MaxNet: Faster Flow Control Convergence

Bartek P. Wydrowski, *Member, IEEE,* Lachlan L. H. Andrew, *Member, IEEE,*
Iven M. Y. Mareels, *Fellow, IEEE*.

***Abstract*— MaxNet is a distributed congestion control architecture in which only the most severely bottlenecked link on the end-to-end path generates the congestion signal that controls the source rate. This is unlike SumNet networks, such as the current Internet or REM, where all of the bottlenecked links on the end-to-end path add to the congestion signal (by packet marking or dropping). Previously, we have shown that MaxNet results in Max-Min like rate allocation and is stable for arbitrarily large networks. In this paper we analyze the small-signal convergence speed of MaxNet. We show that MaxNet is able to converge faster than the SumNet architecture. Faster convergence results in less delay jitter, higher utilisation and lower buffer sizes. Furthermore, we show that MaxNet decouples the control, so that each pole position depends only on parameters of one bottleneck link and of the sources controlled by that bottleneck, enabling optimal pole placement.**

***Index Terms*— Network Flow Control, Congestion Control, Stability, Scalability.**

I. Introduction

THE problem of network flow control is to control source rates so that link capacities are utilised. For Internet-like networks, where links and sources can only have local information, the challenge is to control the source rates in a fully distributed manner. There are many different possible source rate allocations that fulfill the requirement of utilising the capacity. It has been shown [1] that the source rate allocation achieved by TCP Reno maximises a utility function. Most other flow control schemes also maximise (different) utility functions [1]. In [2], we showed that it is possible to achieve a different, fairer rate allocation by altering the way the network signals congestion information.

Models of Internet-like networks control the source rate by a scalar feedback congestion signal. This signal is generated by aggregating the congestion prices of links on the end-to-end connection path of the source. For networks which maximise a utility function, including TCP Reno networks, the signal is aggregated by summing all of the link prices on the path. We refer to these networks as SumNets. In [2], we introduced MaxNet, where the aggregation function is Max, and only the maximum link price along the connection path controls the source rate. In [2], we showed that MaxNet results in Max-Min fairness for sources with homogeneous demand functions.

In [3], the gains required at sources and links were found such that MaxNet remains stable for arbitrary capacities, delays and routing. Importantly we showed that MaxNet can achieve this robust stability without having to explicitly measure the number of bottleneck links on the end-to-end path.

Whilst the steady state rate allocation regime, whether it optimizes utility or fairness, is an important property, the time taken to converge to each steady state regime is a separate and also important issue. Convergence time of a network congestion control scheme is a key property that impacts on the Quality of Service of a network. A slow response results in long traffic transients which are responsible for packet delay, delay-jitter, under-utilisation and buffer-overflow. Reducing the duration and overshoot of transients improves these performance measures and makes smaller buffer sizes possible.

In this paper we consider the convergence time of MaxNet. The first part of our investigation involves a local analysis for which we develop a small-signal linearized model. Using this model we determine the position of poles which determine the convergence time of MaxNet. We compare the convergence time of MaxNet with a lower bound of the performance of SumNet. We prove that a faster pole placement is possible with MaxNet than with SumNet. This makes it possible for MaxNet to achieve better QoS performance. In the second part of this paper we investigate the global performance by simulating the full non-linear system, and relate the local analysis to these simulation results.

In Section II we start by briefly describing the framework in which both MaxNet and SumNet architectures fit, and Section III recalls the steady state rate allocation properties. Sections IV and V describes the control model whose root loci is discussed in Section VI. In Section VII MaxNet convergence time is determined

The authors are with the ARC Special Research Centre for Ultra-Broadband Information Networks. Department of Electrical and Electronic Engineering, The University of Melbourne, Australia. This work has been supported by the Australian Research Council.

{b.wydrowski,l.andrew,i.mareels}@ee.mu.oz.au

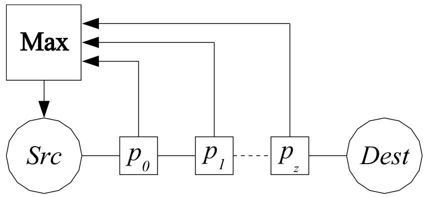

Fig. 1. MaxNet logical feedback loop.

and in Section VIII a bound on SumNet convergence time is obtained. In Section IX numerical results for the convergence time of a practical system are obtained and the performance of MaxNet and SumNet is compared.

II. CONTROL ARCHITECTURE

The following is a brief overview of the MaxNet and SumNet control algorithms. For a fuller description, see [2], [4]. In both MaxNet and SumNet networks, a single congestion signal, q_i, is communicated to source i summarising the prices, p_l, of all links, l, on the end-to-end transmission path, L_i. In MaxNet the congestion signal is the maximum of all link prices,

$$q_i = \max\{p_l : l \in L_i\}, \tag{1a}$$

as illustrated in Figure 1. In contrast, SumNet uses the sum

$$\hat{q}_i = \sum_{l \in L_i} p_l. \tag{1b}$$

(Throughout this paper, variables with a hat pertain to SumNet, and the corresponding variables without a hat pertain to MaxNet.)

The behaviour of source i is governed by an explicit demand function, $D_i(\cdot)$, such that its transmit rate is

$$x_i = D_i(q_i) \tag{2}$$

for a congestion signal q_i. The Active Queue Management (AQM) algorithm in a router sets the price of an outgoing link according to the well studied integrator process [4]:

$$p_l(t+1) = p_l(t) + (y_l(t) - c_l)\varphi_l, \tag{3}$$

where $y_l(t) = \sum_{i:l \in L_i} x_i(t)$ is the aggregate arrival rate for link l at time t, φ_l is the control gain and c_l is the target capacity of link l which is related to its physical capacity, C_l by the target utilisation, $0 < \mu_l < 1$, such that $c_l = \mu_l C_l$.

III. STEADY STATE RATE ALLOCATION

In steady state, when $p_l(t+1) = p_l(t)$, the rate allocation to source i, x_i will depend on the magnitude of its demand function relative to those of other sources sharing the maximum bottleneck link, l, of source i. If T_l is the set of sources traversing link l and link l has the highest price on i's path, then the rate allocation to i is

$$x_i = C_l \frac{D_i(q_i)}{\sum_{k \in T_l} D_k(q_k)}. \tag{4}$$

MaxNet can achieve a Max-Min fair rate allocation. A vector of rates, x, is defined as Max-Min fair if, for every feasible rate vector r with $r_i > x_i$ for some source i, there exists a source k such that $x_k \leq x_i$ and $x_k > r_k$. Put simply, a rate vector is Max-Min fair if it is feasible and no flow can be increased while maintaining feasibility without decreasing a smaller or equal flow.

The following proposition is proved in [2].

Proposition 1: Let $D(\cdot)$ be positive, continuous and decreasing. The rate allocation for a MaxNet network of homogeneous sources with $x_i = D(q_i)$ is Max-Min fair.

As discussed in the introduction and shown in [1], the rate allocation for SumNets optimizes the total utility of all sources in the network, which is generally not the same rate allocation as MaxNet. However, both equilibria are "optimal", in their own senses, and both have rate allocations within the network capacity, resulting in negligible queueing. It is only during transients that flows experience queueing delays, jitter and packet loss. Thus the convergence time can be compared as a separate issue from rate allocation. In the following sections we will develop a linearised model of both networks and later compare their small-signal convergence properties around their respective equilibrium points.

IV. MAXNET CONTROL SYSTEM MODEL

In this section we will describe the MaxNet model from [3] which will be used in subsequent sections to investigate convergence time. This model makes a number of simplifications of the network. The first is to use a fluid-flow approximation of the packet based information flow. The second simplification is that the global non-linear system is linearized about its equilibrium point. MaxNet contains two sources of non-linearity. The first is the Max operation itself and the second is the non-linear demand function, D. These will be linearised separately.

Using these simplifications, the network is represented as a multi-variable control system, shown in Figure 2. Note that Figure 2, for illustration only, shows a large-signal source, and small-signal links and network. The interconnection of sources with links is piecewise linear,

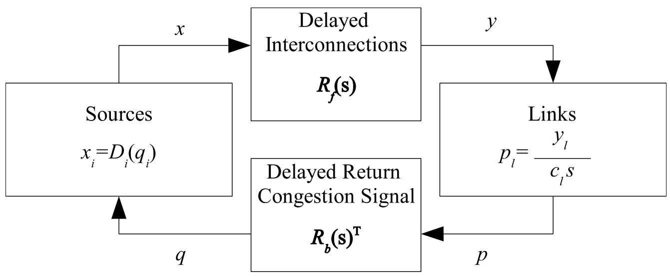

Fig. 2. Flow Control Structure.

due to the Max operation. It is described in the Laplace domain by forward and backward routing matrices. The matrices specify the interconnection and the delay incurred in signal flow from source to link and vice versa. The forward routing matrix is

$$[\overline{R}_f(s)]_{l,i} = \begin{cases} e^{-\tau^f_{i,l}s} & \text{if source } i \text{ uses link } l \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

where $\tau^f_{i,l}s$ is the forward delay between source i and link l. Note that the bar notation in $\overline{R}_f(s)$ indicates it has a row for every link in the network. We later reduce this to a matrix representing only bottleneck links, which does not have the bar.

Let n_i be the bottleneck link that controls source i. Then the backward routing matrix depends on n, and is given by

$$[\overline{R}_b(s;n)]_{l,i} = \begin{cases} e^{-\tau^b_{i,l}s} & \text{if } n_i = l \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

Note that the round-trip time of source i's connection is $\tau_i = \tau^f_{i,l} + \tau^b_{i,l}$. Let L be the number of links in the network. Without loss of generality, order the link prices such that

$$p_1 \geq p_2 \geq \ldots \geq p_L. \quad (7)$$

The backward routing matrix remains static over a period where the variations in link prices do not change the ordering of link prices (7). The overall multi-variable feedback loop in the configuration of Figure 2 is

$$\overline{y}(s) = \overline{R}_f(s)x(s) \quad (8)$$
$$q(s) = \overline{R}_b(s;n)^T\overline{p}(s). \quad (9)$$

We can construct a small signal model as in [4]. Consider small perturbations around equilibrium, $x = x_0 + \delta x$, $\overline{y} = \overline{y}_0 + \delta\overline{y}$, $\overline{p} = \overline{p}_0 + \delta\overline{p}$, $q = q_0 + \delta q$, where subscript 0 denotes a steady state value and prefix δ denotes a perturbation. Note that the bar notation still denotes variables that contain non-bottleneck links and $\delta\overline{p_l}$ is only non-zero for bottleneck links. Note also that when all link prices are distinct, the vector of bottlenecks, n, is unchanged by a sufficiently small perturbation. In this case, the small signal model does not explicitly involve n. This is the first linearisation. Form the vectors $\delta p(s)$, $\delta y(s)$ and the matrices R_f, R_b by eliminating the elements (or rows) corresponding to non-bottleneck links. This gives the reduced small signal model

$$\delta y(s) = R_f(s)\delta x(s) \quad (10)$$
$$\delta q(s) = R_b(s)^T\delta p(s), \quad (11)$$

To achieve stable control for networks of arbitrary dimensions, the gains that sources and links introduce need to be prescribed as detailed in [3]. The second linearisation replaces the demand function of a source by a small-signal gain at the source. That gain, between a perturbation in δq_i and the resulting perturbation in δx_i, is

$$\kappa_i = D'_i(q_i). \quad (12)$$

For robust stability, this gain must be scaled such that

$$\kappa_i = \frac{\alpha_i x_{0i}}{\tau_i}. \quad (13)$$

The selectable parameter $\alpha_i \in (0,1)$ controls the magnitude of the demand function to reflect the source's need for capacity. The term τ_i makes the stability robust to delay. To make stability robust to the number of sources, a gain x_{0i}/c_l is introduced in the closed-loop, with the x_{0i} component put into the source and the $1/c_l$ component in the link as $\varphi_l = 1/c_l$.

Note that (12) implicitly assumes a static demand function. As discussed in [4], the requirement (13) determines the shape of the static demand function. However, recent work in [5] provides dynamic source algorithms which allow arbitrary demand functions, whilst preserving the control gain required for robust stability. They separate the high-frequency gain AC from the DC gain.

In the Laplace domain, the integrator AQM of (3) with the required gain between the coupling of δp_l and δy_l is

$$\delta p_l = \frac{1}{c_l s}\delta y_l. \quad (14)$$

The open-loop transfer function for the small signal MaxNet model is

$$H(s) = \frac{1}{s}R_f(s)\mathcal{K}R_b(s)^T\mathcal{C}, \quad (15)$$

where

$$\mathcal{K} = \text{diag}(\kappa_i), \qquad \mathcal{C} = \text{diag}\left(\frac{1}{c_l}\right). \quad (16)$$

V. SumNet Control System Model

In this section we will describe the model from [4] for a SumNet network. There are many similarities to the MaxNet model, but also some important differences to highlight. Recall that the hat symbol identifies SumNet variables which have a related variable in MaxNet.

For SumNet the forward routing matrix is just the same as the MaxNet case

$$[\overline{R}_f(s)]_{l,i} = \begin{cases} e^{-\tau^f_{i,l}s} & \text{if source } i \text{ uses link } l \\ 0 & \text{otherwise.} \end{cases} \tag{17}$$

The backward routing matrix, which describes the flow of congestion information from each link back to sources, becomes independent of the current transmission rates:

$$[\hat{\overline{R}}_b(s)]_{l,i} = \begin{cases} e^{-\tau^b_{i,l}s} & \text{if source } i \text{ uses link } l \\ 0 & \text{otherwise.} \end{cases} \tag{18}$$

This gives

$$q(s) = \hat{\overline{R}}_b(s)^T \overline{p}(s). \tag{19}$$

The small signal variables also take on the hat notation: $\hat{x} = \hat{x}_0 + \delta\hat{x}$, $\hat{\overline{y}} = \hat{\overline{y}}_0 + \delta\hat{\overline{y}}$, $\hat{\overline{p}} = \hat{\overline{p}}_0 + \delta\hat{\overline{p}}$, $\hat{q} = \hat{q}_0 + \delta\hat{q}$.

For SumNet, the routing matrices can again be reduced to contain only bottleneck links. These reduced matrices are applicable so long as the bottleneckes remain the same throughout the perturbations. A reduced small-signal model can then be written as

$$\delta\hat{y}(s) = \hat{R}_f(s)\delta\hat{x}(s) \tag{20}$$
$$\delta\hat{q}(s) = \hat{R}_b(s)^T\delta\hat{p}(s), \tag{21}$$

where the matrices $\hat{R}_f$, $\hat{R}_b$ and the vectors $\delta\hat{p}(s)$, $\delta\hat{y}(s)$ are obtained by eliminating the rows corresponding to non-bottleneck links.

To achieve stable control for networks of arbitrary dimensions, the gains that sources and links introduce need to satisfy the bounds detailed in [4]. For SumNet, a source i requires a gain $\hat{\kappa}_i$ of

$$\hat{\kappa}_i = \frac{\hat{\alpha}_i \hat{x}_{0i}}{M_i \tau_i}, \tag{22}$$

where M_i is the number of controlling bottleneck links on the end-to-end path, and $\hat{\alpha}_i \in (0, 1)$ is again an adjustable parameter.

Note that MaxNet has the advantage over SumNet of not requiring knowledge of M_i, the number of bottlenecked links on the end-to-end path of source i. SumNets require M_i to be estimated and communicated to the source in order to achieve stability under arbitrary network scaling [4]. Eliminating M_i has several advantages. Firstly, it removes the additional signaling infrastructure required to determine M_i, as proposed for SumNet [4]. To remain stable without this signaling infrastructure, SumNets must assume an upper-bound on M_i and have a slow conservative control policy. With MaxNet, M_i is always 1, which avoids the need for either signaling infrastructure or a conservative control policy.

In link price calculation algorithm under SumNet remains the same (3) as under MaxNet; only the marking process is different. The complete SumNet open-loop transfer function for the small signal model is then

$$\hat{H}(s) = \frac{1}{s} R_f(s)\hat{\mathcal{K}}\hat{R}_b(s)^T\mathcal{C}, \tag{23}$$

where

$$\hat{\mathcal{K}} = \text{diag}(\hat{\kappa}_i), \qquad \mathcal{C} = \text{diag}\left(\frac{1}{c_l}\right). \tag{24}$$

VI. Root Loci

Despite their non-linear nature, the small signal convergence behaviour of MaxNet and SumNet can be characterised by the positions of the dominant poles of their linearisation. The case described below is for MaxNet; the equations for SumNet are the same, with hats on the appropriate variables.

The closed-loop transfer function is

$$T(s) = G(s)(Is + G(s))^{-1}, \tag{25}$$

where

$$G(s) = sH(s). \tag{26}$$

The poles of $T(s)$ are values of s satisfying either of the equations

$$\det(I + H(s)) = 0 \tag{27a}$$
$$\text{eig}(H(s)) = -1. \tag{27b}$$

For non-zero poles, equivalent conditions are

$$\det(Is + G(s)) = 0 \tag{28a}$$
$$\text{eig}(G(s)) = -s. \tag{28b}$$

The root loci of MaxNet and SumNet have many similarities, but some important differences. The open loop transfer function of each has L poles at zero. In MaxNet, these correspond directly to the sources controlled by the L links. In SumNet, there is intrinsic coupling between the links, and it is not helpful to think of poles as belonging to particular links.

For very small, but positive, loop gain, the poles at the origin move left on the real line. Meanwhile, L infinite sets of poles appear with real part $-\infty$, and with imaginary parts uniformly spaced [6]. These poles move right in the complex plane as the loop gains are increased. Importantly, L of these poles move along the real axis. For MaxNet, it is once again possible

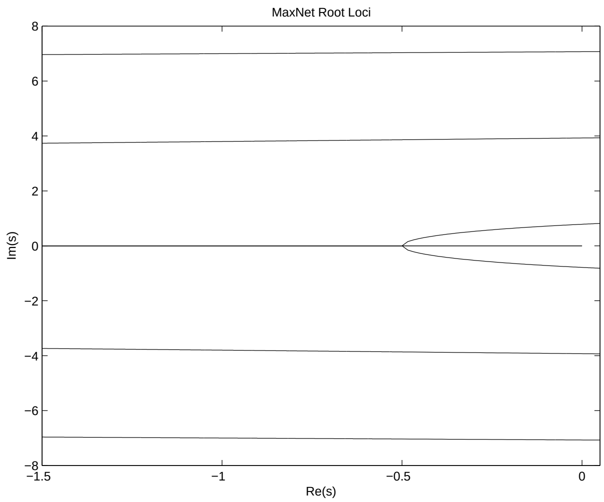

Fig. 3. MaxNet Root Loci.

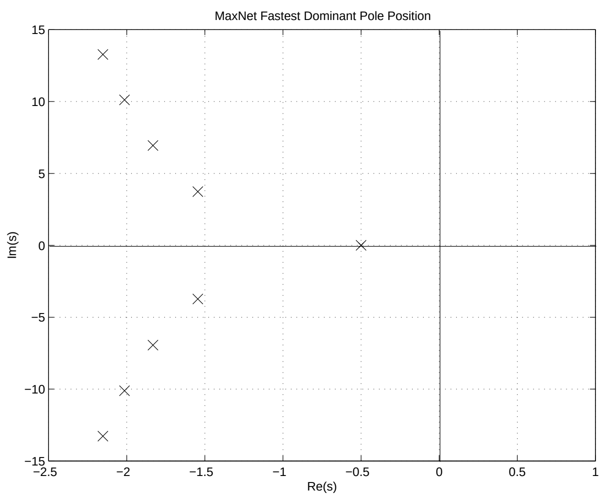

Fig. 4. MaxNet Fastest Dominant Pole Position.

to associate each pole with a specific link, while for SumNet, the poles can only associated with eigenvalues of a less structured matrix.

The point at which the rightmost of the poles coming from infinity meets the leftmost of the poles coming from zero is called a breakpoint. At this point, the two poles become a complex-conjugate pair, and start moving at right angles to the real axis, before going right again to eventually cross the imaginary axis and cause instability. As the gains increase further, subsequent pairs of real poles will meet at their respective break points, and also eventually become unstable. Under MaxNet, the pairs of poles which meet at break points always belong to the same link.

The value of the maximum real pole is minimised at the break point, when two real solutions of (27) coincide. At that point, s_l^*, not only are the left and right hand sides equal, but their derivatives are also equal [7].

VII. MaxNet Convergence Time

This section will derive bounds on the fastest possible convergence time of MaxNet; that is, the most negative value the real part of the dominant pole as the feedback gain is varied. These results hold for MaxNet networks with arbitrary topology, delay, number of sources and capacity.

Lemma 1: For sufficiently small gain, each link, l, introduces a pair of real poles. The minimum value achieved (by increasing the gain) of the maximum of these poles is the break point, s_l^*, which lies between $-1/t_{\max_l}$ and $-1/t_{\min_l}$, where $t_{\max_l}$ and $t_{\min_l}$ are the maximum and minimum round trip times (RTTs) of all of the sources being controlled by link l.

Proof: Since $G(s)$ is lower triangular under MaxNet, the eigenvalues are simply the diagonal elements, each of which corresponds to a particular link. Thus (28b) decouples, and we get one equation per link. From (15), poles associated with link l satisfy

$$-s = \sum_{k \in m_l} \frac{a_i e^{-\tau_i s}}{c_l}, \tag{29}$$

where

$$a_i = \frac{\alpha_i x_{0i}}{\tau_i}. \tag{30}$$

Each of these equations clearly has a real solution for sufficiently small a_i, establishing the first part of the lemma.

Differentiating (29) to find the break point, s_l^*, yields the condition

$$1 = \sum_{k \in m_l} \frac{a_i \tau_i e^{-\tau_i s_l^*}}{c_l} = \sum_{k \in m_l} b_i, \tag{31}$$

where

$$b_i = \frac{a_i \tau_i e^{-\tau_i s_l^*}}{c_l}. \tag{32}$$

Substituting (32) into (29) gives

$$-s_l^* = \sum_{k \in m_l} \frac{b_i}{\tau_i}. \tag{33}$$

Since $1 = \sum_{k \in m_l} b_i$, then (33) is a weighted sum of $1/\tau_i$. A weighted sum is between the maximum and minimum elements in the sum, giving

$$-\frac{1}{\tau_{\min_l}} \le s_l^* \le -\frac{1}{\tau_{\max_l}}. \tag{34}$$

■

Proposition 2: At the break point, s_l^* is the dominant pole associated with link l.

Proof: Except for the pole at the origin, all poles of (25) start with infinitely negative real part for low loop gain. Thus it suffices to show that, as the loop gain is increased, no complex pole crosses the line $\mathrm{Re}(s) = s_l^*$ before the real pole starting at $-\infty$ does.

Substituting $s = -\sigma + j\omega$ into (29) yields the implicit equation for pole positions at link l

$$\sum_{k \in m_l} a_k e^{\sigma \tau_k} (\cos(\omega \tau_k) - j \sin(\omega \tau_k)) = \sigma c_l - j \omega c_l. \quad (35)$$

Taking the real part of (35) gives

$$\sum_{k \in m_l} a_k e^{\sigma \tau_k} \cos(\omega \tau_k) = \sigma c_l. \quad (36)$$

Consider a line on the complex plane where $\mathrm{Re}(s) = \sigma$. If we fix the operating point for parameters x_{0k} and τ_k, then by (35), a is element-wise minimized when $\omega = 0$. Since complex poles begin at negative infinity for $a = 0$, and for the minimum $a_{\min}$ that satisfies (35) there is only a real pole on the line $\mathrm{Re}(s) = \sigma$, it follows that the real pole is the first to cross this line as the gain is increased. Complex poles, with $\omega \neq 0$, that cross this line have an element-wise higher a, and are therefore to the left of the real pole when the gain is $a_{\min}$. Thus the real pole at the break point will be the dominant pole for that link, since no complex poles have crossed to its right. ■

Remark 1: A key conclusion from this analysis is that because the links are independent, it is possible to adjust the control gains such that all links are simultaneously at their break points. That implies that the fastest operation of MaxNet is governed by poles satisfying (34).

VIII. Bound on SumNet Convergence Time

This section will show that, at least for the specific case analysed, MaxNet has a faster transient response than SumNet.

Due to the complexity of the SumNet analysis, we will consider a two link SumNet network only, where all sources have a common round trip time, τ, and only one source traverses both links. The assumption of a common round trip time is expected to favour SumNet by reducing the coupling between link. Thus we have no reason to believe that any other SumNet will be able to achieve a faster transient response than the equivalent MaxNet. It is sufficient to consider only the real pole, even though there may be complex poles which are slower, since this gives a lower bound for the transient response time. We do not consider whether there are complex poles further to the right of the fastest real position, since this would only result in a slower response.

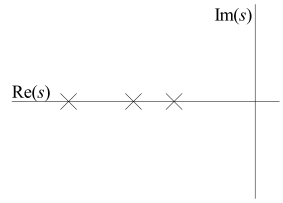

Fig. 5. SumNet Dominant poles with real poles at breakpoint (Case 1).

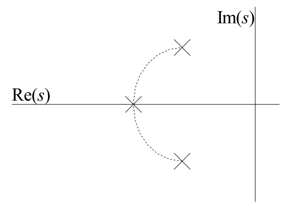

Fig. 6. SumNet Dominant poles with real poles at breakpoint (Case 2).

The SumNet system can be described by a 2×2 open-loop transfer function matrix, $\hat{H}$. Expanding (23) gives the elements of $\hat{H}$ as

$$\hat{H}_{ij}(s) = \frac{1}{s c_j} \sum_{k \in U_i \cap U_j} e^{-(\tau_{ki}^F + \tau_{kj}^B)s} \hat{a}_k, \quad (37)$$

where U_i is the set of sources that uses link i and

$$\hat{a}_k = \frac{\hat{\alpha}_k x_{0k}}{M_k \tau_k}. \quad (38)$$

The following lemmas are proved in the appendix.

Lemma 2: For a two link SumNet, where only one source traverses both links, and all sources have the same RTT τ, the unique break point is at $-1/\tau$.

Lemma 3: Unless $\hat{\alpha}_k = 0$ for all k, $\hat{G}(s)$ for a two link SumNet does not have a repeated eigenvalue for real s.

Combining these two lemmas gives the very significant result that there must be a pole to the right of $-1/\tau$. If the first pole pair has reached the break point, $-1/\tau$, then the other pair must not have met, and the pole which started at the origin must be to the right of $-1/\tau$, as shown in Figure 5. Conversely, if the second pole pair is at the break point, then the first pole pair must already have split, and moved to the right, as shown in Figure 6. Therefore SumNet must have a slower transient response than MaxNet.

IX. Numerical results

In this section we simulate the full non-linear SumNet and MaxNet networks to compare their transient response speeds. The results of this section serve to

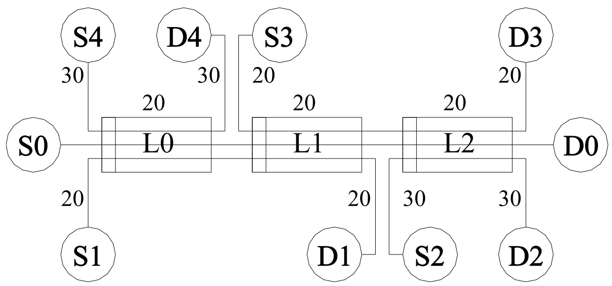

Fig. 7. Network Simulated.

give evidence that the small signal linearized properties proven analytically in the previous sections are relevant to the practical non-linear system.

The system simulated in this section is intended to reflect a physically realisable system. Whilst it may be possible to devise a control strategy where each source measures network properties and tunes its own gains (equation (22) for SumNet or (13) for MaxNet) to optimize transient speed, an online algorithm to achieve this is not trivial. In this paper we consider a practical strategy where all sources use the same demand function. We simulate sources with the same static demand function

$$x_i(t) = x_{\max} e^{-\rho q_i(t)} \tag{39}$$

where ρ is a network wide parameter and $x_{\max}$ is the maximum transmission rate. A similar demand function was introduced in [4], and was shown to be able to satisfy the gain requirements (13). For MaxNet, the parameter ρ relates to the small-signal source gain (13) such that

$$\rho = \frac{\alpha_i x_{0i}}{\tau_i} \tag{40}$$

and for SumNet the equivalent relationship is with (22)

$$\rho = \frac{\hat{\alpha}_i x_{0i}}{M_i \tau_i} \tag{41}$$

Note that ρ may be tuned to improve transient performance. This strategy will in general not result in the fastest possible transient response for MaxNet or SumNet, as the poles are not necessarily placed at their closet position to the break-points. Nevertheless it allows us to demonstrate some important properties.

A small network of 5 sources and 3 links, shown in Figure 7, is simulated using both SumNet and MaxNet congestion signaling. Sources $S0 \dots S4$ transmit to destinations $D0 \dots D4$ respectively.

We model traffic by a fluid flow approximation, that is, the source transmission rate and congestion price are continuous. At each time step in this discrete time simulation, the flow rate values and price feedback move

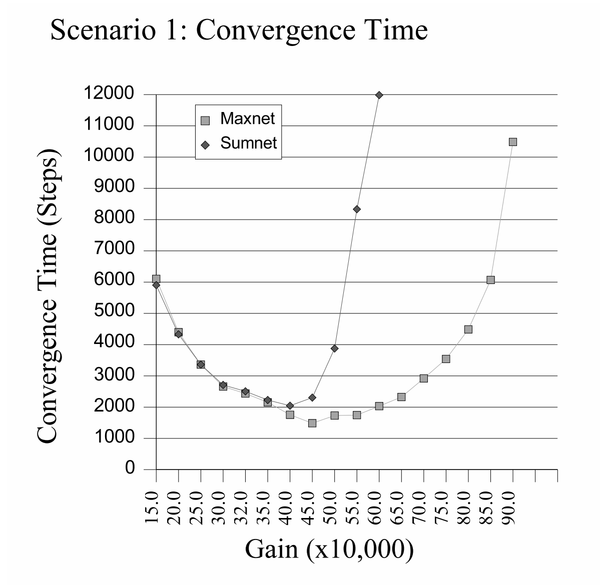

Fig. 8. Simulation scenario 1: SumNet and MaxNet Convergence Time.

one unit along in the forward and backward delay paths between sources and links. Acknowledgements are assumed to traverse the same links in the reverse direction, and consume negligible bandwidth. The numbers near each line in Figure 7 represent delays, in simulation time step units. Note that every source has a RTT of 160 units, and for all sources $x_{\max}$ is set to 15. The MaxNet or SumNet link control law (3) is at the head of the link, represented by the rectangle inside each link in Figure 7.

In the simulation scenarios, we assume that the best-effort congestion controlled traffic is receiving only a portion of the link's physical capacity. This represesents the situation of having higher-priority constant-bit-rate (CBR) traffic occupying some capacity. We simulate two scenarios with different proportions of CBR traffic and different link capacities.

In Scenario 1, the physical link capacities are $c_0 = 5$, $c_1 = 3$ and $c_2 = 5$. To generate a transient, we assume that initially the capacities available to best-effort traffic are 5, 3 and 2 at links $L0$, $L1$ and $L2$ respectively. A transient occurs when the CBR traffic source using $L2$ stops and the available capacities become 5, 3 and 5. Throughout the whole experiment the link gains are, as stipulated in [4], $1/c_l$ such that $\varphi_0 = 1/5$, $\varphi_1 = 1/3$ and $\varphi_2 = 1/5$.

Scenario 2 is the same as scenario 1 except that the physical link capacity of link 1 is $c_1 = 12$, and correspondingly $\varphi_1 = 1/12$. The available capacities again start at 5, 3 and 2, and link 2's available capacity increases to 5.

The transient response metric used is the settle time,

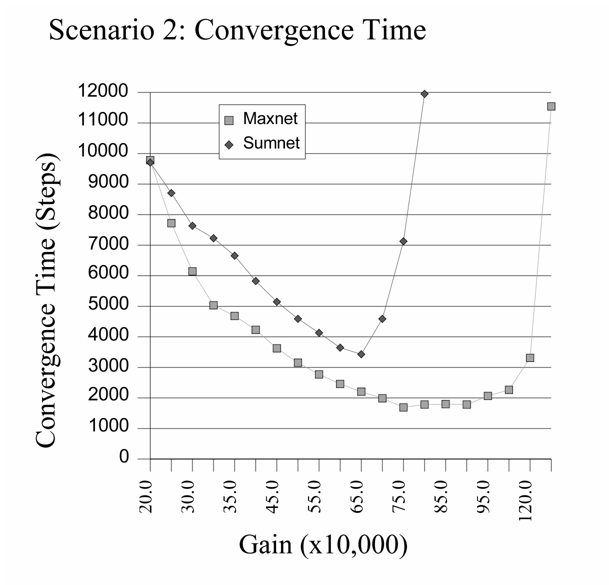

Fig. 9. Simulation Scenario 2: SumNet and MaxNet Convergence Time.

which is the time from the change of capacity to when the last source is within $\pm 1\%$ of its final value. The settle time is measured in simulation time steps. Figures 8 and 9 show this convergence time for both SumNet and MaxNet for gains $\rho = 0.0015$ to $\rho = 0.009$ for scenario 1, and $\rho = 0.002$ to $\rho = 0.012$ for scenario 2.

Note that for both SumNet and MaxNet there is an optimal ρ where the convergence time is the fastest. If ρ is increased beyond this point, a more oscillatory response results and the convergence time increases. Eventually the system becomes unstable and the convergence time is infinite. If ρ is decreased below the optimum, the system becomes over-damped and the convergence time increases. The range for ρ studied was limited by the simulation period, which needs to increase for slower convergence times.

Consider the MaxNet poles for Scenario 1. As a single gain is used for all of the sources, not all of the poles can be positioned on the breakpoint simultaneously. The envelope of the response decays most quickly when the rightmost complex pole pair, which is moving right on the root-loci, has the same real part as the rightmost real pole moving left. For Scenario 1, this optimum point was obtained numerically and is shown in Figure 10. This numerical analysis indicates that the envelope decays fastest when $\rho = 0.0032$. The fastest response in the simulations occured when $\rho = 0.0045$. This slight difference is to be expected, since the oscillatory response to the complex pole pair is usually well within its envelope, decreasing the settline time as defined above. Figures 11 and 12 illustrate the movement of the poles as ρ is smaller or greater than the optimal value. Figure 11 omits the real pole far to the left.

For both scenarios, the fastest convergence time for SumNet is significantly slower than the fastest time for MaxNet. In Scenario 1, SumNet took 1.38 times as long as MaxNet to converge, and in Scenario 2, SumNet took 2.03 times as long. Also note that MaxNet had wide set of gains for which the convergence time was near optimal. This is important as it improves the robustness to small errors in estimates of parameters such as the round trip time. It allows satisfactory performance to be obtained by conservative gains which do not risk causing instability.

The convergence time results for SumNet indicate that right-most real poles of SumNet are more widely spaced than MaxNet. The results suggest that as the gain is increased, a SumNet pole becomes complex earlier than any MaxNet pole; at the same time, a dominant real pole is further to the right than MaxNet. This supports the analysis which determined that SumNet poles could not be made coincident.

Note that increasing the tolerance of the settling condition above $\pm 1\%$ results in the non-linear behaviour having a greater significance on the settling time results.

Figures 13 and 14 show the fastest source rate transient for SumNet and MaxNet in Scenario 1. For SumNet, the fastest transient occurs with $\rho = 0.0040$, with a settling time of 2045 steps. For MaxNet, the fastest transient occurs with $\rho = 0.0045$, with a settling time of 1484 steps.

X. Conclusion

This paper has shown that MaxNet flow control has favourable convergence properties compared with traditional SumNet flow control. For small perturbations from the operating point, MaxNet permits a pole placement that has a faster transient response than that possible with SumNet. Numerical results for the complete nonlinear system confirm the conclusions drawn from the analysis of the linear model. Further work will be undertaken to extend these results to more general SumNet topologies and more general nonlinear stability analysis.

Appendix

The proof of Lemma 2 is as follows.

Proof:

This proof will again use the fact that, at the point a which a pair of real poles meet and become complex conjugates, the derivative of $X = \det(I + \hat{H}(s))$ with respect to s is zero. It also uses the fact that values of s for which $dX/ds = 0$ but $X \neq 0$ are not breakpoints.

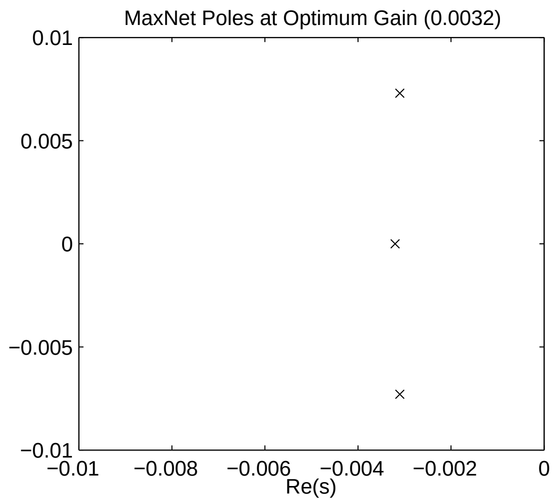

Fig. 10. MaxNet Poles for Scenario 1 at Fastest position.

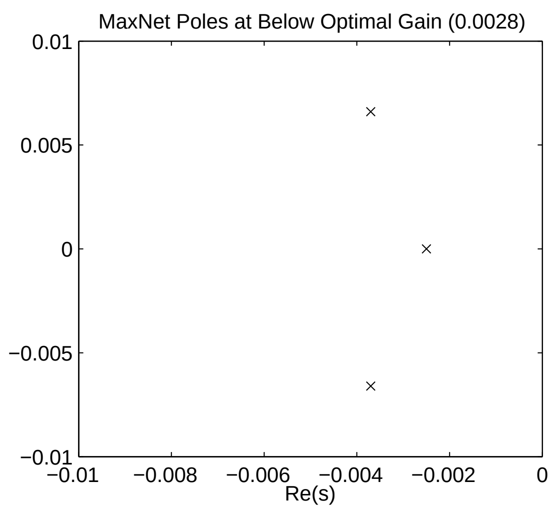

Fig. 11. MaxNet Poles for Scenario 1 gain too low.

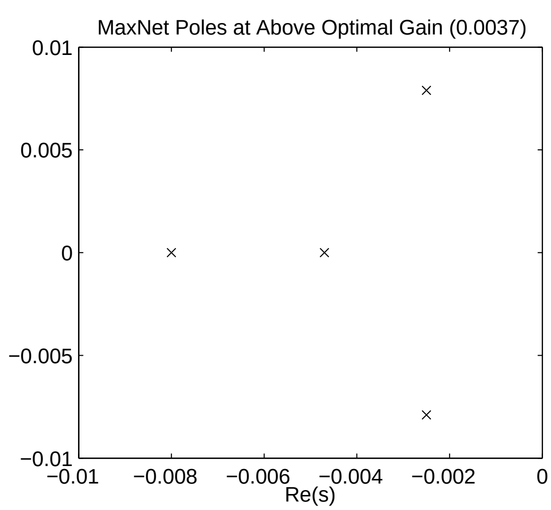

Fig. 12. MaxNet Poles for Scenario 1 gain too high.

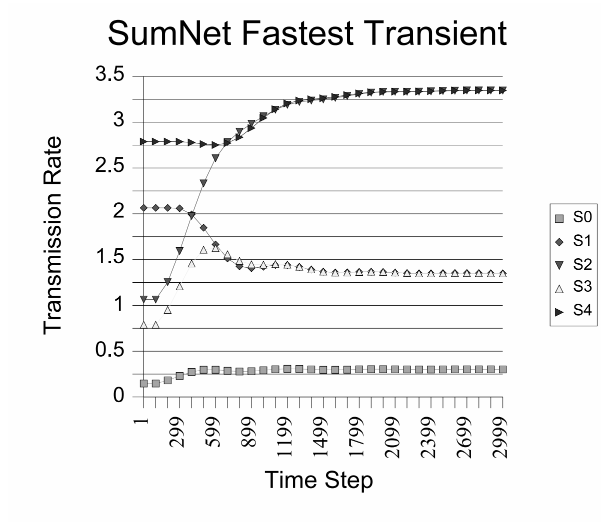

Fig. 13. Scenario 1 SumNet Fastest Transient.

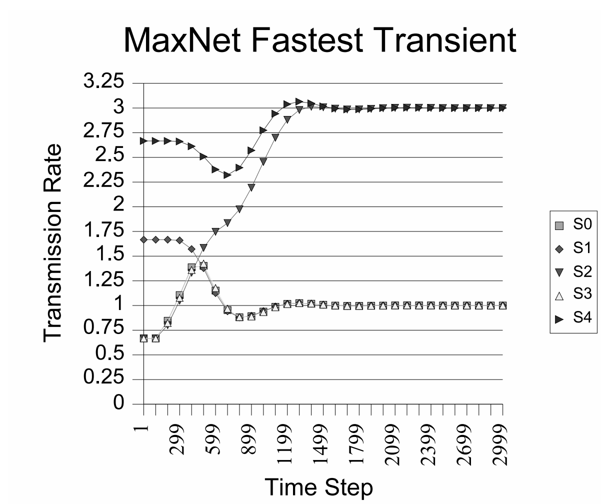

Fig. 14. Scenario 1 MaxNet Fastest Transient.

Since the round trip time of each route is equal, $\tau_{ki}^F + \tau_{ki}^B = \tau$ for all k, (37) implies

$$\hat{H}_{ii}(s) = \frac{1}{sc_i} \sum_{k \in U_i} e^{-\tau s} \hat{a}_k = \frac{e^{-\tau s}}{s} R_i, \qquad (42)$$

for $i = 1, 2$, where $R_i = \sum_{k \in U_i} \hat{a}_k / c_i$.

The fact that only one source traverses both links implies that, for $i \neq j$, the sum in (37) contains a single term. Without loss of generality, let that source be source 1. Then

$$\begin{aligned} \hat{H}_{12}(s)\hat{H}_{21}(s) &= \frac{\hat{a}_1^2}{s^2 c_1 c_2} e^{-(\tau_{11}^F + \tau_{12}^B + \tau_{12}^F + \tau_{11}^B)s} \\ &= \frac{e^{-2\tau s}}{s^2} R_3, \end{aligned} \qquad (43)$$

where $R_3 = \hat{a}_1^2 / (c_1 c_2)$.

From (27a), the 2×2 SumNet poles are at $X = 0$, where

$$X = \hat{H}_{11}(s)\hat{H}_{22}(s) - \hat{H}_{12}(s)\hat{H}_{21}(s) + \hat{H}_{11}(s) + \hat{H}_{22}(s) + 1.$$

Substituting (42) and (43) into this gives

$$X = \frac{e^{-2\tau s}}{s^2}(R_1R_2 - R_3) + \frac{e^{-\tau s}}{s}(R_1 + R_2) + 1. \quad (44)$$

Differentiating (44) to find the break point gives

$$\begin{aligned} \frac{dX}{ds} &= (\tau s + 1)\left(\frac{2e^{-2\tau s}}{s^3}(R_1R_2 - R_3) - \frac{e^{-\tau s}}{s^2}(R_1 + R_2)\right) \\ &= (\tau s + 1)A\frac{2e^{-\tau s}(R_1R_2 - R_3) - s(R_1 + R_2)}{s^3}, \quad (45) \end{aligned}$$

where

$$A \equiv \frac{e^{-\tau s}}{s}.$$

This derivative, (45), is zero when $s = -1/\tau$. The root locus occupies the entire negative real axis, and thus $s^* = -1/\tau$ corresponds to an actual breakpoint. It remains to show that there are no other breakpoints.

Assume, with a view to obtaining a contradiction, that there is another breakpoint, s'. At s', the final factor of (45) must be zero. That implies

$$\frac{e^{-\tau s}}{s} = \frac{R_1 + R_2}{2(R_1R_2 - R_3)}. \quad (46)$$

Substituting (46) into (44) gives

$$X = \frac{3(R_1 + R_2)^2}{4(R_1R_2 - R_3)} + 1.$$

But the left had side is positive, since R_3 is one of the terms in the positive-term sum R_1R_2, and so $X \neq 0$. Thus s' is not a pole, and cannot be a breakpoint. This establishes the result.

■

The proof of Lemma 3 is as follows.

Proof: The eigenvalues of $\hat{G}(s)$ are

$$\hat{G}_{11}(s) + \hat{G}_{22}(s) \pm \sqrt{(\hat{G}_{11}(s) - \hat{G}_{22}(s))^2 + 4\hat{G}_{21}(s)\hat{G}_{12}(s)}. \quad (47)$$

Equating the two solutions to (47) gives the condition for poles being co-incident as

$$0 = (\hat{G}_{11}(s) - \hat{G}_{22}(s))^2 + 4\hat{G}_{21}(s)\hat{G}_{12}(s). \quad (48)$$

When s is real, $\hat{G}_{ij}(s)$ is also real. A real solution to (48) is only possible when $\hat{G}_{21}(s)\hat{G}_{12}(s) \leq 0$. For SumNet, $\hat{G}_{21}(s)\hat{G}_{12}(s) > 0$ for real s, unless $\hat{\alpha}_k = 0$ for all k. Thus (48) cannot be satisfied. ■

REFERENCES

[1] S. H. Low, "A duality model of TCP and queue management algorithms," in *Proc. ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, (Monterey, CA), 2002.

[2] B. Wydrowski and M. Zukerman, "MaxNet: A congestion control architecture for maxmin fairness," *IEEE Commun. Lett.*, vol. 6, pp. 512–514, Nov. 2002.

[3] B. Wydrowski, L. L. H. Andrew, and M. Zukerman, "MaxNet: A congestion control architecture for scalable networks," *IEEE Commun. Lett.*, Accepted for publication.

[4] F. Paganini, J. C. Doyle, and S. H. Low, "Scalable laws for stable network congestion control," in *Proc. IEEE Conf. Decision Contr. (CDC)*, (Orlando, FL), 2001, pp. 185–90.

[5] F. Paganini, Z. Wang, S. Low, and J. Doyle, "A new TCP/AQM for stable operation in fast networks," in *Proc. IEEE Infocom*, (San Francisco, CA), 2003.

[6] J. E. Marshall, H. Górecki, K. Walton, and A. Korytowski, *Time-Delay Systems: Stability and Performance Criteria with Applications*. New York, NY: Ellis Horwood, 1992.

[7] K. Ogata, *Modern Control Engineering*. Englewood Cliffs, NJ: Prentice-Hall, 1970.