Are optimization based Internet congestion control models fragile with respect to TCP structure and symmetry?

Ioannis Lestas*, Glenn Vinnicombe
Department of Engineering
University of Cambridge,
Cambridge CB2 1PZ, UK
(icl20,gv)@eng.cam.ac.uk

***Abstract*— Scalable stability conditions derived so far for optimization based models of congestion control protocols, can be shown mathematically to hold for arbitrary networks provided the underlying protocol is symmetric. In practical implementations, however, deviation from this symmetry is inevitable. It is hence crucial to establish whether these models are fragile with respect to a relaxation of the symmetry assumption. We prove in this paper that this is not the case by presenting scalable, decentralized conditions, that guarantee stability for models of non-symmetric, TCP like protocols, of arbitrary interconnection. These conditions can be seen as local perturbations to the symmetric results and we illustrate how they converge to those derived for symmetric protocols as the degree of non symmetry becomes smaller. Finally, we show the way the decrease rule in TCP is associated with robust stability to non symmetric deviations from the protocol.**

Notation

$\sigma(M)$ denotes the spectrum of a square matrix M, $\rho(M)$ its spectral radius and $|M|$ the elementwise absolute value of the matrix i.e. $|[M_{ij}]| := [|M_{ij}|]$. $Co(S)$ denotes the convex hull of a set S and $\text{diag}(x_i)$ the matrix with elements $x_1, x_2, \ldots$ on the leading diagonal and zeros elsewhere.

I. Introduction

In [1] an optimization based decentralized model is introduced as a means to analyze Internet congestion control protocols. A Lyapunov type proof is given for global stability of the algorithm in the absence of delays. This work is extended in [2], [3], [4] by deriving scalable, decentralized local stability conditions when delays are present. Along the same lines, scalable control laws are also suggested in [5].

The main common feature of this work, which makes the mathematical analysis so elegant, is the symmetry of the underlying protocol. This enables the exploitation of the structure of the interconnections between the dynamical elements of the system. The symmetry lies in the fact that the resources do not discriminate between users when producing congestion signals (prices) and equivalently users do not discriminate between resources when determining their data flow as function of the aggregate prices they receive. Nevertheless in practice such discrimination does occur. Consider, for example, variable packet sizes from users in a packet marking scheme for signalling congestion. Depending then on whether the congested resource at a router is packet rate or bandwidth, routers will scale differently the packet rate from various users when they take it into account to determine the marking probability. Hence in a scenario where price is set by some routers as a function of bitrate and different users use different packet sizes, the stability results in [1], [2], [3], [4], [5] do not hold.

It is known in complex systems that, even though we may have robustness with respect to parameters of interest, the system may be fragile when small failures in the underlying structure occur (e.g. failure of an actuator on an aeroplane). It is therefore important in the networks problem to verify that the models we are considering are not fragile with respect to the protocol structure and symmetry assumed. This is what we prove in this paper by deriving decentralized, scalable stability conditions for arbitrary non-symmetric networks, as an extension of the ones derived for symmetric protocols. In comparison with symmetric results, the conditions derived include a more restrictive condition that affects only non-symmetric paths and scales with the degree of non symmetry.

II. Problem Formulation

We use the notation in [4], [1] and define the following:

x_r is the flow rate associated with route r.

$U_r(x_r)$ is the utility of the user on route r, which is a continuously differentiable, strictly concave, increasing function of the flow x_r.

$T_r = \tau_{lr} + \tau_{rl}$ is the round trip delay of the rth route, with τ_{rl} being the propagation delay from source r to link l and τ_{lr} the return delay from link l to source r.

$$y_l(t) = \sum_{r:r \text{ uses } l} x_r(t - \tau_{rl}) \tag{1}$$

is the aggregate flow through link l.

$p_l = f_l(y_l)$ is the link price per unit flow, which is a non-negative, strictly increasing function of the aggregate flow through the resource l. We assume, in this paper, link prices to be static functions of the flow (this is valid for

*The work of the author was supported by a Gates Scholarship.

low length queues, and large capacities) as in [1].

$$q_r(t) = \sum_{l:l \text{ used by } r} p_l(t - \tau_{lr}) \tag{2}$$

is the aggregate price along route r.

The control law is performed by the users according to
$$\dot{x}_r(t) = k_r x_r(t - T_r)\left(1 - \frac{q_r(t)}{U_r'(x_r(t))}\right) \tag{3}$$

The symmetry of such a congestion control protocol is best appreciated if we consider the underlying signal flow graph (as these are used by Mason in [6], [7]) of the relevant linearized multivariable system. Let $h_i(s)$ and $g_j(s)$ be the the Laplace transforms of the linearized user and link dynamics respectively. In a signal flow graph, transfer functions are represented as weighted edges, connecting nodes which correspond to points where a summation of signals takes place. We consider TCP like

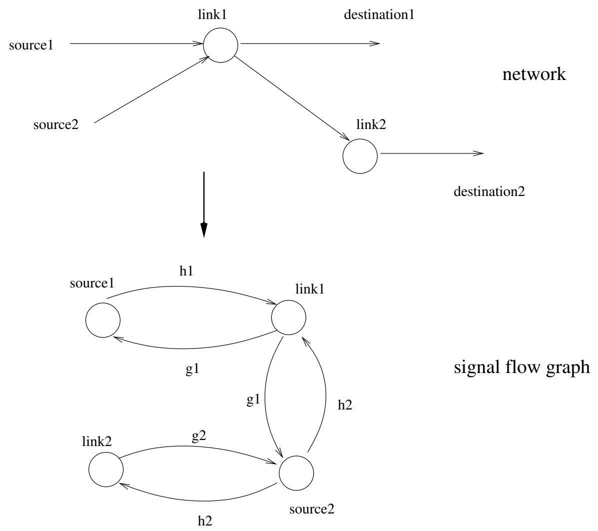

Fig. 1. An example of TCP like protocols and the relevant signal flow graph

protocols where users implement the control law based on the aggregate price from the resources along the route and the resource prices are generated as functions of the aggregate flow through the resource. Figure 1 illustrates the special structure of the signal flow graph for such protocols i.e. it is a directed, bipartite graph (sources are connected only to links). Furthermore when an edge exists between two nodes, an edge between the same nodes but in opposite direction also exists.

In this paper we derive first stability conditions for non-symmetric protocols corresponding to signal flow graphs where some of the directed edges are removed. We then extend these results to the case where the edges are weighted by different scalar factors instead of being eliminated and show how these stability conditions tend to the ones derived for the symmetric case when the weighting factors tend to 1.

III. MAIN RESULT

Taking Laplace transforms we can write (1) in the symmetric protocol as a vector equation

$$\bar{y} = R(s)\bar{x} \tag{4}$$

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

and (2) as
$$\bar{q}(s) = \text{diag}(e^{-sT_i})R^T(-s)\bar{p}(s) \tag{5}$$

We now consider a non symmetric protocol where some of the user dynamics weighted edges in the signal flow graph are removed. We thus define $\bar{R}$, a perturbed R matrix, as

$$\bar{R}_{ij} = R_{ij} \text{ for all } (i,j) \notin A \tag{6}$$

where $A = \{(i,j) : R_{ij} \neq 0, \bar{R}_{ij} = 0\}$ is the set of eliminated edges. We could have equivalently removed price weighted edges by perturbing the R^T matrix (or both price and user edges) with the subsequent analysis being very similar. For small perturbations about the equilibrium flow we get

$$\begin{aligned} y(t) &= \hat{y} + \delta y(t) && (7)\\ \overline{\delta y} &= \bar{R}(s)\overline{\delta x} && (8)\\ \overline{\delta q}(s) &= \text{diag}(e^{-sT_r})R^T(-s)\overline{\delta p}(s) && (9) \end{aligned}$$

and the equilibrium relations

$$\hat{y} = \bar{R}(0)\hat{x}, \quad \hat{q} = R^T(0)\hat{p} \tag{10}$$

Linearization of the source law (3) and the static link price gives

$$\overline{\delta x}_r(s) = -k_r \frac{\hat{x}_r}{\hat{q}_r}\frac{1}{s+\alpha_r}\overline{\delta q}_r(s),$$
$$\text{where} \quad \alpha_r = -\frac{\hat{x}_r}{\hat{q}_r}U_r''(\hat{x}_r) \tag{11}$$
$$\overline{\delta p}_l(s) = f_l'(\hat{y})\overline{\delta y}_l(s) \tag{12}$$

Breaking the loop at the source leads to the following return ratio

$$G(s) = \text{diag}\left(k_r \frac{\hat{x}_r}{\hat{q}_r}\frac{e^{-sT_r}}{s+\alpha_r}\right)R^T(-s)\text{diag}(f_l')\bar{R}(s) \tag{13}$$

Since the system is open loop stable it is sufficient to show that the eigenvalues of $G(j\omega)$ do not encircle the -1 point,

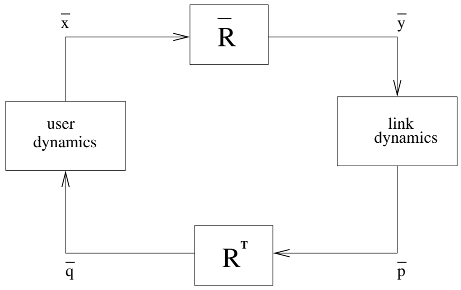

Fig. 2. Block diagram of interconnected network.

using the generalized Nyquist criterion ([8]). The stability condition in proposition 1 follows from Lemma 1 in the appendix.

Proposition 1: The interconnection described by (1-6),(10) is asymptotically stable around its equilibrium if there exists B such that all inequalities below are satisfied

$$f_l'(\hat{y}_l) \leq \frac{\hat{p}_l}{\hat{y}_l + \sum_r \hat{x}_r I_{((l,r)\in A)}} B \quad \forall l \tag{14}$$

$$k_r T_r \leq \frac{1}{B} \quad \forall r \tag{15}$$

$$0 \leq k_r \leq \frac{\alpha_r}{B|\lambda(S)|} \quad \forall (l,r) \in A \tag{16}$$

where $\alpha_r = -\frac{\hat{x}_r}{\hat{q}_r} U_r''(\hat{x}_r) > 0$, set S is defined as $S = \{(v_1^* + v_2^*)v_1 : |v_1| + |v_2| = 1, \quad v_1, v_2 \in \mathbb{C}\}$ and

$$\lambda(S) = \min_{x \in S \bigcap \mathbb{C}_-} \Re(x)$$

Remark 1: The first two conditions are the same as the ones derived in [4] for the symmetric case apart from the fact that (14) is slightly more strict. Condition (16) is due to the non-symmetric nature of the network.

Remark 2: Set S and constant $\lambda(S)$ are fixed ($\lambda(S) = -\frac{1}{8}$ as shown in the appendix) for the case considered, where elements of R are eliminated to form $\bar{R}$. Proposition 1 is stated in this general form because if elements of R are weighted instead of being eliminated, the same proposition holds but with the set S being replaced by a different perturbation set that scales with the degree of non-symmetry (this is discussed in more detail below).

Proof: [of Proposition 1] In order to use Lemma 1 we need to reduce $G(j\omega)$ to a similar form $F\hat{R}^T(-j\omega)\hat{\bar{R}}(j\omega)$ s.t. $\rho(|\hat{R}|^T|\hat{R}|) \leq 1$ and F is a diagonal matrix. This can be achieved by appropriate scaling by observing that

$$\begin{aligned}
&\rho\left(\operatorname{diag}\left(\frac{1}{\hat{q}_r}\right) R^T(-j\omega)\right. \\
&\qquad \left.\operatorname{diag}\left(\frac{\hat{p}_l}{\hat{y}_l + \hat{x}_r I_{((l,r)\in A)}}\right) R(j\omega)\operatorname{diag}(\hat{x}_r)\right) \leq \\
&\left\|\operatorname{diag}\left(\frac{1}{\hat{q}_r}\right) R^T(-j\omega)\operatorname{diag}(\hat{p}_l)\right\|_\infty \\
&\qquad \left\|\operatorname{diag}\left(\frac{1}{\hat{y}_l + \hat{x}_r I_{((l,r)\in A)}}\right) R(j\omega)\operatorname{diag}(\hat{x}_r)\right\|_\infty \leq 1
\end{aligned} \tag{17}$$

The final inequality follows from the fact that each term has row sum ≤ 1 using the equilibrium relations (10) and noting that $\hat{y}_l + \sum_r \hat{x}_r I_{((l,r)\in A)} = [R(0)\hat{x}]_l$. Now

$$\rho(G(j\omega)) = \rho\left(\operatorname{diag}\left(Bk_r \frac{e^{-j\omega T_i}}{j\omega + \alpha_r}\right) \hat{R}^T(-j\omega)\hat{\bar{R}}(j\omega)\right)$$

where $\quad \hat{R}^T(-j\omega) = \operatorname{diag}\left(\sqrt{\frac{\hat{x}_r}{\hat{q}_r}}\right) R^T(-j\omega)\operatorname{diag}\left(\sqrt{\frac{f_l'}{B}}\right)$

and $\quad \hat{\bar{R}}(j\omega) = \operatorname{diag}\left(\sqrt{\frac{f_l'}{B}}\right) \bar{R}(j\omega)\operatorname{diag}\left(\sqrt{\frac{\hat{x}_r}{\hat{q}_r}}\right)$

Using (17) and condition (14) it is deduced that $\rho(\hat{R}^T\hat{R}) \leq 1$ and similarly $\rho(|\hat{R}|^T|\hat{R}|) \leq 1$ (max row sums remain the same). Using Lemma 1 a sufficient condition for local stability with positive gains k_r, is that both (14) and the condition below are are satisfied:

$$\begin{aligned}
&-1 \notin Co\left(0 \cup \left\{Bk_r \frac{e^{-j\omega T_r}}{j\omega + \alpha_r} :\right.\right. \\
&\left.\omega \in \mathbb{R}_+,\ \alpha_r = -\frac{\hat{x}_r}{\hat{q}_r} U_r''(\hat{x}_r),\ (l,r) \notin A \quad \forall l\right\} \cup \\
&\left\{Bk_r \frac{e^{-j\omega T_r}}{j\omega + \alpha_r} S : \omega \in \mathbb{R}_+,\right. \\
&\qquad\qquad \left.\left.\alpha_r = -\frac{\hat{x}_r}{\hat{q}_r} U_r''(\hat{x}_r), \exists l \text{ s.t.} (l,r) \in A\right\}\right)
\end{aligned} \tag{18}$$

where S is the set defined in Lemma 1. For the symmetric case, set S in Lemma 1 becomes $[0,1]$ hence condition (18) is satisfied for all $\alpha_r > 0$ provided that $k_r T_r < \pi/2$ for all r, as this is illustrated in [4]. Note, however, that in the non symmetric case the boundary of S crosses the negative real axis at $-1/8$, therefore the factor multiplying S at $\omega = 0$ must be bounded, i.e. we need a condition at least as strict as (16). It turns out, as shown below, that bounds (14 - 16) are sufficient for (18). Condition (15) ensures that

$$\begin{aligned}
&\Re(x) > -1 \quad \text{for all} \quad x \in \\
&Co\left(\left\{Bk_r T_r \frac{e^{-j\omega T_r}}{j\omega T_r + \alpha_r T_r} : \omega, \alpha_r \in \mathbb{R}_+, k_r T_r \leq \frac{1}{B}\right\}\right)
\end{aligned} \tag{19}$$

and conditions (15), (16) ensure that

$$\begin{aligned}
&\Re(x) > -1 \quad \text{for all} \quad x \in \\
&Co\left(\left\{Bk_r \frac{e^{-j\omega T_r}}{j\omega + \alpha_r} S : \omega, \alpha_r \in \mathbb{R}_+,\right.\right. \\
&\qquad\qquad \left.\left. k_r T_r \leq \frac{1}{B}, Bk_r \leq \frac{\alpha_r}{|\lambda(S)|}\right\}\right)
\end{aligned} \tag{20}$$

Therefore (18) is satisfied. ■

To see statement (20) note that

$$\begin{aligned}
&Co\left\{Bk_r \frac{e^{-j\omega T_r}}{j\omega + \alpha_r} S : \omega, \alpha_r \in \mathbb{R}_+, k_r T_r \leq \frac{1}{B}, Bk_r \leq \frac{\alpha_r}{|\lambda(S)|}\right\} \subseteq \\
&\qquad Co\left\{\frac{e^{-jxBk_r T_r}}{jx + p} S : x \in \mathbb{R}_+, p \geq |\lambda(S)|, Bk_r T_r \leq 1\right\}
\end{aligned} \tag{21}$$

Figure 4 shows the extreme case where $Bk_r T_r = 1$ and $p = |\lambda(S)| = 1/8$ and it is clear that -1 is not included in the hull. (19) and (20) can be easily verified analytically, by deriving parameterized expression of the real part of elements in the sets. This is omitted here and illustrate instead their validity in figures 3 and 4 respectively.

As noted in Remark 3 in the appendix, if instead of eliminating edges we scale them by factors in $[0,1]$, the perturbation set S tends to $[0,1]$, as these factors tend to 1. Therefore condition (16) becomes less strict since $\lambda(S) \to 0$ and the conditions in Proposition 1 tend to those for the symmetric case, as these are stated in [4].

It should be emphasized that the importance of the bounds in proposition 1 lie on the one hand in the

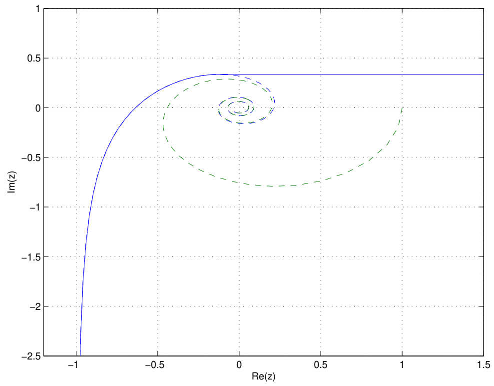

Fig. 3. $Co\left(\left\{\frac{e^{-jx}}{jx+p} : x,p \in \mathbb{R}_+\right\}\right)$

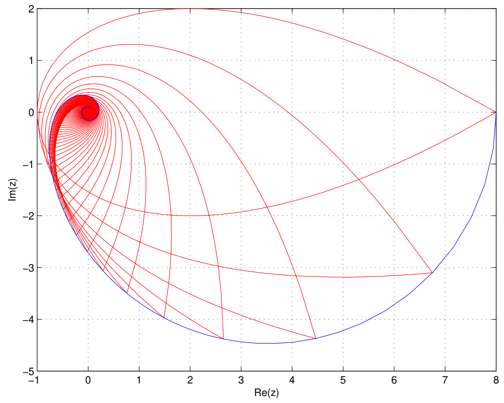

Fig. 4. $Co\left(\left\{\frac{e^{-jx}}{jx+1/8}S : x \in \mathbb{R}_+\right\}\right)$

fact that they are decentralized and hold for arbitrary interconnections like the results for symmetric protocols. Once we deviate from symmetry, the symmetric bounds still hold, i.e. feedback gain depends on delay and the nature of the price functions. An extra bound is, however, also introduced, that depends on the degree of non symmetry as well as the nature of the utility functions of the users. Notice that this bound only affects users behaving in a non-symmetric way. It also turns out that along a certain path, the worst case non symmetric behaviour determines the bound of the corresponding user gain (follows from remark 3 in the appendix).

Finally, it is important to mention that examples can easily be generated, where stability of the algorithm in [1] fails for large enough gains in a nonsymmetric network, when no delays are present (omitted from the paper due to length limits). This shows the importance of bound (16) which gives an upper bound on the gain k_r, even in the absence of delays.

IV. Window based TCP-like algorithms

We now investigate how proposition 1 can be interpreted within a TCP like algorithm, where the source maintains a window `cwnd` of sent but not yet acknowledged packets. We parameterize the control law at the source, such that the window is incremented by $a\texttt{cwnd}^n$ for each unmarked acknowledgement and decremented by $b\texttt{cwnd}^m$ for each marked acknowledgement (assume $m > n$). TCP in its current form, as given in [9], uses increments of $1/\texttt{cwnd}$ and decrements of $\texttt{cwnd}/2$. Neglecting dynamics at the links, a continuous approximation of the TCP algorithm is described in [10]. It takes the form below :

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

where x_r is a continuous time approximation of the sending rate, and

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

where q_r is now interpreted as the probability a packet is marked along route r. In terms of x_r this becomes

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

When the marking probabilities are small, $q_r(t)$ takes the form of a weighted aggregate price which does not affect the proof of proposition 1, since row sums in the open loop transfer function can still be scaled in the same way (see [4] for the detailed analysis). Hence, the same stability conditions hold. Linearizing about the equilibrium and taking Laplace transforms we get the following transfer function

$$\delta\bar{x}_r = -\frac{a\hat{x}_r(\hat{x}_rT_r)^n}{\hat{q}_r}\frac{1}{(sT_r + a'_r)}\delta\bar{q}_r$$
$$\text{where} \quad a'_r = b(m-n)(\hat{x}_rT_r)^m\hat{q}_r > 0$$

Comparing with (12), bounds (15-16) take the form

$$a(\hat{x}_rT_r)^n < \frac{1}{B} \qquad \forall r \tag{22}$$
$$a < b(m-n)(\hat{x}_rT_r)^{m-n}\hat{q}_r\frac{1}{B|\lambda(S)|} \qquad \forall(l,r)\in A \tag{23}$$

As mentioned above, in current TCP $n = -1$ and $m = 1$, so both bounds illustrate that large equilibrium windows favour stability, but make TCP sluggish in its response. This is the reason why $n = 0$ is being proposed in [4] and investigated through practical implementation in [11], as a scalable TCP variant. An interesting observation is that bound (23) which is due to the nonsymmetry, tends to be violated for small values of b and m which correspond to the decrease rule. It might be tempting to reduce b, as this reduces the coefficient of variation in the sending rate (see [12]). Nevertheless this turns out to reduce robust stability to protocol non symmetries.

V. Conclusions

We have proved that the widely used optimization models for Internet congestion control are not fragile with respect to the symmetry of the underlying protocol. This was established by deriving scalable decentralized conditions for the stability of non symmetric TCP like protocols in arbitrary interconnections. These conditions preserve the bounds for symmetric protocols and are extended by a condition that affects only non symmetric paths and scales with the degree of non symmetry. Finally, we show that a more aggressive decrement rule in TCP favours robust stability to deviations from the protocol symmetry.

APPENDIX

Lemma 1: Let $R \in \mathbb{C}^{m\times n}$ satisfy $\rho(|R|^T|R|) \leq 1$ and $\bar{R}$ be such that

$$\bar{R}_{ij} = R_{ij} \text{ for all } (i,j) \notin A$$
$$\text{where } A = \{(i,j) : R_{ij} \neq 0, \bar{R}_{ij} = 0\}$$

and also $\forall i \in \{1,2,\ldots,m\} \quad \exists j \in \{1,2,\ldots,n\}$ s.t.$\bar{R}_{ij} \neq 0$

Then given $F = \mathrm{diag}(f_1,\ldots,f_n),\ \ f_i \in \mathbb{C}$, we can bound the spectrum of $FR^*\bar{R}$ as follows:

$$\sigma(FR^*\bar{R}) \subset Co\left(0 \cup \{f_l : (i,l) \notin A \quad \forall i \in \{1,2,\ldots,m\}\} \cup \{f_kS : \exists i \in \{1,2,\ldots,m\} \text{ s.t. } (i,k) \in A\}\right)$$
$$\text{where } \ S = \{(v_1^* + v_2^*)v_1 : |v_1| + |v_2| = 1, \quad v_1, v_2 \in \mathbb{C}\}$$

Proof:

$$\sigma(FR^*\bar{R}) = \sigma(\bar{R}FR^*)$$

if we ignore zero eigenvalues. This is not a problem since the bounding region in Lemma 1 always includes zero.

$$\rho(|R||R^T|) \leq 1 \Rightarrow \quad v^*|R||R|^T v \leq 1 \ \forall v \in \mathbb{C}^m \text{ s.t. } v^*v = 1$$
$$\text{since} \quad \rho(|R||R|^T) = \||R|\|_2^2 = \sup_{v\in\mathbb{C}^m, v\neq 0} \frac{\||R|^T v\|_2^2}{\|v\|_2^2}$$

expanding $|R||R^T|$ we get

$$\sum_j (|v_1|R_{1j}| + v_2|R_{2j}| + \ldots|)^2 \leq 1 \ \forall v \in \mathbb{C}^m \text{ s.t. } v^*v = 1$$

And since this is true for all such v

$$\sum_j (|v_1 R_{1j}| + |v_2 R_{2j}| + \ldots)^2 \leq 1 \ \forall v \in \mathbb{C}^m \text{ s.t., } v^*v = 1 \quad (24)$$

We then bound the spectrum with the field of values of the corresponding matrix.

$$\sigma(\bar{R}FR^*) \subset F(\bar{R}FR^*) := \{v^*\bar{R}FR^*v : v \in \mathbb{C}^m \ \ v^*v = 1\} \quad (25)$$

For convenience in the analysis we consider first an $\bar{R}$ that is identical to R apart from a single non-zero element R_{pq} which vanishes i.e $\bar{R}_{pq} = 0$. The result can be then easily generalized to the form in the Lemma 1, where multiple elements of R vanish.

$$\begin{aligned} v^*\bar{R}FR^*v &= \sum_k f_k(v^*\bar{R}_{\bullet k}R^*_{k\bullet}v) \\ &= \sum_{k\neq q} |v^*R_{\bullet k}|^2 f_k + f_q\left(\sum_{i=1,i\neq p}^m v_i^* R_{iq}\right)\left(\sum_{i=1}^m v_i R_{iq}^*\right) \\ &= \sum_{k\neq q} |v^*R_{\bullet k}|^2 f_k + \\ &\quad \left(\sum_i |v_i R_{iq}|\right)^2 f_q \frac{\left(\sum_{i=1,i\neq p}^m v_i^* R_{iq}\right)\left(\sum_{i=1}^m v_i R_{iq}^*\right)}{\left(\sum_i |v_i R_{iq}|\right)^2} \end{aligned} \quad (26)$$

The complex number multiplying $\left(\sum_i |v_i R_{iq}|\right)^2 f_q$ turns out to belong to a set which is invariant with respect to the size of R and v as well as the number of vanishing terms on the same column of R, as this is shown in Lemma 2.

Lemma 2: Let

$$S_{n,A(n)} := \left\{ \left(\sum_{i=1}^n v_i\right)^* \left(\sum_{i\notin A(n)} v_i\right) : v_i \in \mathbb{C}, \sum_{i=1}^n |v_i| = 1 \right\}$$
$$S := \{(v_1^* + v_2^*)v_1 : v_1, v_2 \in \mathbb{C}, \ |v_1| + |v_2| = 1\}$$
$$\text{where} \quad A(n) \subset \{1,2,\ldots,n\}$$

Then $S_{n,A(n)} = S$ for all $n \geq 2, A(n)$.

Proof: Note first that $(v_1^* + v_2^*)v_1 = |v_1|^2 + v_1 v_2^*$. Hence

$$S = \{r^2 + r(1-r)e^{j\theta} : r \in [0,1], \theta \in [0,2\pi]\}$$

So S is an uncountable union of circles on the complex plane, which are symmetric about the real axis and with centre in $[0,1]$. As r varies from 1 to 0, the radii of the circles vary continuously from 0, vanishing again only at $r = 0$, and their centre also translates continuously along the real line from 1 to 0. S is hence a simply connected set. More formally the continuous map

$$h : [0,1] \times [0,2\pi] \to S \quad ; \quad h(r,\theta) = r^2 + r(1-r)e^{j\theta}$$

is a homotopy and each circle in S

$$h_r : [0,2\pi] \to S \quad ; \quad h_r(\theta) = h(r,\theta)$$

is null-homotopic to the constant curves $h_1(\theta)$ and $h_0(\theta)$.

The boundary of S can easily be derived analytically (derivation is omitted here due to space limits). A point z on the boundary satisfies

$$Y^2 = \frac{(1+2X)^3}{27} - X^2, \quad \text{where} \quad X = \Re(z), \quad Y = \Im(z)$$

It crosses the real axis at $X = -1/8, 1$ and for $X \in [-1/8, 1]$, (X,Y) defines the boundary of a convex set which is symmetric about the real line (see figure 6). The fact that S is simply connected, convex and includes zero implies that $\mu S \subset S$ for $\mu \in [0,1)$. Hence

$$\{(v_1^* + v_2^*)v_1 : v_1, v_2 \in \mathbb{C}, \ |v_1| + |v_2| = \mu, \ \mu \in [0,1)\} \subset S \quad (27)$$

We now continue with the proof of lemma 2. Let

$$k = \sum_{i\in A(n)} v_i, \quad l = \sum_{i\notin A(n)} v_i, \qquad \text{then}$$

$$\begin{aligned} S_{n,A(n)} &= \left\{ \left(\sum_{i=1}^n v_i\right)^* \left(\sum_{i\notin A} v_i\right) : v_i \in \mathbb{C}, \sum_{i=1}^n |v_i| = 1 \right\} \\ &= \left\{ (k+l)^* l : v_i \in \mathbb{C}, \sum_{i=1}^n |v_i| = 1 \right\} \end{aligned}$$

$$\text{but since} \quad |k|+|l| = \left|\sum_{i\in A(n)} v_i\right| + \left|\sum_{i\notin A(n)} v_i\right| \le \sum_i |v_i| = 1$$

So $S_{n,A(n)} \subseteq S$ from (27).
Conversely, let $S = \{(k^*+l^*)k : k,l \in \mathbb{C},\ |k|+|l| = 1\}$. Then we could find v_i such that

$$k = \sum_{i\in A(n)} v_i, \quad l = \sum_{i\notin A(n)} v_i, \quad \sum_i |v_i| = 1$$

So $S \subseteq S_{n,A(n)}$ and thus $S = S_{n,A(n)}$. ■

Therefore continuing from (26)

$$\begin{aligned}\sigma(\bar{R}FR^*) &\subset \sum_{k\neq q} |v^* R_{\bullet k}|^2 f_k + \left(\sum_i |v_i R_{iq}|\right)^2 f_q S \\ &\subset \left(\sum_{k\neq q} |v^* R_{\bullet k}|^2 + \left(\sum_i |v_i R_{iq}|\right)^2\right) \\ &\quad Co\left(\{f_k, f_q S : k \in \{1,2,\ldots,q-1,q+1,\ldots,n\}\}\right) \\ &\subset Co\left(\{f_k, f_q S : k \in \{1,2,\ldots,q-1,q+1,\ldots,n\}\}\right)\end{aligned}$$

where the last inclusion follows from (24). This result can be trivially generalized to the case where R and $\bar{R}$ differ with more than one elements, to give the form in Lemma 1. This is because the perturbation set S is the same for all f_i 'behaving' in a non symmetric way and independent of the number of elements deleted along the same column of R, as this was proved in Lemma 2. ■

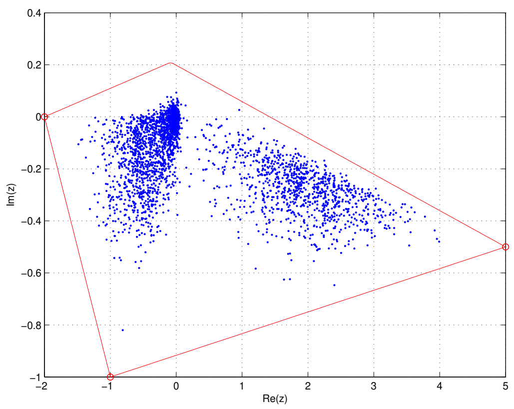

Fig. 5. An example illustrating Lemma 1.

Figure 5 shows an example of Lemma 1. $R \in \mathbb{C}^{3\times 3}$ with $R_{12} = 0$. $\bar{R}$ is equal to R apart from R_{22} which vanishes. $F = \text{diag}(f_1, f_2, f_3)$ with $f_1 = -1, f_2 = -1-j, f_3 = 5-0.5j$. We plot the eigenvalues of 1000 random matrices $FR^*\bar{R}$. The boundary of the hull $Co(\{f_1, f_3, f_2 S\})$ is also plotted and shown to include the eigenvalues, even though $Co(\{0, f_1, f_2, f_3\})$, which is a bound for $\sigma(FR^*R)$ (this is proved in [3]), does not.

Remark 3: If instead of eliminating one element of R we scale it by a factor $\mu \in [0,1]$ so as to form $\bar{R}$, Lemma 1 is still true but with the perturbation set S being replaced by

$$S(\mu) := \{(v_1^* + v_2^*)(\mu v_1 + v_2) : v_1, v_2 \in \mathbb{C},\ |v_1|+|v_2| = 1\}$$

and $I_{((l,r)\in A)}$ in condition (14) replaced by $(1-\mu)I_{((l,r)\in A)}$. Note that the corresponding area enclosed by the boundary of $S(\mu)$ shrinks such that $S(\mu) \to [0,1]$, as $\mu \to 1$ (see figure 6). In a similar way to Lemma 2, we can prove invariance of $S(\mu)$ to the size of R. Furthermore, if more than one elements of a column of R are being scaled, it is conjectured that the corresponding perturbation set only depends on the smallest and largest scaling factors.

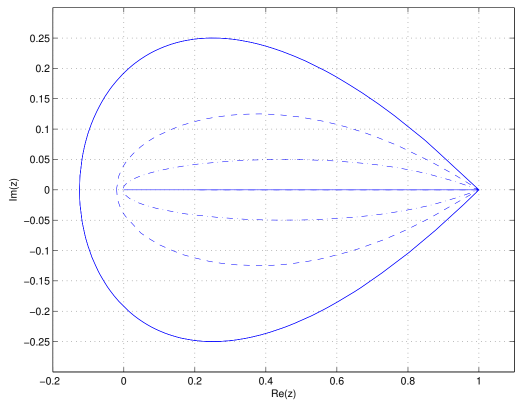

Fig. 6. The boundary of $S(\mu) = \{(v_1^* + v_2^*)(\mu v_1 + v_2) : v_1, v_2 \in \mathbb{C},\ |v_1| + |v_2| = 1\}$ for $\mu = 0, 0.5, 0.8, 1$. Observe that $S(\mu) \to [0,1]$, as $\mu \to 1$.

References

[1] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, no. 3, pp. 237–252, March 1998.

[2] R. Johari and D. Tan, "End-to-End Congestion Control for the Internet: Delays and Stability," *IEEE/ACM Transactions on Networking*, vol. 9(6), pp. 818–832, 2001.

[3] G. Vinnicombe, "On the stability of end-to-end congestion control for the Internet," Cambridge University Engineering Department, Tech. Rep. CUED/F-INFENG/TR.398, 2000.

[4] ——, "On the stability of networks operating TCP-like congestion control," in *IFAC*, 2002.

[5] F. Paganini, J. Doyle, and S. Low, "Scalable laws for stable network congestion control," in *Proceedings of Conference on Decision and Control*, December 2001.

[6] S. Mason, "Feedback theory: Some properties of signal flow graphs," in *Proc. IRE*, vol. 41, Sep. 1953, pp. 1144–1156.

[7] ——, "Feedback theory: Further properties of signal flow graphs," in *Proc. IRE*, vol. 44, July 1956, pp. 920–926.

[8] C. A. Desoer and Y. Yang, "On the generalized Nyquist stability criterion," *IEEE Transactions on Automatic Control*, vol. 25, pp. 187–196, 1980.

[9] V. Jacobson, "Congestion avoidance and control," *ACM Computer Communication Review; Proceedings of the Sigcomm '88 Symposium in Stanford, CA, August, 1988*, vol. 18, 4, pp. 314–329, 1988.

[10] F. Kelly, "Mathematical modelling of the internet," in *Bjorn Engquist and Wilfried Schmid (Eds.), Mathematics Unlimited – 2001 and Beyond@ Springer*, 2001.

[11] T. Kelly, "On engineering a stable and scalable TCP variant," Cambridge University Engineering Department, Tech. Rep. CUED/F-INFENG/TR.435, 2002.

[12] A. Misra and T.J.Ott, "Jointly Coordinating ECN and TCP for rapid adaptation to varying bandwidth," in *Proceedings of IEEE MILCOM*, 2001.