Probability in the Engineering and Informational Sciences, **23**, 2009, 675–698. Printed in the U.S.A.
doi:10.1017/S0269964809990076

STOCHASTICALLY SCALABLE FLOW CONTROL

THOMAS VOICE
Cambridge University Statistical Laboratory
Cambridge, UK
E-mail: tdvoice@gmail.com

Recent advances in the mathematical analysis of flow control have prompted the creation of the *Scalable TCP* (STCP) and *Exponential RED* (E-RED) algorithms. These are designed to be *scalable* under the popular deterministic *delay stability* modeling framework. In this article, we analyze stochastic models of STCP and STCP combined with E-RED link behavior. We find that under certain plausible network conditions, these probabilistic models also exhibit scalable behavior. In particular, we derive parameter choice schemes for which the equilibrium coefficients of variation of flow rates are bounded, however large, fast, or complex the network. Our model is shown to exhibit behavior similar to the mean field convergence that has recently been observed in TCP.

1. INTRODUCTION

The Internet has grown in size and speed by several orders of magnitude since its inception, and it continues to grow exponentially over time. Much research has been carried out in order to try to understand how to regulate the flow of traffic over such large and expanding networks. Since there are limitations on what systems can be simulated or experimented on, mathematical modeling and analysis of flow control has become an important area of interest (see Srikant [21]). Recent advances in this field combined with problems with current Internet TCP—in particular, its slow rate of adaptation for high speed connections—have prompted the creation of the TCP variant, *scalable TCP* (STCP) (see Kelly [8], Kelly [11], and Vinnicombe [22]), and the active queue management (AQM) scheme, *exponential RED* (E-RED) (see Liu, Başar, and Srikant [13]).

Both STCP and E-RED are designed to be *scalable* within one of the most popular mathematical frameworks for analysis, *delay stability* (see, e.g., Johari and Tan [7], Kelly, Maulloo, and Tan [9], Massoulié [14], and Vinnicombe [22]), that is,

© 2009 Cambridge University Press 0269-9648/09 $25.00
https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

decentralized parameter choice schemes exist that ensure local stability for both STCP with static link responses (found by Kelly [11]) and E-RED with generalized TCP-like source behavior (found by Liu et al. [13]). Furthermore, these schemes apply to networks with arbitrary topology, capacity, and propagation delays. STCP and E-RED also have attractive properties within an optimization and fairness framework. In particular, under STCP with static link responses, when flow rates are in equilibrium, they are approximately in a weighted proportionally fair distribution. When E-RED is combined with STCP, then the equilibrium, flow rate distribution is precisely weighted proportionally fair (see Kelly et al. [9]). Thus, STCP and E-RED are potentially strong candidates for future large network flow control.

However, these analyses have only made use of deterministic fluid models, in which continuous time dynamical systems represent the behavior of *average* traffic flow rates. The average here is meant in an implicitly random or pseudorandom sense. If a congestion control is to be fully scalable, then the *variances* of traffic flow rates must remain manageable, regardless how large or fast the network is. In this article, we study the stochastic behavior STCP and E-RED, looking for bounds on flow rate variance that scale favorably as networks become larger, faster, and more complex. For the sake of tractability, we limit ourselves to studying the equilibrium distributions of flow rates. For a scalable algorithm, it is desirable that the steady state variance of a flow rate is approximately proportional to the square of its mean. This is equivalent to the equilibrium coefficients of variation of flow rates being approximately equal to a constant that only depends on algorithm parameters, not on algorithm state. This is the property that we will refer to as *stochastic scalability*. So far, the issue of stochastic scalability for large networks has remained relatively unstudied. In the literature, various stochastic models of flow control have been analyzed, and corresponding steady state flow rate distributions have been examined. For examples see, Baccelli, McDonald, and Reynier [1], Kelly et al. [9], McDonald and Reynier [15], Ott [17,18], and Ott and Swanson [19]. However, much of this work has focused on modelling a network that consists of a single link. In this article, we will examine the equilibrium variance of flow rates in networks of arbitrary topology. We analyze networks operating STCP with static link responses, in keeping with small buffer behavior as analyzed by Raina, Wischik and Towsley [20], and STCP with large buffer links using the E-RED scheme. Following Kelly et al. [9], the model we use is a system of linear stochastic differential equations, approximating the evolution of the vector of traffic flow rates as a multidimensional Ornstein–Uhlenbeck process. We sketch a proof of a convergence result as justification for this model choice.

The E-RED AQM scheme as proposed by Liu et al. [13] is compatible with a number of possible source behaviors; however, we will only be considering that of STCP. This is not only because of its excellent deterministic properties but also because, under a detailed Markov model of a single source, the equilibrium flow rate variances of STCP were found by Ott [17] to be approximately proportional to the squares of their means.

From our analysis, we discover attractive stochastic scalability properties for both STCP with static link responses and STCP with E-RED. In particular, for both

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

algorithms, we derive bounds that show that the equilibrium coefficients of variation of flow rates are close to 0.5 times the square root of an algorithm parameter. These bounds become tighter as the number of flows through each link becomes large. Thus, under a regime in which flows are numerous and diverse, these algorithms are stochastically scalable. This requirement is in keeping with the mean field-theoretic models of TCP used by Baccelli et al. [1] and by McDonald and Reynier [15]. Indeed, our results complement those two articles, in which it was shown that for a single link model of TCP, in the mean field limit, packet dropping probabilities become deterministic. When a network is in steady state, this corresponds to packet dropping probabilities becoming constant. Our results show a similar convergence in our models and, furthermore, give a lower bound on the rate at which it occurs.

Despite all of the positive results regarding the scalability of STCP, there is a problem with STCP under the small buffer static link response regime. In the case of high-speed small buffer networks, many proposed TCP variants, including STCP, were shown by Gu, Towsley, Hollot, and Zhang [5] to underutilize links. The solution proposed by Gu et al. [5] is a new TCP variant called E-TCP, which may be seen as an extension of STCP. It not only possesses the same attractive qualities of STCP discussed earlier, but also does not underutilize high-speed networks. We will discuss this further and explain how our main results relate to E-TCP in Section 3.2.

The rest of this paper is organised as follows: In Section 2 we give a description of STCP source behaviour, and the stochastic processes we shall be using to model it. We sketch a proof of a limiting convergence result to justify this model. We then describe the deterministic dynamic equations we use to describe small buffer links and large buffer links operating E-RED. Section 3 will contain the main results of this paper, which are scalable bounds on the variances of traffic flow rates in equilibrium. This section also contains a rough outline and an extension of our results to cover networks operating E-TCP. We shall draw our conclusions in Section 4.

2. MODELING SCALABLE TCP

In this section we introduce the mathematical systems that we use to model Scalable TCP and motivate our approach. We begin with a brief description of the Scalable TCP protocol, for more detail, see Kelly [11].

2.1. Scalable TCP

Consider a source transmitting data through a network operating Scalable TCP. The data are transmitted as a series of small pieces called *packets*. Once a packet reaches its destination, a much smaller *acknowledgment packet* is sent back to its origin. The time from the sending of a packet to the receipt of an acknowledgment is the *round-triptime* for that route. The source maintains a variable, $\mathbf{w}$, called the *congestion window*, which is precisely the number of packets the source transmits every round-triptime. Whenever an acknowledgment packet is received, the source increases $\mathbf{w}$

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

by adding a constant a. If the network is congested, router buffers can overflow and packets could be lost. Whenever a packet is lost (i.e., the source fails to receive an acknowledgment, the source decreases $\mathbf{w}$ by subtracting $b\mathbf{w}$, for some constant b). For the particular scheme proposed by Kelly [11], $a = 0.01$ and $b = 0.125$.

Routers can also operate AQM schemes, which notify sources of congestion before their buffers overflow. This could be done by routers randomly dropping packets with probability equal to a function of the buffer queue length. Alternatively, packets can carry an explicit congestion notification (ECN) bit. This ECN bit is set to zero on transmission and is set to 1, or *marked*, as it passes through a congested router. Since both methods are modeled the same way, we will refer to them collectively as *packet marking*.

2.2. Network Model

There are many different approaches to modeling TCP and its variants, but they all use a similar framework for modeling the network topology. We will also use such a framework.

We consider a communications network consisting of an interconnection of a set of sources, $\mathbb{S}$, with a set of links $\mathbb{J}$. A link $l \in \mathbb{J}$ represents some kind of physical link in the network (e.g., a cable, a wireless connection, or a router). A source $r \in \mathbb{S}$ represents an ordered pair of physical nodes in the network, called an origin–destination pair, and a user who wishes to transmit information from the origin node to the destination node.

Associated with each source is a *route*, each route being a sequence of links through which packets from that source are flowing. If a source r transmits along a link l, then we write $l \in r$. Later we will have use of the connectivity matrix R, which is defined to have entries $R_{r,l} = 1$ if $l \in r$ and $R_{r,l} = 0$ otherwise. We will use the fairly weak assumption that R has full column rank, this can be guaranteed by each link having at least one source that only passes through that link.

In our model each source r has associated with it a congestion window size $\mathbf{w}_r(t) \geq 0$, which we formally define as the number of packets transmitted by r during the interval $[t - T_r, t]$, for T_r, the round-triptime of r. We will also require the *flow rate* $\mathbf{x}_r(t)$. Here, the flow rate, $\mathbf{x}_r(t)$, is the number of packets transmitted by r per unit time, which we formally define as $\mathbf{x}_r(t) = \mathbf{w}_r(t)/T_r$.

Each link l has associated with it a packet marking probability $\mathbf{p}_l(t) \geq 0$, which varies over time, depending on the total flow rate through that link $\mathbf{z}_l(t)$. Here, we will approximate

$$\mathbf{z}_l(t) = \sum_{r:l\in r} \mathbf{x}_r(t). \tag{1}$$

In [20] it was found that links with relatively small buffers could be modeled statically, with $\mathbf{p}_l(t)$ set equal to a function of $\mathbf{z}_l(t)$. However, for links with large buffers, it was found that because of the dynamics of the buffer queue length, $\mathbf{p}_l(t)$ could be better

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

modeled with an evolving differential equation. The exact nature of this equation is determined by what, if any, AQM is in operation at l.

Each source r experiences a total packet marking probability, $\mathbf{q}_r(t)$. We will assume that packet marking probabilities are small and thus $\mathbf{q}_r(t)$ may be approximated by

$$\mathbf{q}_r(t) = \sum_{l \in r} \mathbf{p}_l(t). \tag{2}$$

Note that when packet dropping is used for congestion feedback, as in the current Internet, $\mathbf{q}_r(t)$ should always be kept small for the sake of efficiency.

2.3. Linear Stochastic Differential Equation Model

The main results of this article rely on a multidimensional Ornstein–Uhlenbeck model of network flow behavior. In this subsection, we sketch the proof of a motivating convergence result. For this discussion we only deal with the isolated one-dimensional case, examining a single traffic source operating Scalable TCP. We consider the situation in which this source experiences a constant packet marking probability p and has an approximately constant round-triptime, T. Packets are marked randomly, as a Bernouilli source, independent of the state of the algorithm. This is precisely the case modeled by Ott's work [17–19] and is also the situation in equilibrium for the limit of the mean field model of TCP derived by Baccelli et al. [1] and by McDonald and Reynier [5]. In Ott's work [17–19], the behavior of TCP-like congestion control algorithms was examined for the limiting case when the packet marking probability becomes very small. However, the convergence we explore occurs as the values of a and b tend to zero.

For a source operating under these assumptions, for $n = 1, 2, \ldots$ we let $\mathbf{W}_n$ be the congestion window size on receiving the nth update, where updates occur on receiving an acknowledgment or detecting a drop. So, we can approximate

$$\mathbf{W}_{n+1} = \mathbf{W}_n + (a(1-p) - bp\mathbf{W}_n) + (a + b\mathbf{W}_n)\mathbf{b}_n, \tag{3}$$

where $\mathbf{b}_n = p$ if there is a successful acknowledgment at the nth feedback and $\mathbf{b}_n = p - 1$ otherwise. Thus, $\mathbb{E}(\mathbf{b}_n) = 0$ for all n. Let $\mathbf{n}(t)$ be the number of updates that have occurred by time t. So, $\mathbf{w}(t) = \mathbf{W}_{\mathbf{n}(t)}$.

We take the ratio $\alpha = a/b$ to be fixed and study what happens as b becomes small (for later convenience we always assume $b < 0.5$). We seek to show that the process $\mathbf{v}(t) = (\mathbf{w}(t/b) - w)b^{-(1/2)}$ converges to the solution of a linear stochastic differential equation, where w is the equilibrium point $w = a(1-p)/bp$. So we will assume that the distribution of $\{\mathbf{v}(t) : t \leq 0\}$ is fixed as b becomes small. Furthermore, we initially make the assumption that $\sup_{t \leq 0} |\mathbf{v}(t)|$ is bounded and consider the evolution of $\mathbf{v}(t)$ over a fixed time interval $[0, \tau]$. This means that we are considering the behavior of $\mathbf{w}(t)$ for $t \in [0, \tau/b]$, with initial conditions that mean that for $t \leq 0$, $\mathbf{w}(t) - w$ is bounded above and below, with bounds that are $O(\sqrt{b})$.

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

From these assumptions we can show that $\mathbf{w}(t)$ is bounded above and away from zero, and these bounds are independent of b. To prove this, we begin first with an upper bound of $\mathbf{w}(t)$. The greatest value that $\mathbf{w}(t)$ can take occurs if $\mathbf{w}(\cdot)$ increases with every feedback update and $\mathbf{w}(0) = \sup_{u\in[-T,0]} \mathbf{w}(u)$. Then, as $\mathbf{w}(t)$ is increasing, we have

$$\mathbf{n}(t) \leq \int_0^t \frac{\mathbf{w}(u)}{T}\, du. \tag{4}$$

Since a constantly increasing $\mathbf{w}$ yields $\mathbf{w}(t) = \mathbf{w}(0) + a\mathbf{n}(t)$, we derive a general upper bound of

$$\mathbf{w}(t) \leq e^{a\tau/bT} \sup_{u\in[-T,0]} \mathbf{w}(u) = e^{\alpha\tau/T} \sup_{u\in[-T,0]} \mathbf{w}(u) \leq K \tag{5}$$

for some constant K. From this we also get that for all $t \in [0, \tau/b]$, $\mathbf{n}(t) \leq \tau K/bT$. The lower bound for $\mathbf{w}(t)$ follows from observing that $\log(1-b) \geq -2b$ for $b < 0.5$ and so

$$\mathbf{w}(t) \geq \mathbf{w}(0)(1-b)^{\mathbf{n}(t)} \geq \mathbf{w}(0)e^{-2\tau K/T}. \tag{6}$$

We now consider bounds on $\mathbf{v}(t)$. From Eq. (3) and the fact that $a(1-p) = bpw$, we get that if there is a single update between times $t < t'$,

$$\mathbf{v}(bt') = (1-bp)\mathbf{v}(bt) + \sqrt{b}\big(\alpha + \mathbf{W}_{\mathbf{n}(t)}\big)\mathbf{b}_{\mathbf{n}(t)}. \tag{7}$$

This means, for a general t, that $\mathbf{v}(t)$ is equal to

$$(1-bp)^{\mathbf{n}(t/b)}\mathbf{v}(0) + \sum_{i=0}^{\mathbf{n}(t/b)} (1-bp)^{\mathbf{n}(t/b)-i}\sqrt{b}(\alpha + \mathbf{W}_i)\mathbf{b}_i. \tag{8}$$

Note, that for any $n \leq \tau K/bT$, $(1-bp)^n \geq e^{-2p\tau K/T}$. For $n = 1, 2, \ldots, \lfloor \tau K/bT \rfloor$, we define

$$\mathbf{M}_n = (1-bp)^{-i}\sqrt{b}\sum_{i=0}^{n}\big(\alpha + \mathbf{W}_i\big)\mathbf{b}_i. \tag{9}$$

Since the value of $\mathbf{W}_i$ is independent of $\mathbf{b}_i$, $\mathbf{M}$ is a martingale and the covariance of M_n and $M_{n+1} - M_n$ is zero. Hence, by induction, $\mathrm{Var}(\mathbf{M}_n)$ is always less than or equal to $e^{2p\tau K/T} nb(\alpha(1-p) + pK)^2$. From Doob's L_2 maximal martingale inequality, we get

$$\Big\| \sup_{t\in[0,\tau/b]} |\mathbf{v}(t)| \Big\|_2 \leq \|\mathbf{v}(0)\|_2 + 2e^{2p\tau K/T}\frac{\tau K}{T}(\alpha(1-p) + pK)^2. \tag{10}$$

This means that $\| \sup_{t\in[0,\tau/b]} |\mathbf{w}(t) - w| \|_2$ is $O(\sqrt{b})$ and tends to zero as b gets smaller.

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

Let us define the process $\mathbf{N}_1, \mathbf{N}_2, \ldots$ to be such that

$$\mathbf{N}_n = (1 - bp)^n \mathbf{v}(0) + \sum_{i=0}^{n} (1 - bp)^{n-i} \sqrt{b}(\alpha + w)\mathbf{b}_i. \tag{11}$$

As we have shown for $\mathbf{M}_n$, by induction, the random variables

$$\sum_{i=0}^{n} (1 - bp)^{-i} \sqrt{b}(\mathbf{W}_i - w)\mathbf{b}_i \quad \text{for } n = 1, 2, \ldots, \lfloor \tau K/bT \rfloor \tag{12}$$

form a martingale that, for each n, has expectation 0 and variance less than

$$\sum_{i=0}^{n} (1 - bp)^{-i} b\mathrm{Var}(\mathbf{W}_i) \le nbe^{2p\tau K/T} \left\| \sup_{t \le \tau/b} |\mathbf{w}(t) - w| \right\|_2^2. \tag{13}$$

The right-hand side of Eq. (13) is $O(nb^2) = O(b)$, and so, by Doob's L_2 maximal martingale inequality, we get that

$$\left\| \sup_{t \in [0,\tau]} |\mathbf{v}(t) - \mathbf{N}_{\mathbf{n}(t/b)}| \right\|_2 = O(\sqrt{b}). \tag{14}$$

It can be shown from the central limit theorem that as b tends to zero, $\mathbf{N}_{\lfloor wt/bT \rfloor}$ will converge in distribution to

$$e^{-pwt/T} \mathbf{v}(0) + \int_0^t e^{-pw(t-u)/T} (\alpha + w) \sqrt{p(1-p)w/T}\, d\mathbf{B}(u), \tag{15}$$

where $\mathbf{B}(t)$ is a Brownian motion.

From the definition of window sizes, for $t < t'$, $\mathbf{n}(t)$ is bounded above by $\sup_{u \le \tau/b} \mathbf{w}(u) \lceil (t' - t)/T \rceil$ and bounded below by $\inf_{u \le \tau/b} \mathbf{w}(u) \lfloor (t' - t)/T \rfloor$. Thus,

$$\left\| \mathbf{n}(t'/b) - \mathbf{n}(t/b) - \frac{w(t' - t)}{bT} \right\|_2 \le \left(\frac{t' - t}{bT} + 1 \right) \left\| \sup_{t \le \tau/b} |\mathbf{w}(t) - w| \right\|_2, \tag{16}$$

which is $O(1/\sqrt{b})$. This means that $b(\mathbf{n}(t'/b) - \mathbf{n}(t/b))$ converges in L_2 to $w(t' - t)/T$.

This is enough to prove the convergence in distribution of $\mathbf{v}(t)$ to Eq. (15); we have that $\mathbf{v}(t)$ tends to $\mathbf{N}_{\mathbf{n}(t/b)}$, and with high probability we can bound $b\mathbf{n}(t/b)$ to be close to wt/T for any t. With the continuity properties of Eq. (15), this is enough to prove convergence of finite-dimensional distributions. The tightness of $\mathbf{v}(t)$ follows from the fact that $\| \sup_{t \le \tau} |\mathbf{v}(t)| \|_2$ is bounded, and so convergence in distribution follows.

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

We have assumed that $\sup_{t\le 0}|\mathbf{v}(t)|$ is bounded in order to derive this result. In general, provided $\sup_{t\le 0}|\mathbf{v}(t)|$ has finite expectation, the same convergence in distribution holds. This can be shown using the previous result and Chebyshev's inequality.

Thus, we have sketched the proof that as b tends to zero, $\mathbf{v}(t)$ converges in distribution to the solution to the linear stochastic differential equation

$$d\mathbf{v}(t) = -p\frac{w}{T}\mathbf{v}(t)\,dt + (\alpha + w)\sqrt{p(1-p)w/T}\,d\mathbf{B}(t). \tag{17}$$

This leads us to the following linear stochastic differential equation (SDE) model for $\mathbf{w}(t)$:

$$d\mathbf{w}(t) = \frac{w}{T}\big(a(1-p) - bp\mathbf{w}(t)\big)\,dt + (a + bw)\sqrt{p(1-p)w/T}\,d\mathbf{B}(t). \tag{18}$$

Note that we have written the deterministic part of Eq. (18) so that it agrees with Eq. (3) and makes explicit the dependence on p.

The steady state distribution of $\mathbf{w}(t)$ under Eq. (18) is Gaussian with variance

$$\mathrm{Var}(\mathbf{w}(t)) = \frac{b}{2}(\alpha + w)^2(1-p) = bw^2\frac{1}{2(1-p)}. \tag{19}$$

2.4. Link Behavior

Motivated by the work of Raina et al. [20], we will apply two models to link behavior as follows.

For networks with small buffers at links, we take the link packet marking probability to be a static function of flow through that link. The appropriate probability can be determined by using queuing-theoretic models of the link buffer under fixed flow rates (for more on queuing theory, see Cooper [4]). We approximate this by setting

$$\mathbf{p}_l(t) = \left(\frac{\mathbf{z}_l(t)}{C_l}\right)^{\beta_l} \tag{20}$$

for some constants C_l and β_l. We let $\beta_{\mathbb{J}}$ be $\beta_{\mathbb{J}} = \max_{l\in\mathbb{J}}\beta_l$.

For networks with large buffers at links, dynamic differential equations models of $\mathbf{p}_l(t)$ are more appropriate. These can take many forms, depending on which, if any, AQM scheme is in operation. We consider the E-RED AQM given by Liu et al. [13], which can be modeled by setting, for each $l \in \mathbb{J}$,

$$d\mathbf{p}_l(t) = \kappa_l\mathbf{p}_l(t)\left(\frac{\mathbf{z}_l(t)}{C_l} - 1\right)dt \tag{21}$$

for some constants C_l and κ_l. We let $\kappa_{\mathbb{J}}$ be $\kappa_{\mathbb{J}} = \max_{l\in\mathbb{J}}\kappa_l$.

In both of these models, C_l can be interpreted as the capacity of link l, measured in packets per second. However, under the E-RED scheme, link l maintains its packet

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

marking probability as an internal dynamic variable that is updated according to Eq. (21). The advantage of this is that C_l can take any desired level of bandwidth utilization, and, on average, $\mathbf{z}_l$ will be less than or equal to C_l at equilibrium.

The deterministic part of Eq. (18) is a linearization in $\mathbf{w}$ of a common deterministic model of Scalable TCP:

$$\frac{d\mathbf{w}(t)_r}{dt} = \frac{\mathbf{w}_r(t)}{T_r}\big(a(1-\mathbf{q}_r(t)) - b\mathbf{q}_r(t)\mathbf{w}_r(t)\big); \tag{22}$$

Equations (22) and (20) and Eqs. (22) and (21) form two systems of ordinary differential equations that each have a globally asymptotically stable equilibrium point. For either system we will refer to this equilibrium point as (w,p). Under the system (22) and (20), the choice of p is not unique; however, the choice of w always is. Here, we have ignored the possibility of links that remain underutilized under system (22) and (21), even at system equilibrium. Such links will have marking probabilities that forever tend towards zero and so can be removed from the system without loss of generality. For more details on deterministic modeling and stability proofs, see Kelly et al. [9], Liu, Başar, and Srikant [12], and Srikant [21].

2.5. Linear Network Model

The above exposition suggests that for small b, the steady state vector of window sizes is likely to be close to the deterministic equilibrium w. In order to aid the tractability of our networkwide model, we assume this is true and linearize our models of link behavior about p. The main result in this article, showing the stochastic scalability properties of Scalable TCP, is consistent with this assumption.

Following on from our model discussion, we define $\mathbf{B}$ to be an $|\mathbb{S}|$-dimensional Brownian motion. For each source r and link l we will write x_r, q_r, and z_l for the values of $\mathbf{x}_r$, $\mathbf{q}_r$, and $\mathbf{z}_l$ at (w,p), respectively. From Eq. (18) we define the following stochastic differential equation network model for $\mathbf{w}_r(t)$:

$$\begin{aligned} d\mathbf{w}(t)_r = {} & \frac{w_r}{T_r}\big(a_r - (a_r + b_r\mathbf{w}_r(t))\mathbf{q}_r(t)\big)\,dt \\ & + (a_r + b_r w_r)\left(\mathbf{q}_r(t)(1-\mathbf{q}_r(t))\frac{w_r}{T_r}\right)^{1/2} d\mathbf{B}_r(t). \end{aligned} \tag{23}$$

We linearize our model of STCP with static link responses, Eqs. (20) and (23), about (w,p), taking, for each source r,

$$\mathbf{w}_r(t) = w_r + \sqrt{b_r w_r(\alpha_r + w_r)}\,\mathbf{u}_r \tag{24}$$

to get

$$d\mathbf{u} = A\mathbf{u}\,dt + F\,d\mathbf{B}(t), \tag{25}$$

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

where for all sources, r and s with $r \neq s$,

$$A_{r,r} = -\frac{b_r}{T_r}\left(\alpha_r(1-q_r) + \frac{1}{T_r}w_r(\alpha_r + w_r)\sum_{l\in r}\frac{\beta_l p_l}{z_l}\right), \tag{26}$$

$$A_{r,s} = -\frac{\sqrt{b_r b_s}}{T_r T_s}\left(w_r(\alpha_r + w_r)w_s(\alpha_s + w_s)\right)^{1/2}\sum_{l\in r\cap s}\frac{\beta_l p_l}{z_l}, \tag{27}$$

and

$$F_{r,r} = \left(\frac{b_r}{T_r}\alpha_r(1-q_r)\right)^{1/2}, \quad F_{r,s} = 0. \tag{28}$$

Here, we have used the fact that $w_r q_r = \alpha_r(1-q_r)$, so $(\alpha_r + w_r)q_r = \alpha_r$.

We linearize our model of STCP with E-RED, Eqs. (23) and (21), about (y, p), taking, for each source r and link l,

$$\mathbf{w}_r(t) = w_r + \sqrt{b_r w_r(\alpha_r + w_r)}\mathbf{u}_r(t), \tag{29}$$

$$\mathbf{p}_l(t) = p_l + \left(\frac{\kappa_l p_l}{C_l}\right)^{1/2}\mathbf{v}_l(t) \tag{30}$$

to get

$$\begin{pmatrix} d\mathbf{u} \\ d\mathbf{v}\end{pmatrix} = \acute{A}\begin{pmatrix}\mathbf{u}\\ \mathbf{v}\end{pmatrix}dt + \acute{F}\begin{pmatrix} d\mathbf{B}(t)\\ 0\end{pmatrix}, \tag{31}$$

where for all sources r and links l,

$$\acute{A}_{r,r} = -\frac{b_r}{T_r}\alpha_r(1-q_r), \acute{A}_{r,l} = -\left(\frac{\kappa_l p_l b_r w_r(\alpha_r + w_r)}{C_l}\right)^{1/2}\frac{1}{T_r}R_{r,l}, \tag{32}$$

$$\acute{A}_{l,r} = -\acute{A}_{r,l}, \tag{33}$$

$$\acute{F}_{r,r} = \left(\frac{b_r}{T_r}\alpha_r(1-q_r)\right)^{1/2}, \tag{34}$$

and all other entries of $\acute{A}$ and $\acute{F}$ are zero.

2.6. Ornstein–Uhlenbeck Solution

The behavior of linear stochastic differential equations is well known (see Oksendal [16] for more detail), and the solution of any equation of the form

$$d\mathbf{u} = A\mathbf{u}\,dt + F\,d\mathbf{B}(t), \tag{35}$$

is a multidimensional Ornstein–Uhlenbeck process, centered around zero. The stationary distribution of $\mathbf{u}$ under Eq. (35), is a multivariate Normal distribution with

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

mean 0. The covariance matrix Σ of the system in equilibrium can be calculated analytically (as was done by Kelly et al. for a similar case in [9]). Furthermore, Σ is all that is needed to fully determine the equilibrium distribution, as it is Gaussian.

The real symmetric covariance matrix Σ is

$$\Sigma = \int_{-\infty}^{0} e^{-uA} F F^T e^{-uA^T}\, du. \tag{36}$$

Using integration by parts (a technique used by Bain and Key [2]), we can see that Σ satisfies the equation

$$\Sigma A^T + A\Sigma = -C, \tag{37}$$

where $C = FF^T$.

In the deterministic models, the equilibrium points of the algorithms we consider are asymptotically stable, and, so, A and $\acute{A}$ have eigenvalues with strictly negative real parts. It can be shown that if A has only eigenvalues with negative real parts, then Eq. (37) has Σ as its unique positive solution (for proof, see Barnett [3]). Equation (37) can be used to calculate Σ and Hammarling [6] derived fast accurate numerical methods for this purpose. Thus, Eq. (37) can be used to calculate the covariance matrices of the invariant distributions of Eqs. (25) and (31). These covariance matrices will be referred to as Σ and $\acute{\Sigma}$, respectively.

From now on, we will mainly be considering the case in which the systems of linear stochastic differential equations we describe are in stochastic equilibrium. Accordingly, for $r \in \mathbb{S}$ and $l \in \mathbb{J}$, we will use the notational convention that when we do not make the dependence on t explicit for $\mathbf{y}_r$, $\mathbf{x}_r$, $\mathbf{q}_r$, $\mathbf{p}_l$, and $\mathbf{z}_l$, this means that they are in their stationary distributions.

3. MAIN RESULTS

For any specific network configuration, Eq. (37) can be solved numerically. However, even for moderately large networks, such calculations quickly become intractable. Thus, analytic results that hold for general networks are highly desirable. In this section we derive scalable bounds on the steady state variances of flow rates under Scalable TCP with small buffers and Scalable TCP with E-RED, and we discuss how these results fit in with other scalable properties of these algorithms.

For each $r \in \mathbb{S}$,

$$\mathrm{Var}(\mathbf{w}_r(t)) = b_r w_r(\alpha_r + w_r)\mathrm{Var}(\mathbf{u}_r(t)) = b_r w_r^2 \frac{1}{1-q_r}\mathrm{Var}(\mathbf{u}_r(t)), \tag{38}$$

and we have

$$\mathrm{Var}(\mathbf{x}_r) = b_r x_r^2 \frac{1}{1-q_r}\Sigma_{r,r} \quad \text{or} \quad \mathrm{Var}(\mathbf{x}_r) = b_r x_r^2 \frac{1}{1-q_r}\acute{\Sigma}_{r,r}. \tag{39}$$

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

Thus, if we want our algorithm to be scalable—that is, if we would like the coefficients of variation of $\mathbf{x}_r$ to be independent of the size, speed, or topology of the network—then it is desirable for Σ and $\acute{\Sigma}$ to be approximately invariant.

In order to provide strong results, we will need to restrict our attention to large, high-speed networks; that is, we will assume that the number of flows in the network is large. More specifically, we consider the situation in which there are large numbers of flows passing through each link, requiring that the quantity γ be close to zero, where

$$\gamma = \max_{r\in\mathbb{S}, l\in r} \frac{x_r}{z_l}. \tag{40}$$

We will also assume that

$$\max_{r\in\mathbb{S}} \frac{1}{1-q_r} \leq 2, \tag{41}$$

which is reasonable and is implied by the assumptions that $\alpha_r \leq 1$ and $w_r \geq 1$, both of which are required for normal network operation.

THEOREM 1: *Under the equilibrium distribution of Eq.* (25), *provided that* $\gamma\sqrt{|\mathbb{S}|} < 1/(16\beta_{\mathbb{J}}^2)$, *the following bounds are satisfied, for all* $r \in \mathbb{S}$:

$$b_r x_r^2 \frac{1}{2(1-q_r)}(1-\gamma\sqrt{\mathbb{S}}\zeta_r) \leq \mathrm{Var}(\mathbf{x}_r) \leq b_r x_r^2 \frac{1}{2(1-q_r)}(1+\gamma\sqrt{\mathbb{S}}\zeta_r), \tag{42}$$

where ζ_r *is the constant*

$$\zeta_r = 4\beta_{\mathbb{J}}(1+8\beta_{\mathbb{J}})\left(\frac{b_r\alpha_r(1-q_r)}{T_r}\right)^{-(1/2)} \max_{s\in\mathbb{S}}\left(\frac{b_r\alpha_r}{T_s}\right)^{1/2}. \tag{43}$$

For the proof, see the Appendix.

THEOREM 2: *Under the equilibrium distribution of Eq.* (31), *the following bounds are satisfied, for all* $r \in \mathbb{S}$:

$$b_r x_r^2 \frac{1}{2(1-q_r)}(1-\gamma\sqrt{|\mathbb{S}|}\zeta_r) \leq \mathrm{Var}(\mathbf{x}_r) \leq b_r x_r^2 \frac{1}{2(1-q_r)}(1+\gamma\sqrt{|\mathbb{S}|}\zeta_r), \tag{44}$$

where ζ_r *is the constant*

$$\zeta_r = 4\sqrt{2}\kappa_{\mathbb{J}} \max_{r,s\in\mathbb{S}}\left(\frac{b_r\alpha_r}{T_r}\right)^{-(3/2)}\left(\frac{b_s\alpha_s}{T_s}\right)^{1/2}. \tag{45}$$

For the proof, see the Appendix.

To illustrate the power of these results, let us consider what happens when the number of sources increases uniformly across a network. More formally, we consider a set of links $\mathbb{J}$ and a series of sets of sources $\mathbb{S}(1), \mathbb{S}(2), \ldots$ transmitting over those links. Each set $\mathbb{S}(n)$ is made by replacing every $r \in \mathbb{S}(1)$ with $N_r(n)$ identical copies of

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

r so that $|\mathbb{S}(n)| \to \infty$ as $n \to \infty$. We assume that $N_r(n) = \Omega(|\mathbb{S}(n)|)$ for all $r \in \mathbb{S}(1)$. For each $r \in \mathbb{S}(n)$, let $s(r) \in \mathbb{S}(1)$ be the source of which r is a copy. We allow link capacities to change, so that in the nth network, link j has capacity $C_j(n)$. All other parameters should remain constant. We still assume that $2 \geq 1/(1 - q_r)$ for all $r \in \mathbb{S}(n)$ for all n.

LEMMA 1: *Under the equilibrium distribution of Eq.* (25), *for* $r \in \mathbb{S}(1)$, *for any sequence* $r(n) \in \mathbb{S}(n)$ *with* $s(r(n)) = r$,

$$\mathrm{Var}(\mathbf{x}_{r(n)}) = b_r x_{r(n)}^2 \frac{1}{2(1 - q_{r(n)})} \left(1 + O(|\mathbb{S}(n)|^{-(1/2)})\right) \tag{46}$$

$$\leq b_r x_{r(n)}^2 \left(1 + O(|\mathbb{S}(n)|^{-(1/2)})\right), \tag{47}$$

and for all $j \in \mathbb{J}$,

$$\mathrm{Var}(\mathbf{p}_j) = p_j^2 O(|\mathbb{S}(n)|^{-1}). \tag{48}$$

For the proof, see the Appendix.

LEMMA 2: *Under the equilibrium distribution of Eq.* (31), *for* $r \in \mathbb{S}(1)$, *for any sequence* $r(n) \in \mathbb{S}(n)$ *with* $s(r(n)) = r$,

$$\mathrm{Var}(\mathbf{x}_{r(n)}) = b_r x_{r(n)}^2 \left(1 + O(|\mathbb{S}(n)|^{-(1/2)})\right) \tag{49}$$

$$\leq b_r x_{r(n)}^2 \left(1 + O(|\mathbb{S}(n)|^{-(1/2)})\right) \tag{50}$$

and

$$\mathrm{Var}(q_{r(n)}) = q_{r(n)}^2 O(|\mathbb{S}(n)|^{-(1/2)}). \tag{51}$$

For the proof, see the Appendix.

Baccelli et al. [1] and McDonald and Reynier [15] proposed a mean field model of TCP congestion control that they proved was the mean field limit of a Markov model of a single link network. The convergence considered was similar to the series of networks $\mathbb{S}(n)$ used earlier. It was found that in the limit, the distributions of the packet marking probability at the link and the flow rates at sources converged toward the distributions obtained when the packet marking probability was defined deterministically. In steady state, a deterministic packet marking probability will be equal to a constant. If we compare the Lemmas 1 and 2 to Eq. (19) from the constant packet marking analysis, we find that a similar convergence occurs in our model. Furthermore, Lemmas 1 and 2 provide a lower bound on the rate at which this convergence occurs.

3.1. Parameter Choice Schemes

The results in Theorems 1 and 2 suggest that when flows are numerous and diverse, the coefficient of variation of $\mathbf{x}_r$ is close to $b_r/2$ at equilibrium, for each $r \in \mathbb{S}$. Thus,

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

for each $r \in \mathbb{S}$, it would seem to be beneficial to choose b_r as small as possible, so that $\mathrm{Var}(\mathbf{x}_r)$ is also as small as possible. However, if b_r is small, then $\mathbf{x}_r$ will react slowly to congestion information. So, there is a trade-off between the equilibrium variance and the rate of convergence of the algorithm.

Furthermore, equilibrium variance and rate of convergence are not the only issues to take into account when choosing parameters for Scalable TCP; fairness and stability are also very important considerations. If we define $\omega_r = a_r/(b_r T_r)$ for each $r \in \mathbb{S}$, then the equilibrium point of the deterministic models of Scalable TCP are approximately weighted proportionally fair, with weights ω_r. Thus, a_r and b_r should be chosen so that ω_r represents the importance (or willingness to pay) of user r. Furthermore, as proven by Kelly [8] and Kelly [11], the delayed fluid model of Scalable TCP, with links obeying Eq. (20), is locally stable, provided $a_r < \pi/(2\beta_{\mathbb{J}})$ for all $r \in \mathbb{S}$. Thus, chosen values of a_r should not go above this bound.

The interaction among the requirements of fairness, delay stability, equilibrium variance, and rate of convergence was discussed by Kelly [10], and the following parameter choice scheme was proposed. For all sources $r \in \mathbb{S}$,

$$a_r = \omega_r T_r \tilde{b}_r,\ b_r = \tilde{b}_r \quad \text{if } T_r \le \frac{\tilde{a}}{\omega_r \tilde{b}_r}, \tag{52}$$

$$\text{otherwise} \quad a_r = \tilde{a}, \quad b_r = \frac{\tilde{a}}{\omega_r T_r}, \tag{53}$$

where $\tilde{a}$, ω_r, and $\tilde{b}_r$ are constants and $\tilde{a} < \pi/(2\beta_{\mathbb{J}})$. This scheme ensures that the delay stability and fairness requirements given earlier are met, and, furthermore, $b_r \le \tilde{b}_r$, so this scheme also allows sources to determine an approximate upper limit on the steady state coefficient of variation of their flow rates.

When the network has links with large buffers using the E-RED scheme modeled by Eq. (21), the delay stability condition, derived by Liu et al. [13], is that for each $l \in \mathbb{J}$,

$$\kappa_l T_{\mathbb{S}} < \frac{1}{2}, \tag{54}$$

where $T_{\mathbb{S}} = \max_{r \in \mathbb{S}} T_r$. Clearly, this places no limits on source parameters, so a suitable scheme would be to set $a_r = \omega_r T_r b_r$ for all $r \in \mathbb{S}$, allowing each source to choose a preferred value of b_r. As for the previous case, this schemes ensures that the fairness requirements are met and, furthermore, that they also allow sources to determine an approximate upper limit on the coefficient of variation of their flow rates in equilibrium.

Note, that although each source might need to know its round-triptime in order to choose its parameters, this might be easily observed from the return of acknowledgment packets. Thus, we have given local, decentralized, potentially scalable parameter choice schemes.

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

3.2. E-TCP

In [5], Gu et al. identified a potential scalability problem for TCP variants acting over small buffer high-speed networks. They show that as a network evolves, if, for some link l, the capacity divided by the number of routes using l (C_l/N_l) tends to infinity, then the equilibrium utilization at l (z_l/C_l) tends to zero. Since this is the kind of evolution currently observed in the Internet, this could turn into a serious problem in the future. The proposed solution is a new TCP variant called E-TCP.

Under E-TCP, when a successful acknowledgment packet is received by a source, it increases its congestion window $\mathbf{w}$ by a constant a. When a marked or dropped packet is detected, the source reduces $\mathbf{w}$ by $ab\mathbf{w}/(a + q_0 b\mathbf{w})$, for constants b and q_0. Here, q_0 represents an acceptable level of packet marking that the source will always allow. If $q = 0$, then this becomes an implementation of STCP.

We can apply the same modeling approach to E-TCP as to STCP, which yields the stochastic differential equation for the single source case,

$$d\mathbf{w}(t) = -\frac{w}{T}h'(w)\mathbf{w}(t)p\,dt + (a + h(w))\left((1-p)p\frac{w}{T}\right)^{1/2} d\mathbf{B}_r(t), \qquad \textbf{(55)}$$

where $\mathbf{B}(t)$ is a Brownian motion and

$$h(w) = \frac{abw}{a + q_0 bw}. \qquad \textbf{(56)}$$

The invariant distribution of the solution to Eq. (55) is Gaussian with variance

$$\mathrm{Var}(\mathbf{w}) = \frac{(a + h(w))^2(1-p)}{2h'(w)} = \frac{h(w)^2}{2h'(w)(1-p)} = bw^2\frac{1}{2(1-p)}. \qquad \textbf{(57)}$$

This is the same as for the Scalable TCP case and would suggest that E-TCP has promising stochastic scalability properties. Furthermore, the delay stability requirement for this system given by Gu et al. [5] is that $a_r < \pi/(2\beta_{\mathbb{J}})$ for all $r \in \mathbb{S}$. Thus, it would seem that the decentralized, scalable parameter choice scheme discussed in Section 3.1 applies to E-TCP as well as to STCP.

However, when the same method of analysis used in Theorem 1 is applied to the E-TCP case, the results are weak. We are still able to bound how close $\mathrm{Var}(\mathbf{x}_r)$ is to $b_r x_r^2/2(1-q_r)$, however, rather than just relying on x_r/z_l to be small for $l \in r$, this bound requires terms such as $(a_r + b_r w_r q_0)x_r/z_l$ and $(a_s + b_s w_s q_0)(a_r + b_r w_r q_0)x_r/z_l$ to be small for $l \in r \cap s$. This is a weaker result, as it requires the number of flows per link to scale in proportion to flow speed squared. Thus, the scope of network evolutions over which E-TCP is a good choice might be more limited than at first thought. It remains an open question whether stronger stochastic scalability results exist for E-RED or whether a more suitable alternative can be designed.

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

4. CONCLUSION

In the past, deterministic fluid flow models—in particular, those that take into account propagation delays— have been a major focus of analysis for the mathematical modeling of congestion control in large networks. This analysis has been very successful and algorithms have been found that have many desirable properties. In particular, the Scalable TCP and E-RED algorithms have scalable, decentralized conditions for local delay stability. However, corresponding stochastic scalability results have been lacking.

In this article we examined the stochastic behavior of Scalable TCP with static link responses and Scalable TCP with E-RED AQM at links. By sketching the proof of a limiting convergence result, we have motivated a continuous time stochastic differential equation model of networkwide congestion control. Using this model, we derived bounds on stationary distribution variances of flow rates that apply to both algorithms. We have also shown that in the mean field limit, these bounds become tight and give the same values for variance as when packet marking probabilities are constant.

Furthermore, the bounds we have derived are scalable. When there are large numbers of sources throughout the network, the coefficient of variation of the flow rate of source r will be bounded by $0.5\sqrt{b_r}$, where b_r is the value r uses for its decrease parameter. This is a highly scalable bound, since a source can ensure that the coefficient of variation of its flow rate in equilibrium is small without any knowledge of the rest of the network.

We have also given a brief discussion of how these results fit in with other results regarding the delay stability of Scalable TCP. We have shown that in the case of STCP in a small buffer environment, our bounds are compatible with a previously proposed delay stable, scalable parameter choice scheme, given by Kelly [11]. For STCP with E-RED, a delay stable, scalable parameter choice scheme is easily derived.

Since Scalable TCP has potential link utilization problems over high-speed small buffer networks, we have explored how our analysis extends to the newly proposed algorithm E-TCP. We find that some of the conclusions we made for STCP still apply and that the scalable parameter choice scheme discussed in Section 3.1 remains attractive. However there remain open questions as to whether E-TCP will really prove to be stochastically scalable in practical applications.

The parameter schemes we discuss neatly deal with the complex relationships among fairness, stability in the presence of propagation delays, stochastic considerations and rate of convergence. Moreover, they show a potential for STCP, E-TCP and E-RED to achieve a compromise that will have attractive properties in all of these areas, however great the scale of network on which they operate.

References

1. Baccelli, F., McDonald, D. & Reynier, J. (2002). A mean-field model for multiple TCP connections through a buffer implementing RED. In *Proceedings of IFIP Performance 2002*, Romo.
2. Bain, A. & Key, P. (2002). Modelling the performance of in-call probing for multilevel adaptive applications. Technical report Msr-tr-2002-06, Microsoft Research.

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

3. Barnett, S. (1975). *Introduction to mathematical control theory*. Oxford: Oxford University Press.
4. Cooper, R. (1972). *Introduction to queueing theory*. New York: Macmillan.
5. Gu, Y., Towsley, D., Hollot, C. & Zhang, H. (2007). Congestion control for small buffer high speed networks. In *Proceedings of IEEE Infocom*. Available from ftp://gaia.cs.umass.edu/pub/gu06-14.pdf.
6. Hammarling, S.I. (1982). Numerical solution of stanle non-negative definite Iyapunov equation. *IMA Journal of Numerical Analysis* 2: 303–323.
7. Johari, R. & Tan, D. (2001). End-to-end congestion control for the internet: Delays and stability. *IEEE/ACM Transactions on Networking* 9: 818–832.
8. Kelly, F.P. (2003). Fairness and stability of end-to-end congestion control. *European Journal of Control* 9: 159–176.
9. Kelly, F.P., Maulloo, A.K. & Tan, D.K.H. (1998). Rate control in communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* 49: 237–252.
10. Kelly, T. (2002). On engineering a stable and scalable TCP variant. Technical report CUED/FINFENG/TR.435 Cambridge: Cambridge University Engineering Department.
11. Kelly, T. (2003). Scalable TCP: Improving performance in highspeed wide area networks. *Computer Communication Review* 33: 83–91.
12. Liu, S., Basar, T. & Srikant, R. (2003). Controlling the internet: A survey and some new results. In *Proceedings of IEEE Conference on Decision and Control*.
13. Liu, S., Basar, T. & Srikant, R. (2005). Exponential RED: A stabilizing AQM scheme for low and high speed TCP. IEEE *Transactions on Networking* 13(5): 1068–1081.
14. Massoulié, L. (2000). Stability of distributed congestion control with heterogeneous feedback delays. Technical report Msr-tr-2000-111, Richmond, VA: Microsoft Research.
15. McDonald, D. & Reynier, J. (2006). Mean field convergence of a model of multiple TCP connections through a buffer implementing RAD. *Annals of Applied Probability* 16(1): 244–294.
16. Oksendal, B. (1992). *Stochastic differential equations: An introduction with applications*. Berlin: Springer.
17. Ott, T.J. (2005). Transport protocols in the TCP paradigm and their performance. *Telecommunication Systems* 30(4): 351–385.
18. Ott, T.J. (2006). Rate of convergence for the "square root formula." *Advances in Applied Probability* 38(4): 1132–1154.
19. Ott, T.J. & Swanson, J. (2007). Asymptotic behavior of a generalized TCP congestion avoidance algorithm. *Journal of Applied Probability* 44: 618–635.
20. Raina, G., Wischik, D. & Towsley, D. (2005). Control theory for buffer sizing. *Computer Communication Review* 35(3): 79–82.
21. Srikant, R. (2004). *The mathematics of internet congestion control*. Boston: Birkhauser.
22. Vinnicombe, G. (2002). On the stability of networks operating TCP-like congestion control. In *Proceedings of IFAC*.

APPENDIX

Here, we include any proofs left out of the main body of the article.

PROOF OF THEOREM 1: The matrix A can be decomposed as $A = -(D + E)$, where D is a diagonal matrix with entries

$$D_{r,r} = b_r \alpha_r (1 - q_r) T_r^{-1} \tag{A.1}$$

and E is a real symmetric matrix. Since E is the product of a matrix and its hermitian conjugate, E is positive definite.

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

Now, let $\mathcal{D}$ be the diagonal linear operator that acts on $|\mathbb{S}| \times |\mathbb{S}|$ matrices such that, for all $r, s \in \mathbb{S}$,

$$\mathcal{D}(Q)_{r,s} = Q_{r,s}(D_{r,r} + D_{s,s})^{1/2}, \tag{A.2}$$

and let $\mathcal{M}$ be the linear operator such that, for all $r, s \in \mathbb{S}$,

$$\mathcal{M}(Q)_{r,s} = \frac{1}{2}\sum_{k\in\mathbb{S}}\big(E_{r,k}Q_{k,s} + E_{s,k}Q_{k,r} + Q_{r,k}E_{k,s} + Q_{s,k}E_{k,r}\big). \tag{A.3}$$

Thus, Eq. (37) is equivalent to $\mathcal{D}^2(\Sigma) + \mathcal{M}(\Sigma) = C$, where $C = FF^T$. If we let $\hat{\Sigma}$ be defined as $\mathcal{D}(\Sigma)$, we get

$$\hat{\Sigma} + \mathcal{D}^{-1}\mathcal{M}\mathcal{D}^{-1}(\hat{\Sigma}) = \mathcal{D}^{-1}(C), \tag{A.4}$$

where adjacency of linear operators denotes composition.

The operator $\mathcal{M}$ is positive definite with respect to the inner product

$$\langle Q, P\rangle = \mathrm{tr}(Q^\dagger P), \tag{A.5}$$

for, given a matrix Q,

$$\langle Q, \mathcal{M}(Q)\rangle = \frac{1}{2}\mathrm{tr}\big(Q^\dagger(EQ + Q^TE + QE + EQ^T)\big) \tag{A.6}$$

$$= \frac{1}{2}\mathrm{tr}\big((Q^\dagger + Q^*)E(Q + Q^T)\big) \geq 0, \tag{A.7}$$

where $(\cdot)^*$ denotes complex conjugate. Let $\|\cdot\|_f$ denote the Frobenius norm, the inner product norm with respect to $\langle\cdot,\cdot\rangle$.

Now, let $\mathcal{P}(\cdot)$ denote the projection operator onto the diagonal elements of matrices. So, for a matrix Q, $\mathcal{P}(Q)_{r,r} = Q_{r,r}$ and $\mathcal{P}(Q)_{r,s} = 0$ if $r \neq s$. Let $\mathcal{P}_\perp(\cdot)$ denote the orthogonal projection to $\mathcal{P}(\cdot)$, so $\mathcal{P}(Q) + \mathcal{P}_\perp(Q) = Q$ for all Q. Let $\rho = \mathcal{P}(\hat{\Sigma})$ and let $\varrho = \mathcal{P}_\perp(\hat{\Sigma})$. From Eq. (A.4), we get

$$\rho + \mathcal{P}\mathcal{D}^{-1}\mathcal{M}\mathcal{D}^{-1}(\rho) = \mathcal{D}^{-1}(C) - P\mathcal{D}^{-1}\mathcal{M}\mathcal{D}^{-1}(\varrho), \tag{A.8}$$

$$\varrho + \mathcal{P}_\perp\mathcal{D}^{-1}\mathcal{M}\mathcal{D}^{-1}(\varrho) = -\mathcal{P}_\perp\mathcal{D}^{-1}\mathcal{M}\mathcal{D}^{-1}(\rho). \tag{A.9}$$

We can derive $\mathcal{P}\mathcal{D}^{-1}\mathcal{M}\mathcal{D}^{-1}(\rho) = \mathcal{P}\mathcal{D}^{-1}\mathcal{M}\mathcal{D}^{-1}\mathcal{P}(\rho) = \mathcal{E}(\rho)$, where for any matrix Q, for all $r \in \mathbb{S}$,

$$\mathcal{E}(Q)_{r,r} = \frac{E_{r,r}}{D_{r,r}}Q_{r,r} \tag{A.10}$$

and all other entries of $\mathcal{E}(Q)$ are zero. Thus, taking $\mathcal{X} = \mathcal{D}^{-1}\mathcal{M}\mathcal{D}^{-1}$,

$$\big((I + \mathcal{E})\rho - \mathcal{D}^{-1}(C)\big) = \mathcal{P}\mathcal{X}\mathcal{P}_\perp(I + \mathcal{P}_\perp\mathcal{X}\mathcal{P}_\perp)^{-1}\mathcal{P}_\perp\mathcal{X}\mathcal{P}(\rho). \tag{A.11}$$

Note that for all $r \in \mathbb{S}$,

$$\frac{E_{r,r}}{D_{r,r}} = \frac{w_r(\alpha_r + w_r)}{T_r\alpha_r(1 - q_r)}\sum_{l\in r}\frac{\beta_l p_l}{z_l} = \frac{x_r}{q_r(1 - q_r)}\sum_{l\in r}\frac{\beta_l p_l}{z_l} \leq 2\gamma\beta_{\mathbb{J}}. \tag{A.12}$$

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

We use $\|\cdot\|_*$ to denote the maximum entry modulus of a matrix, so for any Q, $\|Q\|_* = \max_{r,s} |Q_{r,s}|$. From the above, we get

$$\|\rho\|_* \le \|\mathcal{D}^{-1}(C)\|_* + \|\mathcal{P}\mathcal{X}\mathcal{P}_\perp(I + \mathcal{P}_\perp\mathcal{X}\mathcal{P}_\perp)^{-1}\mathcal{P}_\perp\mathcal{X}\mathcal{P}(\rho)\|_*. \tag{A.13}$$

Now, for each $r \in \mathbb{S}$, let $e(r)$ be the matrix that has $e(r)_{r,r} = 1$ and all other entries zero. For some r',

$$\|\mathcal{P}\mathcal{X}\mathcal{P}_\perp(I + \mathcal{P}_\perp\mathcal{X}\mathcal{P}_\perp)^{-1}\mathcal{P}_\perp\mathcal{X}\mathcal{P}(\rho)\|_* = |\langle e(r'), \mathcal{P}\mathcal{X}\mathcal{P}_\perp(I + \mathcal{P}_\perp\mathcal{X}\mathcal{P}_\perp)^{-1}\mathcal{P}_\perp\mathcal{X}\mathcal{P}(\rho)\rangle|. \tag{A.14}$$

For any positive operator $\mathcal{Y}$ and any matrices M and N from the Cauchy–Schwarz inequality, we get that $|\langle M, \mathcal{Y}(N)\rangle|$ is bounded by $\|M\|_f\|N\|_f$ times the maximum eigenvalue of $\mathcal{Y}$. Since $\mathcal{P}_\perp\mathcal{X}\mathcal{P}_\perp$ is positive, all eigenvalues of $(I + \mathcal{P}_\perp\mathcal{X}\mathcal{P}_\perp)^{-1}$ are bounded by 1. Thus, the maximum eigenvalue of $\mathcal{P}\mathcal{X}\mathcal{P}_\perp(I + \mathcal{P}_\perp\mathcal{X}\mathcal{P}_\perp)^{-1}\mathcal{P}_\perp\mathcal{X}\mathcal{P}$ is bounded by the maximum eigenvalue of $\mathcal{P}\mathcal{X}\mathcal{P}_\perp\mathcal{X}\mathcal{P}$. The maximum eigenvalue of $\mathcal{P}\mathcal{X}\mathcal{P}_\perp\mathcal{X}\mathcal{P}$ is bounded by its maximum absolute row sum. Since all entries of $\mathcal{P}\mathcal{X}\mathcal{P}_\perp\mathcal{X}\mathcal{P}$ are positive, its maximum absolute row sum is equal to

$$\max_{r\in\mathbb{S}} \langle e(r), \mathcal{X}\mathcal{P}_\perp\mathcal{X}(e)\rangle, \text{ for } e = \sum_{r\in\mathbb{S}} e(r). \tag{A.15}$$

By positivity of $\mathcal{X}$ and $\mathcal{P}_\perp$, $\langle e(r), \mathcal{X}\mathcal{P}_\perp\mathcal{X}(e)\rangle = \langle \mathcal{P}_\perp\mathcal{X}(e(r)), \mathcal{X}(e)\rangle$.

For all $s \in \mathbb{S}$, $s \ne r$,

$$\mathcal{P}_\perp\mathcal{X}(e(r))_{r,s} = \mathcal{P}_\perp\mathcal{X}(e(r))_{s,r} = \big(2D_{r,r}(D_{r,r} + D_{s,s})\big)^{-(1/2)}E_{r,s} \tag{A.16}$$

and all other entries of $\mathcal{P}_\perp\mathcal{X}(e(r))$ are zero. For all $r, s \in \mathbb{S}$,

$$\mathcal{X}(e)_{r,s} = \frac{1}{\sqrt{2}}(D_{r,r} + D_{s,s})^{-(1/2)}\big(D_{r,r}^{-(1/2)} + D_{s,s}^{-(1/2)}\big)E_{r,s}. \tag{A.17}$$

Thus,

$$\langle \mathcal{P}_\perp\mathcal{X}(e(r)), \mathcal{X}(e)\rangle = \sum_{s\ne r}(D_{r,r} + D_{s,s})^{-1}D_{r,r}^{-(1/2)}\left(D_{r,r}^{-(1/2)} + D_{s,s}^{-(1/2)}\right)E_{r,s}^2 \tag{A.18}$$

$$\le 2\sum_{s\in\mathbb{S}} D_{r,r}^{-1}E_{r,s}^2D_{s,s}^{-1}$$

$$\le 2\sum_{s\in\mathbb{S}} \frac{w_r(\alpha_r + w_r)w_s(\alpha_s + w_s)}{T_r\alpha_r(1-q_r)T_s\alpha_s(1-q_s)}\left(\sum_{l\in r\cap s}\frac{\beta_l p_l}{z_l}\right)^2 \tag{A.19}$$

$$\le 8\frac{1}{q_r}\sum_{l\in r}\frac{\beta_l p_l}{z_l}\sum_{s:l\in s}\frac{x_s}{q_s}\sum_{j\in s}\frac{\beta_j p_j}{z_j}R_{r,j}x_r \le 8\beta_{\mathbb{J}}^2\gamma. \tag{A.20}$$

The above discussion implies

$$\|\mathcal{P}\mathcal{X}\mathcal{P}_\perp(I + \mathcal{P}_\perp\mathcal{X}\mathcal{P}_\perp)^{-1}\mathcal{P}_\perp\mathcal{X}\mathcal{P}(\rho)\|_* \le 8\beta_{\mathbb{J}}^2\gamma\|\rho\|_f \le 8\beta_{\mathbb{J}}^2\gamma\sqrt{|\mathbb{S}|}\|\rho\|_*. \tag{A.21}$$

Hence, from Eq. (A.13),

$$\|\rho\|_* \le (1 - 8\beta_{\mathbb{J}}^2\gamma\sqrt{|\mathbb{S}|})^{-1}\|\mathcal{D}^{-1}(C)\|_*, \tag{A.22}$$

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

and from Eq. (A.11),

$$\|\rho - \mathcal{D}^{-1}(C)\|_* \leq \|\mathcal{E}(\rho)\|_* + 8\beta_{\mathbb{J}}^2\gamma\sqrt{|\mathbb{S}|}\|\rho\|_* \tag{A.23}$$

$$\leq 2\beta_{\mathbb{J}}\gamma(1 + 4\beta_{\mathbb{J}}\sqrt{|\mathbb{S}|})(1 - 8\beta_{\mathbb{J}}^2\gamma\sqrt{|\mathbb{S}|})^{-1}\|\mathcal{D}^{-1}(C)\|_* \tag{A.24}$$

$$\leq 4\beta_{\mathbb{J}}(1 + 8\beta_{\mathbb{J}})\gamma\sqrt{|\mathbb{S}|}\|\mathcal{D}^{-1}(C)\|_*, \tag{A.25}$$

with the last result following from the assertion $\gamma\sqrt{|\mathbb{S}|} < 1/(16\beta_{\mathbb{J}}^2)$ and the fact that $\sqrt{\mathbb{S}} \geq 1$. Multiplying this relation by $\mathcal{D}^{-1}$ gives us that, for all $r \in \mathbb{S}$, $\Sigma_{r,r}$ must lie within $\gamma\sqrt{\mathbb{S}}\zeta_r/2$ of $1/2$, where

$$\zeta_r = 4\beta_{\mathbb{J}}(1 + 8\beta_{\mathbb{J}})\left(\frac{b_r\alpha_r(1-q_r)}{T_r}\right)^{-(1/2)}\max_{s\in\mathbb{S}}\left(\frac{b_r\alpha_r}{T_s}\right)^{1/2}, \tag{A.26}$$

as in the statement of the theorem. Since $\mathrm{Var}(\mathbf{x}_r) = b_r x_r^2[1/(1-q_r)]\Sigma_{r,r}$, the result follows. ■

Proof of Theorem 2: As in the proof of Theorem 1, we use the inner product $\langle\cdot,\cdot\rangle$, defined for matrices Q and Q' as

$$\langle Q, Q'\rangle = \mathrm{tr}(Q^\dagger Q'), \tag{A.27}$$

and the associated inner product norm $\|\cdot\|_f$.

From Eq. (37), with $C = \acute{F}(\acute{F})^\dagger$, and from the symmetry of $\acute{\Sigma}$, we can deduce

$$\acute{\Sigma}_{r,s} = \frac{-\sum_{l\in s}\acute{A}_{s,l}\acute{\Sigma}_{l,r} - \sum_{l\in r}\acute{A}_{r,l}\acute{\Sigma}_{l,s} - C_{r,s}}{\acute{A}_{r,r} + \acute{A}_{s,s}}. \tag{A.28}$$

For $l, j \in \mathbb{J}$,

$$\sum_{r:l\in r}\acute{A}_{l,r}\acute{\Sigma}_{j,r} + \sum_{r:j\in r}\acute{A}_{j,r}\acute{\Sigma}_{l,r} = 0. \tag{A.29}$$

Additionally, for all $r \in \mathbb{S}$, $l \in \mathbb{J}$,

$$0 = \acute{A}_{r,r}\acute{\Sigma}_{l,r} + \sum_{s\in\mathbb{S}}\acute{A}_{l,s}\acute{\Sigma}_{s,r} + \sum_{j\in\mathbb{J}}\acute{A}_{r,j}\acute{\Sigma}_{l,j}. \tag{A.30}$$

Let $\mathcal{P}$, $\mathcal{P}_\dashv$, and $\mathcal{P}_\vdash$ be linear operators acting on the space of $(|\mathbb{S}| + |\mathbb{J}|) \times (|\mathbb{S}| + |\mathbb{J}|)$ matrices, such that for any matrix Q, $\mathcal{P}(Q)$ is the upper left $|\mathbb{S}| \times |\mathbb{S}|$ submatrix of Q, $\mathcal{P}_\dashv(Q)$ is the lower left $|\mathbb{J}| \times |\mathbb{S}|$ submatrix of Q, and $\mathcal{P}_\vdash(Q)$ is the lower right $|\mathbb{J}| \times |\mathbb{J}|$ submatrix of Q.

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

Thus, for all $r, s \in \mathbb{S}$, and $l, j \in \mathbb{J}$,

$$\mathcal{P}(Q)_{r,s} = Q_{r,s}, \qquad \mathcal{P}_{\dashv}(Q)_{l,r} = Q_{l,r}, \qquad \mathcal{P}_{\vdash}(Q)_{l,j} = Q_{l,j}. \tag{A.31}$$

Let $\rho = \mathcal{P}(\acute{\Sigma})$, $\phi = \mathcal{P}_{\dashv}(\acute{\Sigma})$, $\psi = \mathcal{P}_{\vdash}(\acute{\Sigma})$, and $E = \mathcal{P}_{\dashv}(\acute{A})$. Define linear operators $\mathcal{N}$ and $\mathcal{M}$ such that for any $|\mathbb{J}| \times |\mathbb{S}|$ matrix Q,

$$\mathcal{N}(Q) = E^T Q + Q^T E, \;\; \mathcal{M}(Q) = QE^T + EQ^T. \tag{A.32}$$

It can be calculated that for any suitable Q and Q',

$$\mathcal{N}^T(Q) = E(Q + Q^T), \qquad \mathcal{M}^T(Q') = (Q' + Q'^T)E. \tag{A.33}$$

Let $\mathcal{D}$ and $\mathcal{E}$ be diagonal operators such that for any $|\mathbb{J}| \times |\mathbb{S}|$ matrix Q, for all $r \in \mathbb{S}$, $l \in \mathbb{J}$, $\mathcal{D}(Q)_{l,r} = |\acute{A}_{r,r}|^{(1/2)} Q_{l,r}$ and for any $|\mathbb{S}| \times |\mathbb{S}|$ matrix Q, for all $r, s \in \mathbb{S}$,

$$\mathcal{E}(Q)_{r,s} = |\acute{A}_{r,r} + \acute{A}_{s,s}|^{1/2} Q_{r,s}. \tag{A.34}$$

Now, Eqs (A.28)–(A.30) become,

$$\rho = \mathcal{E}^{-2}\mathcal{P}(C) - \mathcal{E}^{-2}\mathcal{N}(\phi), \;\; \mathcal{M}(\phi) = 0, \tag{A.35}$$

$$\mathcal{D}^2(\phi) = \frac{1}{2}\mathcal{N}^T(\rho) - \frac{1}{2}\mathcal{M}^T(\psi). \tag{A.36}$$

This last relation follows from the definition of $\mathcal{N}$ and $\mathcal{M}$ and the symmetry of ρ and ψ.

Let $\mathcal{P}_M$ be the projector on the space of $|\mathbb{S}| \times |\mathbb{J}|$ matrices with image $\{\mathcal{D}(Q) : Q \in \ker(\mathcal{M})\}$, where $\ker(\mathcal{M})$ denotes the kernel of $\mathcal{M}(\cdot)$. For any suitable matrix Q,

$$\|\mathcal{P}_M \mathcal{D}^{-1}\mathcal{M}^T(Q)\|_f = \langle Q, \mathcal{M}\mathcal{D}^{-1}\mathcal{P}_M\mathcal{D}^{-1}\mathcal{M}^T(Q)\rangle = 0, \tag{A.37}$$

as the image of $\mathcal{D}^{-1}\mathcal{P}_M$ lies in the kernel of $\mathcal{M}$. Thus, $\mathcal{P}_M\mathcal{D}^{-1}\mathcal{M}^T = 0$. Furthermore, $\mathcal{P}_M\mathcal{D}(\phi) = \mathcal{D}(\phi)$, since ϕ is in the kernel of $\mathcal{M}(\cdot)$. So, if we apply $\mathcal{P}_M\mathcal{D}^{-1}$ to Eq. (A.36), we get $\mathcal{D}(\phi) = \frac{1}{2}\mathcal{P}_M\mathcal{D}^{-1}\mathcal{N}^T(\rho)$. If we let $\varphi = \mathcal{D}(\phi)$ and make use of Eq. (A.35) we get

$$\varphi = \frac{1}{2}\mathcal{P}_M\mathcal{D}^{-1}\mathcal{N}^T(\mathcal{E}^{-2}\mathcal{P}(C) - \mathcal{E}^{-2}\mathcal{N}\mathcal{D}^{-1}(\varphi)) \tag{A.38}$$

and, so,

$$\varphi = \frac{1}{2}(I + \mathcal{X})^{-1}\mathcal{P}_M\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-2}\mathcal{P}(C), \tag{A.39}$$

where $\mathcal{X} = \mathcal{P}_M\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-2}\mathcal{N}\mathcal{D}^{-1}\mathcal{P}_M$. Since $\mathcal{E}$ is diagonal with positive real entries, $\mathcal{X}$ must be a positive matrix. If we let $\varrho = \mathcal{E}(\rho)$ and substitute our expression for φ into Eq. (A.35), we

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

derive

$$(\varrho - \mathcal{E}^{-1}\mathcal{P}(C)) = -\frac{1}{2}\mathcal{E}^{-1}\mathcal{N}\mathcal{D}^{-1}\mathcal{P}_M(I+\mathcal{X})^{-1}\mathcal{P}_M\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-2}\mathcal{P}(C). \tag{A.40}$$

As in the proof of Theorem 1, we let $\|\cdot\|_*$ denote the maximum entry modulus norm. For $r, s \in \mathbb{S}$, let $e(r,s)$ be the matrix with $e(r,s)_{r,s} = 1$ and all other entries zero. It follows that

$$\|\varrho - \mathcal{E}^{-1}\mathcal{P}(C)\|_* = \frac{1}{2}|\langle e(r',s'), \mathcal{E}^{-1}\mathcal{N}\mathcal{D}^{-1}\mathcal{P}_M(I+\mathcal{X})^{-1}\mathcal{P}_M\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-2}\mathcal{P}(C)\rangle| \tag{A.41}$$

$$= \frac{1}{2}|\langle \mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-1}(e(r',s')), \mathcal{P}_M(I+\mathcal{X})^{-1}\mathcal{P}_M\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-2}\mathcal{P}(C)\rangle|, \tag{A.42}$$

for some $r', s' \in \mathbb{S}$. By positivity, as in the proof of Theorem 1, this is less than or equal to

$$\frac{1}{2}\eta\|\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-1}(e(r',s'))\|_f\|\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-2}\mathcal{P}(C)\|_f, \tag{A.43}$$

where η is the maximum modulus of the eigenvalues of $\mathcal{P}_M(I+\mathcal{X})^{-1}\mathcal{P}_M$. However, since $\mathcal{X}$ is positive and $\mathcal{P}_M$ is a projector, $\eta \le 1$.

Let us seek a bound on $\|\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-1}(Q)\|_f$ for arbitrary matrices Q in terms of $\|Q\|_f$. All antisymmetric matrices lie in the kernel of $\mathcal{N}^T\mathcal{E}^{-1}$, so we can assume, without loss of generality, that Q has no antisymmetric part and is thus symmetric. So,

$$\|\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-1}(Q)\|_f^2 = 2\sum_{l\in\mathbb{J}}\sum_{r,s\in\mathbb{S}}|\acute{A}_{s,s}(\acute{A}_{r,r}+\acute{A}_{s,s})|^{-1}\acute{A}_{l,r}^2Q_{r,s}^2. \tag{A.44}$$

Now, for all $r \in \mathbb{S}$,

$$\sum_{l\in\mathbb{J}}\frac{\acute{A}_{l,r}^2}{|\acute{A}_{r,r}|} = \sum_{l\in r}\frac{\kappa_l p_l x_r}{C_l q_r(1-q_r)} \le 2\gamma\max_{l\in\mathbb{J}}\kappa_l = 2\gamma\kappa_{\mathbb{J}}. \tag{A.45}$$

Thus,

$$\|\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-1}(Q)\|_f^2 \le 4\gamma\kappa_{\mathbb{J}}\max_{r\in\mathbb{S}}|\acute{A}_{r,r}|^{-1}\|Q\|_f^2, \tag{A.46}$$

so $\|\varrho - \mathcal{E}^{-1}\mathcal{P}(C)\|_*$ is less than or equal to $4\gamma\kappa_{\mathbb{J}}\max_{r\in\mathbb{S}}|\acute{A}_{r,r}|^{-1}\|\mathcal{E}^{-1}\mathcal{P}(C)\|_f$. Furthermore,

$$\|\mathcal{E}^{-1}\mathcal{P}(C)\|_f^2 = \sum_{r\in\mathbb{S}}\frac{C_{r,r}^2}{2|A_{r,r}|} \le \frac{1}{2}|\mathbb{S}|\max_{r\in\mathbb{S}}C_{r,r}. \tag{A.47}$$

If follows that

$$\|\varrho - \mathcal{E}^{-1}\mathcal{P}(C)\|_* \le 2\sqrt{2}\gamma\sqrt{|\mathbb{S}|}\kappa_{\mathbb{J}}\max_{r,s\in\mathbb{S}}|A_{r,r}|^{-1}C_{s,s}^{1/2}. \tag{A.48}$$

If we multiply this by $\mathcal{E}^{-1}$, we get that for all $r \in \mathbb{S}$, $\acute{\Sigma}_{r,r}$ is within $\gamma\sqrt{|\mathbb{S}|}\zeta_r/2$ of $1/2$, where

$$\zeta_r = 4\sqrt{2}\kappa_{\mathbb{J}}\max_{r,s\in\mathbb{S}}\left(\frac{b_r\alpha_r}{T_r}\right)^{-(3/2)}\left(\frac{b_s\alpha_s}{T_s}\right)^{1/2}, \tag{A.49}$$

as in the statement of this theorem. Since $\mathrm{Var}(\mathbf{x}_r) = b_r x_r^2[1/(1-q_r)]\acute{\Sigma}_{r,r}$, the result follows. ∎

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

In the proof of Lemma 2, we will need the following result, which continues from the analysis given above.

LEMMA 3: *Under the equilibrium distribution of Eq.* (31),

$$\mathrm{Var}(\mathbf{q}_r) \le q_r^2 \left(\frac{b_r T_r}{\alpha_r}(1-q_r) \right)^{1/2} \gamma \sqrt{|\mathbb{S}|}\kappa_{\mathbb{J}} \max_{r,s\in\mathbb{S}} |\acute{A}_{r,r}|^{-1} C_{s,s}^{-(1/2)}. \tag{A.50}$$

PROOF: Using the definitions and notation of the proof of Theorem 2, we note that from Eq. (A.35) and the derived expression for φ,

$$\varrho = \mathcal{E}^{-1}\mathcal{P}(C) - \frac{1}{2}\mathcal{E}^{-1}\mathcal{N}\mathcal{D}^{-1}\mathcal{P}_M\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-1}(\varrho), \tag{A.51}$$

and, hence, by positivity, $\|\varrho\|_f \le \|\mathcal{E}^{-1}\mathcal{P}(C)\|_f$. Now, let $\mathcal{P}_\perp$ be the projection orthogonal to $\mathcal{P}_M$, so ϱ lies in the kernel of $\mathcal{P}_\perp$ and $\mathcal{P}_\perp\mathcal{D}^{-1}\mathcal{M}^T = \mathcal{D}^{-1}\mathcal{M}^T$. Applying $\mathcal{E}^{-1}\mathcal{N}\mathcal{D}^{-1}\mathcal{P}_\perp\mathcal{D}^{-1}$ to Eq. (A.36) results in

$$\mathcal{E}^{-1}\mathcal{N}\mathcal{D}^{-2}\mathcal{M}^T(\psi) = \mathcal{E}^{-1}\mathcal{N}\mathcal{D}^{-1}\mathcal{P}_\perp\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-1}(\varrho). \tag{A.52}$$

For $r, s \in \mathbb{S}$, let $e(r,s)$ be the matrix with $e(r,s)_{r,s} = 1$ and all other entries equal to zero. There exists some r and s such that

$$\|\mathcal{E}^{-1}\mathcal{N}\mathcal{D}^{-2}\mathcal{M}^T(\psi)\|_* = \langle e(r,s), \mathcal{E}^{-1}\mathcal{N}\mathcal{D}^{-1}\mathcal{P}_\perp\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-1}(\varrho)\rangle, \tag{A.53}$$

$$\le \|\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-1}(e(r,s))\|_f \|\mathcal{D}^{-1}\mathcal{N}^T\mathcal{E}^{-1}(\varrho)\|_f, \tag{A.54}$$

since $\mathcal{P}_\perp$ is a projector. Thus, from Eq. (A.46), and our bound on $\|\varrho\|_f$,

$$\|\mathcal{E}^{-1}\mathcal{N}\mathcal{D}^{-2}\mathcal{M}^T(\psi)\|_* \le 2\sqrt{2}\gamma\sqrt{|\mathbb{S}|}\kappa_{\mathbb{J}} \max_{r,s\in\mathbb{S}} |\acute{A}_{r,r}|^{-1} C_{s,s}^{-(1/2)}. \tag{A.55}$$

By symmetry of ψ, $\mathcal{M}^T(\psi) = 2\psi E$. So, for any $r \in \mathbb{S}$,

$$\mathcal{E}^{-1}\mathcal{N}\mathcal{D}^{-2}\mathcal{M}^T(\psi)_{r,r} = 2\sqrt{2}\acute{A}_{r,r}^{-(3/2)} \sum_{l,j\in\mathbb{J}} \acute{A}_{l,r}\acute{A}_{j,r}\psi_{l,j} \tag{A.56}$$

$$= 2\sqrt{2}b_r^{-(1/2)}T_r^{-(1/2)}\left(\alpha_r(1-q_r)\right)^{-(3/2)} w_r(\alpha_r + w_r)\,\mathrm{Var}(\mathbf{q}_r) \tag{A.57}$$

$$= 2\sqrt{2}b_r^{-(1/2)}T_r^{-(1/2)}\alpha_r^{1/2}(1-q_r)^{-(1/2)}\frac{1}{q_r^2}\,\mathrm{Var}(\mathbf{q}_r). \tag{A.58}$$

The result follows. ■

PROOF OF LEMMA 1: For any $r \in \mathbb{S}(n)$, $j \in r$, since there are $N_{s(r)}(n)$ copies of r in $\mathbb{S}(n)$ and each must have the same equilibrium flow rate, $(x_r/z_j) \le (1/N_{s(r)}(n))$. Thus, if we let $\gamma(n)$ denote the value of γ for $\mathbb{S}(n)$, we get

$$\gamma(n) \le \min_{r\in\mathbb{S}(1)} \frac{1}{N_r(n)}. \tag{A.59}$$

By the choice of the $\mathbb{S}(n)$, for all $r \in \mathbb{S}(1)$, $\liminf_{n\to\infty}(N_r(n)/|\mathbb{S}(n)|) > 0$; thus, $\gamma(n) = O(|\mathbb{S}(n)|^{-1})$ and so $\gamma(n)|\mathbb{S}(n)|^{-(1/2)} = O(|\mathbb{S}(n)|^{1/2})$. For $r \in \mathbb{S}(1)$ and $r(n) \in \mathbb{S}(n)$

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press

with $r = s(r(n))$, we have $b_{r(n)} = b_r$, $\alpha_{r(n)} = \alpha_r$ and $T_{r(n)} = T_r$, so Theorem 1 tell us that

$$\mathrm{Var}(\mathbf{x}_{r(n)}) = b_r x_{r(n)}^2 \frac{1}{2(1-q_{r(n)})}(1 + O(|\mathbb{S}(n)|^{-(1/2)}), \tag{A.60}$$

as required.

For each $j \in J$, under our linear model,

$$\mathrm{Var}(\mathbf{p}_j) = \beta_j^2 \frac{p_j^2}{z_j^2} \sum_{r:j\in r} \mathrm{Var}(\mathbf{x}_r) \tag{A.61}$$

$$\leq \beta_j^2 \frac{p_j^2}{z_j^2} \sum_{r:j\in r} b_r x_r^2 (1 + O(|\mathbb{S}(n)|^{-(1/2)}). \tag{A.62}$$

Thus,

$$\mathrm{Var}(\mathbf{p}_j) \leq p_j^2 \sum_{r:j\in r} \frac{x_r^2}{z_j^2} O(1) \tag{A.63}$$

$$\leq p_j^2 \gamma(n) \sum_{r:j\in r} \frac{x_r}{z_j} O(1)$$

$$= p_j^2 O(|\mathbb{S}(n)|^{-1}), \tag{A.64}$$

as required. ■

Proof of Lemma 2: As in the proof of Lemma 1, $\gamma(n) = O(|\mathbb{S}(n)|^{-1})$, and so Theorem 2 tells us that for $r \in \mathbb{S}(1)$ and $r(n) \in \mathbb{S}(n)$ with $r = s(r(n))$,

$$\mathrm{Var}(\mathbf{x}_{r(n)}) = b_r x_{r(n)}^2 \frac{1}{2(1-q_{r(n)})}(1 + O(|\mathbb{S}(n)|^{-(1/2)}), \tag{A.65}$$

as required. That $\mathrm{Var}(\mathbf{q}_{r(n)}) = q_{r(n)}^2 O(|\mathbb{S}(n)|^{-(1/2)})$ follows directly from Lemma 3. ■

https://doi.org/10.1017/S0269964809990076 Published online by Cambridge University Press