Proceedings of the 2007 American Control Conference
Marriott Marquis Hotel at Times Square
New York City, USA, July 11-13, 2007

On the ergodicity of AIMD networks

Robert Shorten, Chris King, Fabian Wirth, Douglas Leith

***Abstract*— In this paper we study communication networks that employ drop-tail queueing and *Additive-Increase Multiplicative-Decrease (AIMD)* congestion control algorithms. A basic assumption in the study of such networks is that the underlying stochastic process is ergodic. In this paper we present a proof of ergodicity for such networks under very general assumptions.**

I. Introduction

In this paper we study communication networks that employ drop-tail queueing and *Additive-Increase Multiplicative-Decrease (AIMD)* congestion control algorithms. A basic assumption in the study of such networks is that the underlying stochastic process is ergodic. In this paper we present a proof of ergodicity under very general assumptions. Fairness formulae for the average congestion window size, and for throughput are also given. Our starting point is a recently derived algebraic model that captures the essential features of networks in which TCP is deployed. For convenience, a brief derivation of this model is given in the next section. We then show that it is possible to relate important network properties to the characteristics of the non-negative matrices that arise in the study of such communication networks under very general conditions. In particular, our approach allows us to prove that the network is ergodic under very general assumptions. Full proofs of our results are given in [1].

II. Nonnegative matrices and communication networks

Our basic tool used in this paper is a recently derived model of a bottleneck link servicing a fixed number of flows employing AIMD congestion control [2]. Before proceeding, we recall briefly the the notation adopted in this paper: $k \in \mathbb{N}$ denotes the event corresponding to instant that the network becomes congested for the k'th time; $w_i(k)$ is the value of the i'th flows window at the k'th event; $\alpha_i > 0$ is the increase factor of the i'th source; and β_i, The constant $\beta_i \in (0,1)$ is the multiplicative decrease factor of the i'th flow. The maximal round trip time of flow i is denoted by $RTT_{\max,i}$ and we set $\gamma_i := RTT^{-1}_{\max,i}$. Then, assuming small network buffers, and ignoring slow-start, it is shown in [2] that the dynamics of such a network can be modelled in a stochastic setting by studying products of matrices of a certain type. In particular, the dynamics of the entire network of sources at the k'th congestion event are described by

$$W(k+1) = A(k)W(k),\ A(k) \in \{A_1, ..., A_m\}. \qquad (1)$$

where $A(k) = R(k) + st(k)^T$, where $R(k) = diag([\beta_1(k), ..., \beta_n(k)]$, $\left(\sum_{j=1}^n \gamma_j \alpha_j\right) s^T = [\alpha_1, ..., \alpha_n]$, and $t(k)^T = [\gamma_1(1-\beta_1(k)), \ldots, \gamma_n(1-\beta_n(k))]$, The matrices in the switching set are constructed by setting some, but not all, of the β_i to 1. This gives rise to $m = 2^n - 1$ matrices associated with the system (1) that correspond to the different combinations of source drops that are possible.

III. Main results

It follows from (1) that $W(k) = \Pi(k)W(0)$, where $\Pi(k) = A(k)A(k-1)....A(0)$. Consequently, the behaviour of $W(k)$, as well as the network fairness and convergence properties, are governed by the properties of the infinite matrix product $\Pi(k)$. The objective of this section is to analyse the average behaviour of $\Pi(k)$ with a view to making concrete statements about important network properties. We begin by considering networks in which we assume that the manner in which sources are informed of congestion is independent of $W(k)$, and for which the probability of source i experiencing a drop is constant at each congestion event. We then develop results for networks in which this assumption is relaxed and for which the relevant probabilities depend on $W(k)$; namely, the probabilities are place dependent. Finally, under the assumption of ergodicity, we develop a number of results that apply to networks with general capacity constraints.

A. Constant drop-probabilities

Here, we make the following two assumptions that are observed to be empirically true in a variety of networks [2].

Assumption 3.1: The probability that $A(k) = A_i$ in (1) is independent of k and equals ρ_i.

Given the probabilities ρ_i for $i \in \{1, ..., 2^n - 1\}$, one may then define the probability Λ_j that source j experiences a backoff at the k'th congestion event as follows:

$$\Lambda_j = \sum \rho_i \,,$$

where the summation is taken over those i which correspond to a matrix in which the j'th source sees a drop.

Assumption 3.2: We assume that $\Lambda_j > 0$ for all $j \in \{1, ..., n\}$.

Joint first author, Hamilton Institute, NUI Maynooth, Ireland.
Joint first author, Dept. of Maths., Northeastern Univ., USA.
Dept. of Maths, University of Bremen, Germany
Hamilton Institute, NUI Maynooth, Ireland.

Authorized licensed use limited to: The Library NUI Maynooth. Downloaded on December 21, 2009 at 07:21 from IEEE Xplore. Restrictions apply.

Simply stated, Assumption 3.2 states that almost surely all flows must see a drop at some time.

Under the foregoing assumptions we have the following key result.

Theorem 3.1: Consider the stochastic system defined in the above preamble. Let $\Pi(k)$ be the random matrix product arising from the evolution of the first k steps of this system:

$$\Pi(k) = A(k)A(k-1)....A(0).$$

Then, the expectation of $\Pi(k)$ is given by

$$E(\Pi(k)) = (\sum_{i=1}^{m} \rho_i A_i)^k; \tag{2}$$

and the asymptotic behaviour of $E(\Pi(k))$ satisfies

$$\lim_{k\to\infty} E(\Pi(k)) \quad = \quad x_p y_p^T, \tag{3}$$

where the vector x_p is given by $x_p^T = \Theta(\frac{\alpha_1}{\Lambda_1(1-\beta_1)}, \frac{\alpha_2}{\Lambda_2(1-\beta_2)}, ..., \frac{\alpha_n}{\Lambda_n(1-\beta_n)})$, $y_p^T = (\gamma_1, ..., \gamma_n)$. Here $\Theta \in \mathbb{R}$ is chosen such that the capacity constraint is satisfied for all k.

Corollary 3.1: For given $W(0)$ define random variable $\overline{W}(k)$ with:

$$\overline{W}(k) \quad := \quad \frac{1}{k+1}\sum_{i=0}^{k} W(i).$$

Then expectation of $\overline{W}(k)$ is given by:

$$E(\overline{W}(k)) = \frac{1}{k+1}(I + E(A(1)) + \cdots + E(A(1))^k)W(0)$$

And since $E(A(1))^k \to x_p y^T$ as $k \to \infty$,

$$\lim_{k\to\infty} E(\overline{W}(k)) = x_p y^T W(0)$$

Note that the congestion window vector $W(k)$ converges, on average, to the unique value $\bar{W}_{ss} = \Theta x_p$. Window fairness is achieved, on average, when the vector x_p is a scalar multiple of the vector $[1, ..., 1]$; that is, when the ratio $\frac{\alpha_i}{\Lambda_i(1-\beta_i)}$ does not depend on i. Observe that fairness depends upon on the relative drop probability of each flow.

B. Place dependent drop-probabilities

We now relax the assumption that flows are dropped with constant probabilities at congestion events. Instead, we allow the drop probabilities to depend on the current state of the system, that is the set of current window sizes $W(k)$ at the kth congestion event.

Mathematically, our model is a discrete time Markov chain whose state space is the simplex $\mathcal{S} = \{w = (w_1, \ldots, w_N) : w_i \geq 0, \sum w_i = C\}$, where C is the link capacity. The state of the system $W(k)$ evolves according to the rule $W(k+1) = A(k)W(k)$, where $A(k) \in \mathcal{A}$ is chosen randomly using a place-dependent probability distribution on $\mathcal{A}$. Specifically, for each $w \in \mathcal{S}$ there is a probability distribution $p_1(w), p_2(w), \ldots, p_m(w)$ on $\mathcal{A}$ ($m = |\mathcal{A}| = 2^n - 1$), and $A(k)$ is chosen randomly with probability $P(A(k) = A_i) = p_i(W(k))$. With some mild assumptions on the drop probability functions we will show that this model is ergodic. In particular:

Assumption 3.3: The distribution $p_i(w)$ is uniformly Lipschitz on $\mathcal{S}$. That is, we assume there is a constant K such that for all $w, v \in \mathcal{S}$ and all $i = 1, \ldots, m$,

$$|p_i(w) - p_i(v)| \leq K\,||w - v||_1 = K\sum_{j=1}^{n} |w_j - v_j| \tag{4}$$

Assumption 3.4: The probability that all flows are dropped at a congestion event is non-zero in the interior of the simplex $\mathcal{S}$.

Assumption 3.5: The probability that the ith flow is dropped at a congestion event is zero at any point in $\mathcal{S}$ where $w_i = 0$.

The following theorem extends results in [3] and presents and provides an alternative approach to results in [4]. Our principal tools are taken from the exposition presented by Stenflo [5]. These results rely on contractivity properties of the matrices $A \in \mathcal{A}$ on the simplex $\mathcal{S}$. This is guatanteed if the matrices in $\mathcal{A}$ are column stochastic, see [3]. However, it is easy to see that the matrices in the set $\mathcal{A}$ are not column stochastic, unless the γ_i are all equal. To solve this problem note that the matrices in this set are simultaneously similar to a set of column stochastic matrices under the transformation $\Gamma = \mathrm{diag}[\gamma_1, ..., \gamma_n]$. Applying this state transformation we arrive at a situation in which the results in citeSten are applicable. In the following we will assume without loss of generality, that this transformation has been applied.

Theorem 3.2: Assume that the place-dependent drop probabilities $\{p_i(w)\}$ satisfy the assumptions listed above. Then (i) there is an attractive, unique stationary probability measure for the Markov process $\{W(k)\}$, (ii) for any continuous function $f(w)$ on $\mathcal{S}$, the conditional expectations $\mathrm{E}[f(W(k)) \,|\, W(0) = w]$ converge uniformly to a constant as $k \to \infty$, For any continuous function $f(w)$ on $\mathcal{S}$, the time average $\frac{1}{K}\sum_{k=1}^{K} f(W(k))$ converges almost surely to the ensemble average of $f(W)$ with respect to the stationary measure.

Theorem 3.2 may be pproved using the results of Isaac [6], Barnsley et al [7], and Stenflo [5], who established general conditions for ergodicity of Markov chains with place-dependent probabilities. Theorem 3.2 implies that the process $\{W(k)\}$ converges a.s. $k \to \infty$, and that the limiting distribution is independent of the initial conditions. This ergodic property allows us to relate time averages to ensemble averages, and hence to use pathwise calculations to compute average quantities. Using this method we will show that a version of the result (3) derived in the case of constant drop-probabilities continues to hold for the place-dependent model. Our result will involve the average window size for the ith flow computed only at the congestion events where it experiences a drop. To set up the notation, let

Authorized licensed use limited to: The Library NUI Maynooth. Downloaded on December 21, 2009 at 07:21 from IEEE Xplore. Restrictions apply.

$D(k) \subset \{1, \ldots, m\}$ denote the set of flows which experience a drop at the kth congestion event, and define

$$\theta_i(k) = \begin{cases} 1 & \text{if } i \in D(k) \\ 0 & \text{if } i \notin D(k) \end{cases} \tag{5}$$

Theorem 3.3: Under the conditions for Theorem 3.2, the following limits exist and are independent of initial conditions:

$$\langle w_i \rangle = \lim_{k\to\infty} \mathrm{E}[w_i(k) \,|\, i \in D(k)], \quad \lambda_i = \lim_{k\to\infty} \mathrm{E}[\theta_i(k)] \tag{6}$$

Furthermore these quantities are related by

$$\langle w_i \rangle = \frac{\alpha_i}{\lambda_i(1-\beta_i)} \mathrm{E}[T] \tag{7}$$

where $\mathrm{E}[T]$ is the average time between congestion events.

Comment : The quantity $\langle w_i \rangle$ is the average window size for the ith flow computed at the congestion events where this flow experiences a drop, and λ_i is the probability that the ith flow experiences a drop at a randomly selected congestion event. So (24) is directly related to Theorem 3.1 (3) obtained in the case of constant drop probabilities. Furthermore the result is independent of the details of how the flows are dropped at each event, and depends only on the probability of a drop for the ith flow.

Proof: For any $k \geq 1$

$$\mathrm{E}[w_i(k) \,|\, i \in D(k)] = \frac{\mathrm{E}[w_i(k)\,\theta_i(k)]}{\mathrm{E}[\theta_i(k)]} \tag{8}$$

Furthermore

$$\mathrm{E}[w_i(k)\,\theta_i(k)] = \mathrm{E}[w_i(k) \sum_{j\in\Delta_i} p_j(W(k))] \tag{9}$$

where $\Delta_i \subset \{1, \ldots, m\}$ is the list of all subsets for which the ith flow experiences a drop. Applying Theorem 3.2 we conclude that (9) converges uniformly to a value independent of initial conditions as $k \to \infty$. The same argument applies to the denominator in (8), hence the left side of (8) converges to a limit which we define to be $\langle w_i \rangle$. Similar reasoning applies to define λ_i.

Considering a sample path of the process, we see that the ith window sizes w_i, w_i' at two successive events where flow i is dropped are related by

$$w_i' = \beta_i w_i + \alpha_i T_i \tag{10}$$

where T_i is the time between these events. Define the long-run time averages

$$\langle w_i \rangle_K = \frac{\sum_{k=1}^K w_i(k)\theta_i(k)}{\sum_{k=1}^K \theta_i(k)}, \quad \langle T_i \rangle_K = \frac{\sum_{k=1}^K T(k)}{\sum_{k=1}^K \theta_i(k)} \tag{11}$$

where $T(k)$ is the time between the kth and $(k+1)^{th}$ congestion events. Then (10) implies

$$\langle w_i \rangle_K = \frac{\alpha_i}{1-\beta_i} \langle T_i \rangle_K + O(\frac{1}{K}) \tag{12}$$

where the error term $O(\frac{1}{K})$ takes care of the mismatches in the sums at $k = 1$ and $k = K$ (recall that w_i and $T(k)$ are uniformly bounded, so this term is bounded by a constant times $\frac{1}{K}$). Therefore (7) follows from part (iii) of Theorem 3.2, which states that time averages converge to ensemble averages, and hence $\langle w_i \rangle_K$ converges to $\langle w_i \rangle$ and $\langle T_i \rangle_K$ converges to $\frac{\mathrm{E}[T]}{\lambda_i}$. ■

Our next result involves the throughput for the ith flow, which is defined by the pathwise expression

$$\delta_i = \lim_{T\to\infty} \frac{1}{T} \int_0^T w_i(t)\, dt \tag{13}$$

Theorem 3.4: Under the conditions for Theorem 3.2, with probability one the expression (13) exists and is independent of the sample path, and is given by

$$\delta_i = \frac{\lambda_i(1-\beta_i^2)}{2\alpha_i} \mathrm{E}[T]\, \langle w_i^2 \rangle \tag{14}$$

where $\langle w_i^2 \rangle = \lim_{k\to\infty} \mathrm{E}[w_i(k)^2 | i \in D(k)]$. It satisfies the bounds

$$\langle w_i \rangle \frac{1+\beta_i}{2} \leq \delta_i \leq \langle w_i \rangle \frac{1+\beta_i}{2} (1 + \frac{\mathrm{VAR}[T]}{\mathrm{E}[T]^2}) \tag{15}$$

where $\mathrm{VAR}[T]$ is the variance of the time between congestion events.

Proof: We will write $\{\tau_1, \tau_2, \ldots\}$ to denote the times of the congestion events where flow i experiences a reduction, and $\{w_i(1), w_i(2), \ldots\}$ its window sizes at these events. The evolution equation for w between congestion events is

$$w_i(k+1) = \beta_i w_i(k) + \alpha_i(\tau_{k+1} - \tau_k) \tag{16}$$

Elementary calculations along the sample path show that $\int w_i(t)\, dt$ can be expressed as a sum of the squares of window sizes at congestion events where the ith flow is dropped. Ergodicity then relates this time average to the ensemble average, and this gives (14).

We now use upper and lower bounds on $\langle w_i^2 \rangle$ to derive (15). For the lower bound we just use

$$\langle w_i^2 \rangle \geq \Big(\langle w_i \rangle\Big)^2 \tag{17}$$

For the upper bound we square (16) and take the expected value to get

$$\begin{aligned} E[w_i(k+1)^2] = \beta_i 2E[w_i(k)^2] + \alpha_i^2 E[(\tau_{k+1} - \tau_k)^2] \\ +2\alpha_i\beta_i E[w_i(k)(\tau_{k+1} - \tau_k)] \end{aligned} \tag{18}$$

We use the bound $2w_i(k)(\tau_{k+1} - \tau_k) \leq x w_i(k)^2 + x^{-1}(\tau_{k+1} - \tau_k)^2$, which holds for every $x >> 0$; inserting into (18) gives

$$\begin{aligned} E[w_i(k+1)^2] \leq (\beta_i^2 + \alpha_i\beta_i x) E[w_i(k)^2] \\ +(\alpha_i^2 + \alpha_i\beta_i x^{-1}) E[(\tau_{k+1} - \tau_k)^2] \end{aligned} \tag{19}$$

Taking $k \to \infty$ and using the ergodic property gives

$$\langle w_i^2 \rangle \leq \frac{\alpha_i^2 + \alpha_i\beta_i x^{-1}}{1 - \beta_i^2 - \alpha_i\beta_i x} E[T^2] \tag{20}$$

Authorized licensed use limited to: The Library NUI Maynooth. Downloaded on December 21, 2009 at 07:21 from IEEE Xplore. Restrictions apply.

where T is the time between congestion events. Using the optimal value $x = \alpha_i^{-1}(1-\beta_i)$ we find

$$\langle w_i^2 \rangle \leq \frac{\alpha_i^2}{(1-\beta_i)^2} E[T^2] \tag{21}$$

Finally we use $E[T^2] = \mathrm{VAR}[T] + \Big(E[T]\Big)^2$ to get the result. ■

IV. Conclusions

In this paper we have derived basic stability results for a network of n AIMD flows that compete for shared bandwidth via a bottleneck router employing drop-tail queuing.

Acknowledgements

This work was supported by Science Foundation Ireland grant 04-IN3 -I460.

References

[1] R. Shorten, C. King, F. Wirth, and D. Leith, "Modelling tcp in droptail environments," *Automatica*, vol. 43, pp. 441–449, 2007.

[2] R. Shorten, F. Wirth, and D. Leith, "A positive systems model of TCP-like congestion control: Asymptotic analysis," *IEEE/ACM Transactions on Networking*, vol. 14, pp. 616–629, 2006.

[3] F. Wirth, R. Stanojevic, R. Shorten, and D. Leith, "Stochastic equilibria of AIMD communication networks,." SIAM Journal on Matrix Analysis and Applications, vol. 28(3), pp. 703-723, 2006.

[4] A. Leizarowitz, R. Stanojevic, and R. Shorten, "Tools for the analysis and design of communication networks with markovian dynamics,." IEE Proceedings on Control Theory - in press, 2006.

[5] O. Stenflo, "Uniqueness of invariant measures for place dependent iterations of functions," *IMA Journal of Applied Mathematics*, vol. 132, pp. 13–32, 2002.

[6] R. Isaac, "Markov processes and unique probability measures," *Pacific Journal of Mathematics*, vol. 12, pp. 273–286, 1962.

[7] M. Barnsley, S. Demko, J. Elton, and J. Geronimo, "Invariant measures for markov processes arising from iterated function systems with place dependent probabilities," *Ann. Inst. H. Poincaire Probab. Statist.*, vol. 24, pp. 367–394, 1988.

[8] J. Elton, "A multiplicative ergodic theorem for Lipschitz maps," *Stochastic Processes App.*, vol. 231(2), pp. 39–47, 1990.

Appendix : Proof of mathematical results

Proof of Theorem 3.2: Barnsley et al [7] have derived a general condition for ergodicity of Markov chains with place-dependent probabilities. The version we need (compact state space, Lipschitz continuous transition functions) was derived earlier by Isaac [6]. Elton [8] extended these results by proving almost sure convergence of time averages to ensemble averages under the same conditions. The conditions which guarantee ergodicity in our case are contained in the following lemma. These conditions state that the two-step transition probabilities satisfy an average contractivity property.

Lemma 4.1: Assume the drop probabilities satisfy the conditions 3.3-3.5. Then there is $r < 1$ such that for all $v, w \in \mathcal{S}$,

$$\sum_{i,j=1}^{m} p_i(v)\, p_j(A_i v) ||A_j A_i v - A_j A_i w||_1 \leq r||v-w||_1 \tag{22}$$

and there is $\delta > 0$ such that

$$\sum_{C} p_i(v)\, p_j(A_i v)\, p_i(w)\, p_j(A_i w) \quad > \quad \delta \tag{23}$$

where C denotes $i, j : ||A_j A_i v - A_j A_i w||_1 \leq r||v-w||_1$.

Proof: Since each matrix A_i is column stochastic, it is a contraction with respect to the l_1 norm, that is for any vector x we have

$$||A_i x||_1 \leq ||x||_1 \tag{24}$$

Also, let A_{all} be the matrix corresponding to the case where all flows are dropped. Then A_{all} is a strict contraction, so there is $r' < 1$ such that

$$||A_{all} x||_1 \leq r'||x||_1 \tag{25}$$

Define

$$B_\epsilon = \{w \in \mathcal{S} | \min_i w_i < \epsilon\} \tag{26}$$

By our assumptions on the $p_i(w)$, we can find $q_0 > 0$ and $\epsilon >> 0$, such that for every $w \in B_\epsilon$ there is some index i such that $p_i(w) > q_0$ and $A_i w \notin B_\epsilon$. Furthermore, since by assumption the probability to select A_{all} is nonzero in the interior, this probability is bounded away from zero in B_ϵ^c, that is there is $q_1 > 0$ such that

$$P(A(k) = A_{all} | W(k) \notin B_\epsilon) \geq q_1 \tag{27}$$

To prove (22), suppose first that $v \notin B_\epsilon$), then we obtain the bound

$$\begin{aligned} \sum_{i,j=1}^{m} p_i(v)\, p_j(A_i v) ||A_j A_i v - A_j A_i w||_1 \\ \leq \sum_{i=1}^{m} p_i(v) ||A_i v - A_i w||_1 \\ \leq (1 - q_1 + r q_1)||v-w||_1 \end{aligned} \tag{28}$$

where the second inequality follows by separately considering the cases $A_i = A_{all}$ and $A_i \neq A_{all}$. Now consider the case where $v \in B_\epsilon$, then at the first step there is probability q_0 to select a map that leads at the next step to the set B_ϵ^c. Combining with the previous estimate gives the bound

$$\begin{aligned} \sum_{i,j=1}^{m} p_i(v)\, p_j(A_i v) ||A_j A_i v - A_j A_i w||_1 \\ \leq (1-q_0)||v-w||_1 + q_0(1 - q_1 + r q_1)||v-w||_1 \end{aligned} \tag{29}$$

and together these two cases establish (22). The bound (23) follows by again separating the cases where $v \in B_\epsilon$ and $v \notin B_\epsilon$, and making the choices described above.

We now apply Theorem 2.1 of [7], which states that under the conditions of Lemma 4.1 there is an attractive, unique invariant probability measure for the Markov chain (note that Barnsley et al assume average contractivity at each step of the process, however average contractivity over two steps as in Lemma 4.1 is also sufficient to derive the result). In

Authorized licensed use limited to: The Library NUI Maynooth. Downloaded on December 21, 2009 at 07:21 from IEEE Xplore. Restrictions apply.

fact as noted before the proof in our case is simpler, as we have a compact state space and we assume uniform Lipschitz regularity for the drop probability functions. This establishes (i) and (ii). Property (iii) follows by Elton's result [8], which proves that time averages converge almost surely to ensemble averages for this model. ■

Authorized licensed use limited to: The Library NUI Maynooth. Downloaded on December 21, 2009 at 07:21 from IEEE Xplore. Restrictions apply.