Turning again to our linear analysis, we expect that

$$\tau < \pi/(2\gamma\lambda_{\max}) = 3.77$$

will be required for stability. Figure 6 shows simulations for various values of the delay.

From the top plot we see that delays of $\tau = 2$ or $\tau = 3$ affect the response, but retain stability; once we move to $\tau = 4$, however, the equilibrium for the second price becomes unstable and the nonlinear system exhibits a limit cycle.

6 Conclusion

We have explored the application of tools from feedback theory to flow control algorithms based on pricing. Our investigation with simple linear models reveals basic design tradeoffs involved when selecting control parameters, e.g. the price update parameter γ in [8]. An interesting finding is that an ill-conditioned matrix RDR^T can severely affect these dynamic tradeoffs. This factor should thus be considered, in addition to other considerations such as fairness, in decisions on routing and utility function selection.

We have focused on the control law from [8], but in principle similar analysis is possible for other methods with the general structure of Figure 1. A natural question

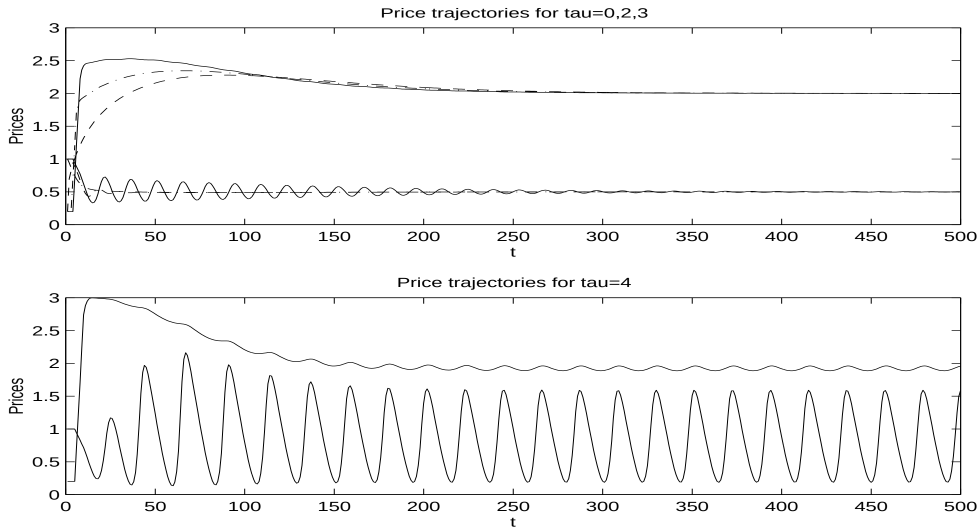

Figure 6: Price trajectories for various delay values. Top: $\tau = 0, 2, 3$; bottom: $\tau = 4$.

is whether control theory can design better algorithms than those already proposed; the answer is not obvious due to the decentralization requirement, which is difficult to handle in control synthesis methods; pursuing it will be a topic of future research.

References

[1] C. Courcoubetis, V. Siris, and G.D.Stamoulis, "Integration of Pricing and Flow Control for Available Bit Rate Services in ATM Networks", *Proc. Globecom '96.*

[2] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance", *IEEE/ACM Trans. on Networking*, 1(4):397-413, Aug. 1993.

[3] V. Jacobson, "Congestion avoidance and control", *Proc. ACM SIGCOMM '88.*

[4] F. P. Kelly, "Mathematical modeling of the Internet", Fourth International Congress on Industrial and Applied Mathematics, Edinburgh, Scotland, July 1999.

[5] F. P. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: Shadow prices, proportional fairness, and stability". *Jour. Oper. Res. Soc.*, 49(3), pp. 237-252, 1998.

[6] H.K. Khalil. *Nonlinear Systems.* Prentice-Hall, 1996.

[7] S. H. Low and D. E. Lapsley, "Optimization flow control – I: basic algorithm and convergence" *IEEE/ACM Trans. on Networking*, Vol 7, No 6. Dec 1999.

[8] S. Athuraliya and S. H. Low, "Optimization flow control – II: random exponential marking" *preprint*, http://www.ee.mu.oz.au/staf/slow/research, May 2000.

[9] S. Low, L. Peterson and L. Wang, "Understanding TCP Vegas: Theory and Practice", *preprint*, http://www.ee.mu.oz.au/staf/slow/research, Feb. 2000.

[10] M. Morari and E. Zafiriou. *Robust Process Control.* Prentice Hall, 1989.

[11] F. Paganini, "On the stability of optimization-based flow control", submitted to the 2001 American Control Conference. http://www.ee.ucla.edu/∼paganini/