

Scalable Laws for Stable Network Congestion Control

Fernando Paganini¹ John Doyle² Steven Low³

Abstract

This paper discusses flow control in networks, in which sources control their rates based on feedback signals received from the network links, a feature present in current TCP protocols. We develop a congestion control system which is arbitrarily scalable, in the sense that its stability is maintained for arbitrary network topologies and arbitrary amounts of delay. Such a system can be implemented in a decentralized way with information currently available in networks plus a small amount of additional signaling.

1 Introduction

Flow control in a communication network like the Internet concerns the adjustment of individual source transmission rates so that network resources are fully utilized, and link capacities are not exceeded. Existing protocols such as TCP regulate their transmission rates by a *congestion window* of outstanding packets; feedback signals from the network (packet acknowledgment or loss) are used to dynamically adjust this window to match available capacity. The additive increase multiplicative decrease (AIMD) algorithm of TCP Tahoe [6] and its enhancements (Reno, etc.) has performed remarkably well as the network scales up several orders of magnitude. However future growth combined with the upcoming diversity of traffic and communication substrates will stretch the limits of this simple algorithm, calling for a more extensive analytical investigation of the problem.

Exciting progress has been made very recently towards a theoretical understanding of TCP, both of its equilibrium, using optimization theory, and its dynamics, using control theory. Key to these developments has been to explicitly model the feedback signal generated by links and communicated to sources; in practice these signals correspond to packet dropping, or packet marking when Explicit Congestion Notification (ECN) bits are available. Interpreting these signals as *prices* has allowed for

the rate assignment problem to be cast in supply/demand terms [8]; in a related interpretation [13], prices correspond to Lagrange multipliers of a certain optimization problem. When applied to the existing variants of TCP, this framework allows for the analysis of the resulting equilibria [12].

In terms of the *dynamics* of TCP, very recently an analytical model has been developed [18] that correlates well with standard simulators, and allows for stability studies [5]. It is found that TCP exhibits oscillatory instabilities as delays grow, and perhaps more surprisingly, as link capacities grow. This motivates the investigation of new protocols where one could provide guarantees of stability, robustness to delay and scalability to large capacities and arbitrary topologies.

Some work in recent years coming from the control community (e.g. [1, 19, 16]) has investigated both classical and modern tools in these problems, however results are mainly confined to single bottleneck networks. For arbitrary networks, the work on price signals has allowed for global stability results [8, 13, 20, 10], but only when ignoring delays. When considering delays, one is forced to slow down the control gains to retain stability (see, e.g. [8, 13, 21]). As noted in [9], this mechanism is already implicit in window-based protocols and it is therefore conceivable that one could obtain stability for arbitrary delays. In this vein, recent work in [7, 17, 22] has derived conditions for delay robustness for the protocols in [8].

In this paper we pose the objective of finding a protocol which can be implemented in a decentralized way by sources and routers, and satisfies some basic objectives: efficient use of network resources in equilibrium, and local stability for arbitrary capacities, delays, and routing. These requirements, laid out in Section 3, lead us to adopt some conditions for the linearized dynamics: integration at the links, and certain scalings on the gain at sources and links. In Section 4 we then prove that a system with these properties will satisfy our requirements, combining ideas in [7] with methods of multivariable robust control. In Section 5 we show how to implement the source control laws by nonlinear functions which around equilibrium provide the correct dynamic linearization, and in Section 6 we give some comments on packet-level implementation. Conclusions are given in Section 7.

¹UCLA Electrical Engineering, Los Angeles, CA 90095-1594. Corresponding author. Email: paganini@ee.ucla.edu.

²Control and Dynamical Systems, Caltech, Pasadena, CA 91125.

³Computer Science & Electrical Engineering, Caltech, Pasadena, CA 91125.

2 Problem Formulation

We are concerned with a system of L communication links shared by a set of S sources. For each link l we have: the capacity c_l , the aggregate rate y_l of all flows through the link, and the price signal p_l . For each source i we have the source rate x_i , and the aggregate price q_i of all links used by source i . We will work with *flow* models, i.e. regard these variables as deterministic, non-negative real-valued. We use vector notation to collect the above variables across all links or sources; thus we define $c, y, p \in \mathbb{R}^L$, and $x, q \in \mathbb{R}^S$.

We now model the network, including the propagation delays, in terms of the relationships (in the Laplace domain, T denotes transpose)

$$y(s) = R_f(s)x(s), \quad (1)$$

$$q(s) = R_b(s)^T p(s). \quad (2)$$

Here R_f and R_b are the delayed forward and backward routing matrices, defined by

$$[R_f(s)]_{l,i} = \begin{cases} e^{-\tau_{i,l}^f s} & \text{if source } i \text{ uses link } l \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

and similarly for $R_b(s)$ with $\tau_{i,l}^f$ replaced by $\tau_{i,l}^b$. Thus, link flows are obtained by aggregation, with respective delays $\tau_{i,l}^f$, of source flows using the link; prices seen by sources are aggregations, with delays $\tau_{i,l}^b$, of link prices used by the source. We define the total round trip time (RTT) by source,

$$\tau_i = \tau_{i,l}^b + \tau_{i,l}^f. \quad (4)$$

This quantity is available to sources in real time. What remains to be specified is: (i) How the links fix their prices based on link utilization; (ii) how the sources fix their rates based on their aggregate price. These operations are up to the designer, but have a main restriction: both must be *decentralized*. For instance the source rate x_i can only depend on the corresponding aggregate price q_i . The objective of this feedback is to allow for flows to track 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 desirable static properties.

This paper is mainly concerned with the dynamic properties around a given equilibrium point x_0, y_0, p_0, q_0 . For more discussion on equilibria see Section 5, but for now we only require that $y_{0l} \leq c_l$ (link capacities are not exceeded), and that non-bottleneck links (those where $y_{0l} < c_l$) have a price $p_{0l} = 0$.

Now consider a small perturbation around equilibrium; $x = x_0 + \delta x$, $y = y_0 + \delta y$, $p = p_0 + \delta p$,

$q = q_0 + \delta q$. Assuming the set of bottlenecks is unchanged by this small perturbation, δp_l is only non-zero for bottleneck links. Therefore for the local analysis to follow, we can write the reduced model

$$\delta \bar{y}(s) = \bar{R}_f(s) \delta x(s), \quad (5)$$

$$\delta q(s) = \bar{R}_b(s)^T \delta \bar{p}(s), \quad (6)$$

where the matrices \bar{R}_f , \bar{R}_b , and the vectors $\delta \bar{p}$, $\delta \bar{y}$ are obtained by eliminating the rows corresponding to non-bottleneck links. For simplicity in the sequel, we drop the “bar” notation and assume the routing matrices such as (3) refer only to bottleneck links.

A subtlety arises when employing the linear equation (5) for incremental flows: a bottleneck link includes effectively a saturation nonlinearity in its outgoing flow, preventing an increase δx in one of its source flows from propagating to downstream links. We eliminate this issue by assuming that the target rate c_l for each link is slightly lower than its actual capacity (a “virtual” capacity, see [11]). This is also desirable since it leads to empty equilibrium queues.

We will also assume that the matrix $R = R_f(0) = R_b(0)$ is of full row rank. This means that there are no algebraic constraints between bottleneck link flows, ruling out, for instance, the situation where all flows through one link also go through another. Note, however, that our R matrices refer now to bottlenecks only, and typically in the above scenario only one of the links would be a bottleneck; so our assumption is quite generic.

3 Control objectives and a proposed law

We now proceed to lay out a series of objectives for the feedback control laws in purely local (linearized) terms. These will lead us to conjecture a candidate local control law, which we argue is the simplest that can satisfy our requirements, and at the same time allows for the required decentralized implementation. In the next section we will prove that it actually achieves our objectives.

A first objective is that the target capacity c_l is matched at equilibrium; this calls for an integrator in the feedback loop. As will become clear in the next section, to have stability we must perform the integration at the lower-dimensional end of the problem, i.e. at the links. So we write

$$\delta \dot{p}_l = \beta_l \delta y_l$$

as in [13], i.e. prices integrating excess capacity. The constant β_l will be chosen later.

The next, main objective is to have local dynamic stability for arbitrary network delays, link capacities and routing topologies.

Regarding delays, consider first the case of a single link and source. The link integrator, plus the network delay will yield a term $e^{-\tau s}/s$ in the loop transfer function, which leads to instability as τ grows, unless the gain is made a function of τ . Indeed, introducing a gain $1/\tau$ in the loop (specifically at the source, which measures RTT), gives a loop transfer function

$$\frac{e^{-\tau s}}{\tau s} \quad (7)$$

which is scale-invariant: namely, its frequency response is a function of $\theta = \tau\omega$, so Nyquist plots for all values of τ would fall on a single curve Γ , depicted below. If the overall gain is kept under control, then we can have stability for all τ . Further, closed loop time responses for different τ 's would be the same except for scale.




Figure 1: Nyquist plot of $e^{j\theta}/j\theta$

Going now to the general network case, we will extend the above by including a gain $1/\tau_i$ at each source. In the multivariable setting, however, we must also compensate for the effect of the routing matrices R_f , R_b ; intuitively, as more links participate in the feedback the gain must be appropriately scaled down. The difficulty is implementing this in a decentralized fashion, without access to the global routing information.

One source of information that can be exploited is that at equilibrium, the aggregate source rates add up to capacity, so $Rx_0 = c$. This motivates the following heuristic: introduce (i) a gain $\beta_l = \frac{1}{c_l}$ at each link; and (ii) a gain x_{0i} at each source, in addition to the $1/\tau_i$ factor. In the case of a single link, and possibly many sources, this gives a loop transfer function of

$$L(j\omega) = \sum_i \frac{x_{0i}}{c_l} \frac{e^{-j\tau_i\omega}}{\tau_i\omega},$$

which is a convex combination of points of the form (7); looking at the figure, it follows that this convex combination will remain stable by a Nyquist argument. In the multiple link case, we are still left with the need of bounding the gain of the backward path R_b ; an analogous strategy is to introduce a gain $\frac{1}{M_i}$ at each source, M_i being the number of bottleneck links in the source's path. For a discussion on its implementation, see Section 5.

Summarizing the above requirements in their simplest possible form, we propose the source control law (between δq_i and δx_i) to be the static gain

$$-\kappa_i := -\frac{\alpha_i x_{0i}}{M_i \tau_i}, \quad (8)$$

where α_i is a parameter. The sign provides negative feedback. As link control, we take an integrator with gain normalized by capacity,

$$\delta p_l = \frac{1}{c_l s} \delta y_l. \quad (9)$$

With this normalization, the price corresponds to a (virtual) queueing delay at the link, similar to the case of TCP Vegas [3, 14].

4 Linear stability results




Figure 2: Overall feedback loop

With the source and link controllers described above, we proceed to study the linearized stability of the closed loop. We express the overall, multivariable feedback loop in the classical configuration of Figure 2, with open loop transfer function

$$L(s) = R_f(s) \mathcal{K} R_b^T(s) C \frac{I_L}{s}, \quad (10)$$

where I_L is the identity matrix of size L , and

$$\mathcal{K} = \text{diag}(\kappa_i), \quad C = \text{diag}\left(\frac{1}{c_l}\right).$$

Note that there are no unstable pole/zero cancellations within $L(s)$; the proposition below provides stability conditions for such multivariable loops with integral control.

Proposition 1 Consider a standard unity feedback loop, with $L(s) = \gamma F(s) \frac{I}{s}$. Suppose:

- (i) $F(s)$ is analytic in $\text{Re}(s) > 0$, and $\|F(s)\| \leq \beta$ in $\text{Re}(s) \geq 0$.
- (ii) $F(0)$ has strictly positive eigenvalues.
- (iii) For all $\gamma \in (0, 1]$, -1 is not an eigenvalue of $L(j\omega)$, $\omega \neq 0$.

Then the closed loop is stable for all $\gamma \in (0, 1]$.

In essence, the above conditions are a “nominal” stability requirement for small γ , that says that we have strictly negative feedback of enough rank to stabilize all the integrators, and a “robustness” argument that says we can perform a homotopy to

$\gamma = 1$ without bifurcating into instability. Details of the proof are omitted for brevity.

Applying this to the $L(s)$ in (10), we take

$$F(s) = R_f(s) \mathcal{K} R_b^T(s) \mathcal{C};$$

we will later add the scaling γ . Note that (i) is automatically satisfied. For (ii), note that

$$\text{eig}(F(0)) = \text{eig}(\mathcal{C}^{\frac{1}{2}} R \mathcal{K} R^T \mathcal{C}^{\frac{1}{2}}),$$

where $R = R_f(0) = R_b(0)$ is the static routing matrix. Assuming R has full row rank, condition (ii) will hold. Here we see the importance of putting the integrators at the links (the lower dimensional portion). If, instead, we tried to integrate at the sources, the resulting feedback matrix at DC would never have enough rank to stabilize the (larger) number of integrators.

What remains is to establish (iii). For this purpose we bring in some more notation:

$$X_0 = \text{diag}(x_{0i}), \quad \mathcal{M} = \text{diag}\left(\frac{1}{M_i}\right),$$

$$\Lambda(s) = \text{diag}(\lambda_i(s)), \quad \lambda_i(s) = \frac{\alpha_i e^{-\tau_i s}}{\tau_i s}.$$

Also, $R_b(s) = R_f(-s) \text{diag}(e^{-\tau_i s})$ follows from (4), so

$$R_b^T(s) = \text{diag}(e^{-\tau_i s}) R_f^T(s),$$

where $R_f^T(s) = R_f^T(-s)$ is the adjoint system. We incorporate this notation to rewrite $L(s)$, for $s \neq 0$, in the more convenient form

$$L(s) = R_f(s) X_0 \mathcal{M} \Lambda(s) R_f^T(s) \mathcal{C}. \quad (11)$$

We now tackle the robustness argument.

Theorem 2 Consider an equilibrium point where rates match capacity, i.e. $c = R_f(0)x_0$. Let $\alpha_i < \frac{\pi}{2}$ and the delays be arbitrary. Then with $L(s)$ as in (11), $-1 \notin \text{eig}(L(j\omega))$, $\omega \neq 0$.

Proof: Since nonzero eigenvalues are invariant under commutation, and also many of the factors in (11) are diagonal, we observe that

$$-1 \in \text{eig}(L(j\omega)) \iff -1 \in \text{eig}(P(j\omega) \Lambda(j\omega)),$$

$$P(j\omega) := \mathcal{M}^{\frac{1}{2}} X_0^{\frac{1}{2}} R_f(j\omega)^* \mathcal{C} R_f(j\omega) X_0^{\frac{1}{2}} \mathcal{M}^{\frac{1}{2}} \geq 0.$$

Claim:

$$0 \leq P \leq I. \quad (12)$$

This amounts to bounding the spectral radius

$$\rho(P) = \rho(\mathcal{M} R_f(j\omega)^* \mathcal{C} R_f(j\omega) X_0)$$

$$\leq \|\mathcal{M} R_f(j\omega)^*\| \cdot \|\mathcal{C} R_f(j\omega) X_0\|.$$

Any induced norm will do, but if we use the l_∞ -induced (max-row-sum) norm, we find that

$$\|\mathcal{C} R_f(j\omega) X_0\|_{\infty\text{-ind}} = \max_i \frac{1}{c_i} \sum_{i \text{ uses } l} |e^{-\tau_i l j\omega} x_{0i}|$$

$$= \max_i \frac{1}{c_i} \sum_{i \text{ uses } l} x_{0i} = 1;$$

note we are dealing with bottlenecks. Also the norm $\|\mathcal{M} R_f^*\|$ is equal to 1, because each row contains exactly M_i elements of magnitude $1/M_i$. So $\rho(P) \leq 1$ as claimed. Indeed, $\rho(P) = 1$ at $\omega = 0$, the eigenvector being the vector of all ones.

Now suppose $-1 \in \text{eig}(P(j\omega) \Lambda(j\omega))$ for some ω . We thus have a vector u , $|u| = 1$ such that

$$y = \Lambda u, \quad u = -P y.$$

Now

$$u^* y = u^* \Lambda u = \sum_i \lambda_i |u_i|^2$$

is a convex combination of the $\{\lambda_i\}$, which are points in the curve Γ of Figure 1, scaled by $\alpha_i < \frac{\pi}{2}$. It is clear from the figure that such convex combinations and scaling cannot reach any point in the half-line $(-\infty, -1]$. However, we also have

$$1 + u^* y = u^* u + u^* y = y^* P (P - I) y \leq 0,$$

using (12). So $u^* y \in (-\infty, -1]$, a contradiction. ■

Remark 1 Some elements of the proof, in particular the use of l_∞ induced norms to prove a spectral radius bound, are inspired by the work of [7] for the control laws in [8]. More recently [22] has extended the stability argument for the laws in [8] in a parallel fashion to our work.

Theorem 2 establishes (iii) in Proposition 1; note that scaling down by γ is equivalent to making the α_i smaller. To summarize, we have:

Theorem 3 Let $R_f(s)$, $R_b(s)$ denote the routing matrices of sources in relation to the bottleneck links. Suppose $R_f(0) = R_b(0)$ has full row rank, and that $\alpha_i < \frac{\pi}{2}$. Then the system with source controller gain (8) and link controller (9) is linearly stable for arbitrary delays and link capacities.

We have modeled a law consisting only of the integrator and delay dynamics, but otherwise with only gains at links and sources. Could this stability argument extend to include additional dynamics?

In this regard, a first comment is that there can be no more pure integrators. Otherwise the Nyquist plot in Figure 1 would branch towards $-\infty$, and convex combinations of such points could reach the critical point. In particular, the second order link laws used in REM [2] would not qualify. Given this, and the requirement of scale-invariance to delay, all additional dynamic terms should be located at the sources, so that they can be ‘‘clocked’’ at the rate of the round-trip time. For instance a term $\frac{1}{\tau_i s + \beta_i}$ could be added to the source controller; restrictions on β_i would be required.

5 Implementation and the equilibrium structure

We have presented a linearized control law with some desirable stability properties. We now discuss how to embed such linear control laws in a global, nonlinear control scheme whose equilibrium would linearize as required.

For control at the links, since the gains are constant it is straightforward to implement the price dynamics as

$$\dot{p}_l = \begin{cases} \frac{y_l - c_l}{c_l} & \text{if } p_l > 0; \\ \max\{0, \frac{y_l - c_l}{c_l}\} & \text{if } p_l(t) = 0. \end{cases} \quad (13)$$

Therefore prices integrate excess capacity in a normalized way, and are saturated to be always non-negative. At equilibrium, bottlenecks with nonzero price will have $y_{l0} = c_l$ as required. Non-bottlenecks with $y_{l0} < c_l$ will have zero price.

For the sources, the problem is somewhat more involved since the linearized gains depend on the equilibrium rate x_{0i} , and also depend on parameters τ_i and M_i . Let us assume those parameters are constant for the moment, and focus on the mapping between q_i and x_i . The simplest thing to try is a static, strictly decreasing function $x_i = f_i(q_i)$.

Indeed this choice of source and link laws makes this control a special case of the class of algorithms proposed in [13]. These algorithms are interpreted as carrying out a distributed computation to maximize the aggregate source utility $\sum_i U_i(x_i)$, subject to capacity constraints, $y_l \leq c_l$ for all l . Here the strictly concave utility function $U_i(\cdot)$ represents the source's demand for rate, and is related to f_i by $U_i' = (f_i)^{-1}$. It follows from the work in [13] that for arbitrary strictly concave utility functions, these laws have a unique equilibrium and are stable in the absence of delay. The question we address here is what constraints our conditions for linear stability in the *presence* of delay impose on f_i , and thus on the allowable utility functions.

Now given an equilibrium point x_0, q_0 , let us impose the linearization requirement

$$\frac{\partial f_i}{\partial q_i} = -\frac{\alpha_i f_i(q_i)}{M_i \tau_i},$$

for some $0 < \alpha_i < \pi/2$, as needed in (8). Let us assume initially that α_i is constant. Then the above differential equation can be solved analytically, and gives the static source control law

$$x_i = f_i(q_i) := x_{\max,i} e^{-\frac{\alpha_i q_i}{M_i \tau_i}}.$$

Here $x_{\max,i}$ is a maximum rate parameter, which can vary for each source, and can also depend on M_i, τ_i (but not on q_i). So we find that an *exponential backoff* of source rates as a function of aggregate

price, can provide the desired control law, together with the link control in (13). The corresponding utility function (for which $f_i = (U_i')^{-1}$) is

$$U_i(x) = \frac{M_i \tau_i}{\alpha_i} x \left[1 - \log \left(\frac{x}{x_{\max,i}} \right) \right], \quad x \leq x_{\max,i}.$$

Are these exponential laws the only possible choice? Among static laws, there is a further degree of freedom in letting the parameter α_i be a function of the operating point. In general, we would allow any mapping $x_i = f_i(q_i)$ that satisfies the differential inequality

$$0 \geq \frac{\partial f_i}{\partial q_i} \geq -\frac{\pi f_i(q_i)}{2 M_i \tau_i}. \quad (14)$$

Essentially, the requirement is that the slope of the source rate function (the "elasticity" in source demand) decreases with delay τ_i ; in other words, the longer the delay, the slower the source control must respond, in order to avoid instability. Similar considerations apply to the number of bottlenecks M_i .

We could also consider certain dynamic control laws at the source; it can be shown, however, that under our stability restrictions, this does not add any further degrees of freedom to the equilibrium structure that can be obtained from static laws.

Finally, we emphasize that while the above implementations will behave as required around equilibrium, we have not proved global convergence to equilibrium in the nonlinear case.

6 Signaling requirements

We briefly discuss here the information needed at sources and links to implement our dynamic laws, and the resulting communication requirements.

Links generate prices by integrating the excess flow $y_l - c_l$ with respect to the virtual capacity; this is easily implemented by maintaining a "virtual queue" variable.

Sources must have access to the round-trip time τ_i , which can be obtained by timing packets and their acknowledgments, and two variables from the network: the aggregate price q_i , and the number of bottlenecks M_i . To communicate q_i , the technique of random exponential marking [2] can be used. Here, an ECN bit would be marked at each link l with probability $1 - \phi^{-p_l}$, where $\phi > 1$ is a global constant. Assuming independence, the overall probability that a packet from source i gets marked is (see [2])

$$1 - \phi^{-q_i},$$

and therefore q_i can be estimated from marking statistics. Alternatively packet dropping can be

used in lieu of marking as in other active queue management schemes [4].

Regarding M_i , in the simplest implementation one would employ an upper bound, which could be a globally known constant, or based on the total number of links in the route, found e.g. from traceroute information in IP. For a more aggressive control one would need to communicate in real-time, how many links are bottlenecks. This can be done in an analogous way as what is done with prices, by having any bottleneck link mark a second ECN bit with probability $1 - \phi^{-1}$. Then the probability of marking during a route is $1 - \phi^{-M_i}$, which allows for the estimation of M_i at the sources in real time.

Remark 2 *Estimation of τ_i and M_i , even if errors are neglected, will yield time-varying source control laws; again, our linearized stability analysis would have to be extended to prove global convergence.*

7 Conclusion

We have obtained decentralized congestion control laws, which provide local stability for arbitrary network routing, delay, and link capacities, improving certain limitations of current TCP. We have also provided an implementation, based on integration at the links and a static control law at the source, that has desirable equilibria and linearizes as required. This law is a special case of those considered in [13], which are interpreted as maximizing aggregate source utility; from this point of view the result in this paper poses constraints on the allowable classes of utility functions that would yield the desired local stability. Finally we have described mechanisms based on marking ECN bits that can provide approximate implementations of these control laws. Packet-level implementations of this protocol are currently under development [15].

References

- [1] E. Altman, T. Basar and R. Srikant. "Congestion control as a stochastic control problem with action delays", *Automatica*, December 1999.
- [2] S. Athuraliya and S. H. Low, "Optimization flow control - II: Implementation", *preprint*, <http://netlab.caltech.edu>, May 2000
- [3] L.S. Brakmo and L. L. Peterson. "TCP Vegas: end to end congestion avoidance on a global Internet". *IEEE Journal on Selected Areas in Communications*, 13(8), October 1995.
- [4] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance", *IEEE/ACM Trans. on Networking*, 1(4):397-413, Aug. 1993.
- [5] C. Hollot, V. Misra, D. Towsley and W.-B. Gong, "A control theoretic analysis of RED," *Proceedings IEEE Infocom 2001*.
- [6] V. Jacobson, "Congestion avoidance and control", *Proc. ACM SIGCOMM '88*.
- [7] R. Johari and D. Tan, "End-to-End Congestion Control for the Internet: Delays and Stability", *Cambridge Univ. Statistical Lab. Research Report 2000-2*.
- [8] 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.
- [9] F. P. Kelly, "Models for a Self-Managed Internet", *Philosophical Transactions of the Royal Society A358* (2000) 2335-2348.
- [10] S. Kunniyur and R. Srikant, "End-to-end congestion control schemes: utility functions, random losses and ECN marks", *Proc. IEEE Infocom*, 2000.
- [11] S. Kunniyur and R. Srikant, "A time-scale decomposition approach to adaptive ECN marking", *Proc. IEEE Infocom*, 2001.
- [12] S. H. Low. A duality model of TCP flow controls. In *Proceedings of ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, <http://netlab.caltech.edu>, September 18-20 2000.
- [13] S. H. Low and D. E. Lapsley, "Optimization flow control - I: basic algorithm and convergence" *IEEE/ACM Trans. on Networking*, Vol 7(6) Dec 1999.
- [14] Steven H. Low, Larry Peterson, and Limin Wang. Understanding Vegas: a duality model, to appear in *Proceedings of ACM Sigmetrics*, June 2001. <http://netlab.caltech.edu/pub.html>.
- [15] S. H. Low, F. Paganini, J. Wang, S. A. Adlakha, and J. C. Doyle. Dynamics of TCP/AQM and a scalable control. Submitted for publication, July 2001. <http://netlab.caltech.edu>.
- [16] S. Mascolo, "Congestion control in high-speed communication networks using the Smith principle", *Automatica*, 1999.
- [17] L. Massoulié. "Stability of distributed congestion control with heterogeneous feedback delays." Technical report, Microsoft Research, 2000. MSR-TR-2000-111.
- [18] V. Misra, W-B Gong, D. Towsley, "Fluid-Based Analysis of a Network of AQM Routers Supporting TCP Flows with an Application to RED", *Proc. ACM/SIGCOMM*, 2000.
- [19] H. Ozbay, S. Kalyanaraman, A. Iftar, "On rate-based congestion control in high-speed networks: design of an H_∞ based flow controller for single bottleneck", *Proc. American Control Conference*, 1998.
- [20] F. Paganini, "On the stability of optimization-based flow control", *Proc 2001 American Control Conference*.
- [21] F. Paganini, "Flow control via pricing: a feedback perspective", *Proceedings 2000 Allerton Conference*, Monticello, IL.
- [22] G. Vinnicombe. On the stability of end-to-end congestion control for the internet. Tech.report, Univ. of Cambridge, 2000. CUED/F-INFENG/TR.398.