

ON THE STABILITY OF NETWORKS OPERATING TCP-LIKE CONGESTION CONTROL

Glenn Vinnicombe

*Department of Engineering
University of Cambridge*

Abstract: We derive decentralized and scalable stability conditions for a fluid approximation of a class of Internet-like communications networks operating a modified form of TCP-like congestion control. The network consists of an arbitrary interconnection of sources and links with heterogeneous propagation delays. The model here allows for arbitrary concave utility functions and the presence of dynamics at both the sources and the links.

Keywords: Robust control, Communication networks, Complex systems

NOTATION

$\sigma(Z)$ denotes the spectrum of a square matrix Z and $\rho(Z)$ its spectral radius. Whenever the meaning is unambiguous, $\{f_i\}$ is used as an abbreviation for $\{f_i : i = 1, 2, \dots\}$. In particular, $\text{Co}\{x_i\}$ denotes the convex hull of the set of points $\{x_1, x_2, \dots\}$ and $\text{diag}(x_i)$ denotes the matrix with the elements x_1, x_2, \dots on the leading diagonal and zeros elsewhere.

1. INTRODUCTION

We consider a communications network, such as the Internet, consisting of an interconnection of users/sources which generate data and resources/links which carry it. The key constraint of the network that we are interested in is its decentralized nature: control information can only be passed along the same routes as the data that is being transmitted, and with the same propagation delay as that data. We wish to investigate the limitations imposed by this structure. Such a structure might arise, for example, if the resources are allowed to communicate with the users by manipulating certain reserved portions of the data packet header (by marking for example). (Kelly *et al.*, 1998) have shown that a certain network utility optimization problem may be solved in a decentralized manner over this structure. In this scheme, each link sets a price per unit flow, based on the aggregate flow through that link,

and the sources set their transmission rates based on the aggregate price they see. In the absence of delays, this scheme is globally stable. Moreover, this stability is maintained even when the sources are allowed to adapt their rates arbitrarily fast to achieve arbitrarily high utilization of the network. In the presence of delays this is no longer true. In this paper we derive a sufficient condition for the local stability of such a scheme, for arbitrary network topologies and heterogeneous round trip times. Moreover, we show that this condition captures a trade off between utilization of the network and its speed of response.

Associated with each source is a route, which is the collection of links through which information from that source is flowing. If x_r is the sending rate of source r then the flow through each link is given by

$$y_l(t) = \sum_{r:r \text{ uses } l} x_r(t - \vec{\tau}_{lr}). \quad (1)$$

where $\vec{\tau}_{lr}$ denotes the propagation delay from source r to link l . Similarly, if p_l is the price per unit flow at link l , then the aggregate price back at a source is given by

$$q_r(t) = \sum_{l:l \text{ used by } r} p_l(t - \overleftarrow{\tau}_{lr}) \quad (2)$$

where $\overleftarrow{\tau}_{lr}$ represents the return delay from link l to source r . We assume throughout that

$$\vec{\tau}_{lr} + \overleftarrow{\tau}_{lr} = T_r \quad \forall l, \quad (3)$$

where T_r is the round-trip delay of the r th route. This assumption is consistent with the price information being communicated back to the source via the recipient, along with acknowledgements.

We assume that the link prices are set according to the law

$$p_l = f_l(z_l), \quad f_l > 0, f'_l > 0 \quad (4)$$

where z_l is an exponentially smoothed flow rate, satisfying

$$\beta_l \dot{z}_l + z_l = y_l, \quad \beta_l \geq 0. \quad (5)$$

$\beta_l = 0$ corresponds to the price being set as an instantaneous function of the rate, $\beta_l > 0$ corresponds to the perhaps more realistic scenario where the rate is estimated either by arrivals over an exponentially weighted window or from queue lengths. (Alternatively, it might be the price information itself that is smoothed, i.e. (4) and (5) might be replaced by $\hat{p}_l = f_l(y_l)$ and $\beta_l \dot{\hat{p}}_l + p_l = \hat{p}_l$. This would give the same linearized transfer function from y_l to p_l and so all of our results will still hold.)

To complete the picture, the sources each have a utility function $U_r(\cdot)$ and attempt to maximize their net utility $U_r(x_r) - q_r x_r$ by setting their rates according to

$$\dot{x}_r(t) = k_r x_r(t - T_r) \left(1 - \frac{q_r(t)}{U'_r(x_r(t))} \right). \quad (6)$$

We shall assume that these utility functions satisfy

$$U'_r(x_r) > 0, \quad U''_r(x_r) < 0 \quad (7)$$

and so are concave. It is shown in (Kelly *et al.*, 1998) that if $T_r = 0$ for each route and $\beta_l = 0$ for each link then the interconnection is globally stable and converges to the unique equilibrium which maximizes the global utility $\sum_r U_r(x_r) - \sum_l \int_0^{y_l} p_l(y) dy$. For a closely related source law a simple condition on the gains k_r was derived in (Johari and Tan, 2000) which guarantees local stability of this equilibrium whenever all the round-trip trip times are equal. This same condition was conjectured to guarantee stability in the case of heterogeneous round trip delays. This conjecture is true ((Massoulié, 2000), (Vinnicombe, 2000)). In this paper we shall derive similar stability conditions for the general source law (6) and for lags at the links, i.e. $\beta_l > 0$. The condition is that if there exists a globally agreed constant B such that $\frac{\hat{y}_l f'_l(\hat{y}_l)}{f_l(\hat{y}_l)} \leq B$ at each link l , and $k_r T_r \leq \frac{1}{B}$ for each source r then the network is locally stable about its equilibrium for any collection of lag time constants β_l and any admissible utility functions. The condition on $f_l(\cdot)$ would be satisfied, for example, if the price is set as

$$f_l(y_l) = \left(\frac{y_l}{C_l} \right)^B$$

which is the probability that an M/M/1 queue with an arrival rate y_l and capacity C_l is of length B or greater.

We apply these results to a fluid flow approximation of a network with sources operating TCP-like algorithms, where the price is interpreted as the probability that a packet is marked and the dynamics at the links may be interpreted as a combination of queueing effects and the deliberate smoothing introduced by systems such as RED. For the usual TCP algorithm, the conclusion is that the network can be guaranteed to be stable only if the number of packets in flight is sufficiently large. We propose a modification of TCP which avoids this problem. (We should point out, though, that our results are for rate-based control and ignore the potentially stabilizing effects of the window-based control used in current TCP. It is argued elsewhere, though, (e.g. (Kelly, 2001b)) that this distinction should disappear in the limiting regime where capacities increase and queueing delays and queue emptying times become small in relation to propagation delays)

2. MAIN RESULT

We can write relationship (1) in terms of Laplace transforms as

$$\bar{y} = R(s)\bar{x}$$

where

$$R_{lr} = \begin{cases} e^{-s\vec{\tau}_{lr}} & \text{if route } r \text{ uses link } l \\ 0 & \text{otherwise} \end{cases}$$

Since

$$\vec{\tau}_{lr} + \overleftarrow{\tau}_{lr} = T_r \quad \forall l,$$

where T_r is the round-trip time on route r , we can write (2) as

$$\bar{q}(s) = \text{diag}(e^{-sT_i}) R^T(-s) \bar{p}(s).$$

These routing relations also hold for small perturbations of course, that is

$$\delta \bar{y} = R(s) \delta \bar{x}$$

and

$$\delta \bar{q}(s) = \text{diag}(e^{-sT_i}) R^T(-s) \delta \bar{p}(s)$$

where $y(t) = \hat{y} + \delta y(t)$ etc. Note that the matrix $R(0)$ also determines the static relationships between equilibrium values, i.e.,

$$\hat{y} = R(0)\hat{x}, \quad \hat{q} = R(0)^T \hat{p}. \quad (8)$$

The source law (6) may be linearized around the equilibrium $q_r = U'_r(x_r)$ to give

$$s \delta \bar{x}_r(s) = k_r \frac{\hat{x}_r}{\hat{q}_r} (U''_r(x_r) \delta \bar{x}_r(s) - \delta \bar{q}_r(s))$$

or

$$\delta \bar{x}_r(s) = -k_r T_r \frac{\hat{x}_r}{\hat{q}_r} \frac{1}{s T_r + \alpha_r} \delta \bar{q}_r(s) \quad (9)$$

where $\alpha_r = -T_r \frac{\hat{x}_r}{\hat{q}_r} U''_r(x_r) > 0$.

The link law (4–5) may be linearized to give

$$\delta \bar{p}_l(s) = f'_l(\hat{y}_l) \frac{1}{s\beta_l + 1} \delta \bar{y}_l(s) \quad (10)$$

First we state a stability result when all the time constants β_l are 0. The proof is a minor modification of that in (Vinnicombe, 2000) for a slightly different source law and, like that of (Vinnicombe, 2000), relies on ideas taken from (Johari and Tan, 2000) to bound the size of the return ratio.

Proposition 1. The interconnection described by (1–7), with $\beta_l = 0$ for all l and $k_r \geq 0$ for all r , is locally asymptotically stable around the equilibrium $y_l, z_l = \hat{y}_l; q_r = U'(x_r)$ if there exists a B such that

- (1) $\frac{\hat{y}_l f'_l(\hat{y}_l)}{f_l(\hat{y}_l)} \leq B$ for all l , and
- (2) $k_r T_r \leq \frac{\pi}{2} \frac{1}{B}$ for all r .

Remark 2. Note that the first condition may be satisfied by $f_l(y_l) = \left(\frac{y_l}{C_l}\right)^B$ giving $f'_l(y_l) = \frac{B}{y_l} f_l(y_l)$. This $f_l(\cdot)$ is precisely the probability that an M/M/1 queue with an arrival rate y_l and capacity C_l is of length B or greater, and so is a fairly natural price in this context. This result shows that such a pricing function remains desirable for more realistic queueing systems.

The result also captures a tradeoff between utilization of the network and speed of convergence. In (Kelly *et al.*, 1998) the prices $f_l(\cdot)$ are regarded as barrier functions in the global maximization of $\sum_r U_r(x_r)$ subject to the constraint that the flow at each link is no greater than the capacity, a larger B would thus correspond to the price remaining small until closer to capacity and then increasing more rapidly around capacity. However, a larger B also requires that the sources react more slowly to fluctuations in the network if stability is to be maintained.

PROOF. [of Theorem 1]. Using the linearizations developed above, it can be seen that the return ratio seen at the source is given by

$$\text{diag} \left(k_i T_i \frac{\hat{x}_i}{\hat{q}_i} \frac{e^{-sT_i}}{sT_i + \alpha_i} \right) R^T(-s) \text{diag} (f'_i) R(s).$$

Since the open-loop system is stable, we just need to show that the eigenvalues of this function, for $s = j\omega$, do not encircle the point -1 . Now, these eigenvalues are identical to those of

$$\text{diag} \left(\frac{e^{-j\omega T_i}}{j\omega T_i + \alpha_i} \right) \hat{R}^T(-j\omega) \hat{R}(j\omega). \quad (11)$$

where

$$\hat{R}(j\omega) = \text{diag} \left(\sqrt{f'_i} \right) R(j\omega) \text{diag} \left(\sqrt{k_i T_i \frac{\hat{x}_i}{\hat{q}_i}} \right).$$

Furthermore

$$\begin{aligned} \bar{\sigma}^2(\hat{R}(j\omega)) &= \rho \left(\hat{R}^T(-j\omega) \hat{R}(j\omega) \right) \leq \\ \frac{\pi}{2} \rho \left(\text{diag} \left(\frac{1}{\hat{q}_i} \right) R^T(-j\omega) \text{diag} \left(\frac{\hat{p}_i}{\hat{y}_i} \right) R(j\omega) \text{diag}(\hat{x}_i) \right) \\ &\leq \frac{\pi}{2} \left\| \text{diag} \left(\frac{1}{\hat{q}_i} \right) R^T(-j\omega) \text{diag}(\hat{p}_i) \right\|_{\infty} \cdot \\ &\quad \left\| \text{diag} \left(\frac{1}{\hat{y}_i} \right) R(j\omega) \text{diag}(\hat{x}_i) \right\|_{\infty} = \frac{\pi}{2} \end{aligned}$$

since the row sums of each of these matrices all equal 1 by virtue of (8).

Now, if λ is an eigenvalue of (11) then there exists a v , $\|v\| = 1$, such that

$$\hat{R}(j\omega) \text{diag} \left(\frac{e^{-j\omega T_i}}{j\omega T_i + \alpha_i} \right) \hat{R}^T(-j\omega) v = \lambda v$$

giving

$$\lambda = v^* \hat{R}(j\omega) \text{diag} \left(\frac{e^{-j\omega T_i}}{j\omega T_i + \alpha_i} \right) \hat{R}^T(-j\omega) v.$$

Since $\|\hat{R}^T(-j\omega) v\|^2 \leq \frac{\pi}{2}$ this implies that

$$\lambda \in \text{Co} \left(0 \cup \left\{ \frac{\pi}{2} \frac{e^{-j\omega T_i}}{j\omega T_i + \alpha_i} \right\} \right).$$

Now, for $\alpha > 0$, $\left(\frac{e^{-jx}}{jx + \alpha}\right)$ always crosses the real axis to the right of the point -1 (see Fig. 1). It follows that the eigenloci cannot cross the real axis at or to the left of the point -1 , and that the closed-loop system is stable by the generalized Nyquist stability criterion ((Desoer and Yang, 1980)). \square




Fig. 1. Convex hull of $\frac{\pi}{2} \exp(-jx)/(jx + \alpha)$, $x, \alpha \in \mathbb{R}_+$

We now extend the previous result to allow for first order dynamics at the links. The following theorem shows that, in this case, stability is guaranteed if the gains are reduced by a factor of $\pi/2$. So, the lags introduced at the links are not having much impact on nominal stability. The same is not necessarily true for the robust stability of the interconnection, as will be discussed in Remark 4.

Theorem 3. The interconnection described by (1–7), with $k_r \geq 0$ for all r , is locally asymptotically stable around the equilibrium $y_l, z_l = \hat{y}_l; q_r = U'(x_r)$ if there exists a B such that

- (1) $\frac{\hat{y}_l f'_l(\hat{y}_l)}{f_l(\hat{y}_l)} \leq B$ for all l , and
- (2) $k_r T_r \leq \frac{1}{B}$ for all r .

PROOF. [of Theorem 3]. The proof is an extension of that of Proposition 1. The return ratio seen at the source is now given by

$$\text{diag} \left(k_i T_i \frac{\hat{x}_i}{\hat{q}_i} \frac{e^{-sT_i}}{sT_i + \alpha_i} \right) R^T(-s) \text{diag} \left(\frac{f'_i}{s\beta_i + 1} \right) R(s).$$

Again, since the open-loop system is stable, we just need to show that the eigenvalues of this function, for $s = j\omega$, do not encircle the point -1 . These eigenvalues are identical to the eigenvalues of

$$\text{diag} \left(\frac{e^{-j\omega T_i}}{j\omega T_i + \alpha_i} \right) \hat{R}^T(-j\omega) \text{diag} \left(\frac{1}{j\omega\beta_i + 1} \right) \hat{R}(j\omega). \quad (12)$$

where \hat{R} is defined as

$$\hat{R}(j\omega) = \text{diag} \left(\sqrt{f'_i} \right) R(j\omega) \text{diag} \left(\sqrt{k_i T_i \frac{\hat{x}_i}{\hat{q}_i}} \right).$$

as in Proposition 1. Here, though,

$$\begin{aligned} \bar{\sigma}^2(\hat{R}(j\omega)) \leq & \left\| \text{diag} \left(\frac{1}{\hat{q}_i} \right) R^T(-j\omega) \text{diag}(\hat{p}_i) \right\|_{\infty} \\ & \left\| \text{diag} \left(\frac{1}{\hat{y}_i} \right) R(j\omega) \text{diag}(\hat{x}_i) \right\|_{\infty} \leq 1 \end{aligned}$$

using (8) again.

Now, if λ is an eigenvalue of (12) then there exists a v , $\|v\| = 1$, such that

$$\begin{aligned} & \text{diag} \left(\frac{1}{j\omega\beta_i + 1} \right) \hat{R}(j\omega) \\ & \cdot \text{diag} \left(\frac{e^{-j\omega T_i}}{j\omega T_i + \alpha_i} \right) \hat{R}^T(-j\omega) v = \lambda v \end{aligned}$$

and hence

$$\lambda = \frac{v^* \hat{R}(j\omega) \text{diag} \left(\frac{e^{-j\omega T_i}}{j\omega T_i + \alpha_i} \right) \hat{R}^T(-j\omega) v}{v^* \text{diag} (j\omega\beta_i + 1) v}.$$

Since $\|\hat{R}^T(-j\omega)v\| \leq 1$ this implies that

$$\lambda \in \frac{\text{Co} \left(0 \cup \left\{ \frac{e^{-j\omega T_i}}{j\omega T_i + \alpha_i} \right\} \right)}{\text{Co} \{j\omega\beta_i + 1\}}.$$

Now, for $\alpha > 0$, $\Re \left(\frac{e^{-jx}}{jx + \alpha} \right) > -1$. In addition, the real part of the denominator is always 1. Consequently this region contains no points on the real axis at or to the left of the point -1 . It follows that the eigenloci cannot cross the real axis at or to the left of the point -1 , and so the closed-loop system is stable by the

generalized Nyquist stability criterion ((Desoer and Yang, 1980)). \square

Remark 4. The theorem does not directly guarantee any level of robust stability, however consideration of the proof reveals that the stability is indeed robust to independent perturbation of the source laws as long as these perturbations do not take their Nyquist locus to the left of the line $\Re(\cdot) = -1$. Appropriate reduction of the gains k_r to less than $1/T_r$ could thus guarantee any desired level of robustness to perturbations of the source laws. The situation at the links is not so good though; gain variations can be tolerated, but small levels of extra phase lag there can cause instability unless limits are placed on the β_l . If the time constants β_l are of the order of the round trip times of routes using that link then one might expect a degree of robustness. This is indeed true, and rigorous stability conditions based on this idea are developed in (Vinnicombe, 2002).

3. TCP-LIKE ALGORITHMS

We consider a class of algorithms where, as in TCP, the source maintains a window cwnd of sent but not yet acknowledged packets. This window is incremented by $a\text{cwnd}^n$ for each unmarked acknowledgement, and decremented by $b\text{cwnd}^m$ for each marked acknowledgement, where $m > n$. Conventional TCP in its congestion avoidance phase uses an increment of $1/\text{cwnd}$ and a decrement $\text{cwnd}/2$ (and with packets being dropped rather than marked). Properties of this algorithm may be analysed using the methods of the previous section provided it is assumed that queueing delays are small relative to propagation delays, which is equivalent to the assumption that queue dynamics may be averaged over periods significantly shorter than round trip times. See (Kelly, 2001b) for justification of this limiting regime, and the derivation of models of this kind. Our conclusion will be that, if $a = 1/B$, $n = 0$, i.e. the increment is changed from $1/\text{cwnd}$ to $1/B$ for some global constant B , then the network will be locally stable provided the price function at each link satisfies $y_l f'_l(y_l)/f_l(y_l) \leq B$. In contrast, the usual version of TCP described above can only be guaranteed to be stable when $\text{cwnd} \geq B$; furthermore, it will be overly sluggish when $\text{cwnd} \gg B$. We emphasize that these conclusions are only valid in the limiting regime described above; in the current Internet queueing delays are typically large.

We let

$$x_r(t) = \text{cwnd}(t)/T_r$$

be a continuous approximation of the sending rate to obtain

$$\frac{d}{dt} \text{cwnd}(t) = \frac{a\text{cwnd}^n(1 - q_r(t)) - b\text{cwnd}^m q_r(t)}{T_r/\text{cwnd}(t - T_r)}$$

where the “price” $q_r(t)$ should now be interpreted as the probability that an acknowledgement received at time t carries a mark. In terms of rates, this becomes

$$T_r \dot{x}_r(t) = x_r(t - T_r) \left(a(x_r(t) T_r)^n (1 - q_r(t)) - b(x_r(t) T_r)^m q_r(t) \right) \quad (13)$$

This equation does not quite fall into the framework of the previous section, because of the $(1 - q_r)$ term. However, the presence of this term does not affect the linearizations. The equation may be linearized about its equilibrium

$$\hat{q}_r = \frac{a(\hat{x}_r T_r)^n}{a(\hat{x}_r T_r)^n + b(\hat{x}_r T_r)^m}$$

to give

$$\begin{aligned} T_r \frac{d}{dt} \delta x_r(t) &= \hat{x}_r \left(-\delta q_r(t) (a(\hat{x}_r T_r)^n + b(\hat{x}_r T_r)^m) + \right. \\ &\quad \left. (a T_r n (\hat{x}_r T_r)^{n-1} (1 - \hat{q}_r) - b T_r m (\hat{x}_r T_r)^{m-1} \hat{q}_r) \delta x_r(t) \right) \\ &= \hat{x}_r \left(-\delta q_r(t) \frac{a(\hat{x}_r T_r)^n}{\hat{q}_r} \right. \\ &\quad \left. - b T_r (m - n) (\hat{x}_r T_r)^{m-1} \hat{q}_r \delta x_r(t) \right) \\ &= - \left(\frac{a \hat{x}_r (\hat{x}_r T_r)^n}{\hat{q}_r} \right) \delta q_r(t) - \alpha_r \delta x_r(t). \quad (14) \end{aligned}$$

where

$$\alpha_r = b(m - n) (\hat{x}_r T_r)^m \hat{q}_r > 0$$

with the transfer function

$$\bar{\delta x}_r = - \frac{a \hat{x}_r (\hat{x}_r T_r)^n}{\hat{q}_r} \frac{1}{(s T_r + \alpha_r)} \bar{\delta q}_r$$

A further complication is that marking does not result in the the prices being additive, unless the marking probabilities at each link are vanishingly small. The stability conditions of Theorem 3 still hold nevertheless ((Kelly, 2001a)) as we now show. If we assume that marking at each resource is independent, then (2) should be replaced by

$$q_r(t) = 1 - \prod_{l:l \text{ used by } r} (1 - p_l(t - \tau_{lr}))$$

where p_l should be interpreted as the marking probability at link l . This relationship linearizes to

$$\delta q_r(t) = \sum_{l:l \text{ used by } r} \frac{1 - \hat{q}_r}{1 - \hat{p}_l} \delta p_l(t - \tau_{lr})$$

and the relationship between the equilibrium values is then given, of course by,

$$\hat{q}_r = 1 - \prod_{l:l \text{ used by } r} (1 - \hat{p}_l)$$

Compared to the setting of Theorem 3, the term $\tilde{R}^T(-s)$ in the return ratio should be replaced by $\tilde{R}^T(-s)$ where

$$\tilde{R}_{lr} = \begin{cases} e^{-s \tau_{lr}} \frac{1 - \hat{q}_r}{1 - \hat{p}_l} & \text{if route } r \text{ uses link } l \\ 0 & \text{otherwise} \end{cases}$$

. The same stability condition now holds provided

$$\left\| \text{diag} \left(\frac{1}{\hat{q}_i} \right) \tilde{R}^T(-j\omega) \text{diag}(\hat{p}_i) \right\|_{\infty} \leq 1.$$

But this last relationship is true because the row sums are given by

$$\begin{aligned} \frac{1 - \hat{q}_r}{\hat{q}_r} \sum_{l:l \text{ used by } r} \frac{\hat{p}_l}{1 - \hat{p}_l} &= \frac{1}{\hat{q}_r} \sum_{l:l \text{ used by } r} \hat{p}_l \frac{1 - \hat{q}_r}{1 - \hat{p}_l} \\ &= \frac{\text{Pr}(\text{packet on route } r \text{ marked exactly once})}{\text{Pr}(\text{packet on route } r \text{ marked at least once})} \leq 1 \end{aligned}$$

Comparison with (9) now shows that the appropriately modified version of Theorem 3 guarantees local stability whenever $a(\hat{x}_r T_r)^n < 1/B$. As mentioned above, conventional TCP has $n = -1$, $a = 1$ and so stability is only guaranteed if the equilibrium congestion window $\hat{x}_r T_r$ is *greater* than B where B is an upper bound on $y_l f'_l(y_l)/f_l(y_l)$ and $p_l(t) = f_l(y_l(t))$ at each link. For routers operating threshold marking, B might correspond to something like the buffer size (in packets) at which packets are marked. (This would be exact if the buffers behave as M/M/1 queues). An appealing alternative is to take $n = 0$ and choose $a = 1/B$, that is choose a fixed increment of something like the reciprocal of the average buffer level at which packets are marked.

4. ACKNOWLEDGEMENTS

I would like to thank Frank Kelly for bringing this problem to my attention and for a number of illuminating conversations on these subjects.

5. CONCLUSIONS

We have derived a simple decentralized stability condition for a network consisting of an interconnection of links and sources where the only required shared information is of a single global constant B . The results suggest that TCP as usually implemented is likely to be prone to instabilities when the congestion window is small, and overly sluggish when it is large. We have suggested a simple scalable modification where the shared information is of the approximate buffer length at which packets are marked. It should be emphasized that these conclusions are only valid in the limiting regime where queueing delays and queue emptying times are small in relation to propagation delays. A further conclusion is that the exponential smoothing introduced by protocols such as RED should not compromise stability (at least as long as the time constants are not too large.)

6. REFERENCES

- Desoer, C. A. and Y. T. Yang (1980). On the generalized Nyquist stability criterion. *IEEE Transactions on Automatic Control* **25**, 187–196.
- Johari, R. and D. Tan (2000). End-to-end congestion control for the Internet: delays and stability. Technical report. Statistical Laboratory, University of Cambridge.
- Kelly, F. (2001a). Stability of a network model of TCP. unpublished notes.
- Kelly, F. P. (2001b). *Mathematics Unlimited – 2001 and Beyond*. Chap. Mathematical modelling of the Internet, pp. 685–702. Springer-Verlag, Berlin.
- Kelly, F. P., A. K. Maulloo and D. K. H. Tan (1998). Rate control in communication networks: shadow prices, proportional fairness, and stability. *Journal of the Operational Research Society* **49**, 237–252.
- Massoulié, L. (2000). Stability of distributed congestion control with heterogeneous feedback delays. Technical Report MSR-TR-2000-111. Microsoft Research.
- Vinnicombe, G. (2000). On the stability of end-to-end congestion control for the Internet. Technical Report CUED/F-INFENG/TR.398. Engineering Department, University of Cambridge. <http://www-control.eng.cam.ac.uk/gv/internet>.
- Vinnicombe, G. (2002). Robust congestion control for the internet. Submitted to SIGCOMM 2002.