SIAM J. Matrix Anal. Appl.
Vol. 28, No. 3, pp. 703–723

STOCHASTIC EQUILIBRIA OF AIMD COMMUNICATION NETWORKS*

F. WIRTH†, R. STANOJEVIĆ†, R. SHORTEN†, and D. LEITH†

Abstract. In this paper tools are developed to analyse a recently proposed random matrix model of communication networks that employ additive-increase multiplicative-decrease (AIMD) congestion control algorithms. We investigate properties of the Markov process describing the evolution of the window sizes of network users. Using paracontractivity properties of the matrices involved in the model, it is shown that the process has a unique invariant probability, and the support of this probability is characterized. Based on these results we obtain a weak law of large numbers for the average distribution of resources between the users of a network. This shows that under reasonable assumptions such networks have a well-defined stochastic equilibrium. ns2 simulation results are discussed to validate the obtained formulae. (The simulation program *ns*2, or *network simulator*, is an industry standard for the simulation of Internet dynamics.)

Key words. positive matrices, infinite products of positive matrices, AIMD congestion control, communication networks, Markov e-chain, law of large numbers

AMS subject classifications. 15A60, 68M12, 15A52

1. Introduction.

The dynamics of communication networks have attracted increased attention in recent years. Networks of devices that employ *additive-increase multiplicative-decrease* (AIMD) congestion control algorithms, such as the widely deployed *transmission control protocol* (TCP), have become the focus of much of this activity. Typically, the approach adopted by the community is to model such networks by means of a fluid analogy and to employ techniques from control theory and convex optimization in their analysis; see the recent book by Srikant [27] and the references therein for an overview of this work. Recently, several authors have proposed an alternative model of TCP dynamics using products of random matrices [2, 3, 24]. The basic approach followed in these papers is to use ideas from hybrid systems theory to model the dynamics of AIMD networks as a switched, or time-varying, discrete time linear system. The approach adopted in [24] allows for techniques from the theory of nonnegative matrices and Markov chains to be employed in the analysis of these networks. The application of these techniques to the study of such networks and the mathematical analysis of the model are the principal contributions of this paper.

Networks of unsynchronized sources and drop-tail queues have been the subject of several other studies [1, 3, 5, 12, 16], and it has been documented by many authors that networks of many AIMD flows exhibit extremely complex behavior. Consequently, it is convenient to analyze such networks from a probabilistic viewpoint, as we shall do in section 4. The novelty of our approach lies in the fact that we use positive matrices to model network behavior. We shall see that this will enable us to use results from

*Received by the editors December 16, 2004; accepted for publication (in revised form) by R. Nabben March 22, 2006; published electronically September 19, 2006. This work was supported by the European Union–funded research training network *Multi-Agent Control*, HPRN-CT-1999-00107, by the Enterprise Ireland grant SC/2000/084/Y, and by the Collaborative Research Center 637 "Autonomous Logistic Processes—A Paradigm Shift and its Limitations," funded by the German Research Foundation.

http://www.siam.org/journals/simax/28-3/62095.html

†NUI Maynooth, Hamilton Institute, Maynooth, Co. Kildare, Ireland (fabian.wirth@nuim.ie, rade.stanojevic@nuim.ie, robert.shorten@nuim.ie, doug.leith@nuim.ie).

the theory of positive matrices to be employed to make predictions concerning the behavior of AIMD networks.

Fluid analogy approaches to the modeling of networks of unsynchronized sources have been the subject of wide study in the TCP community; see [6, 13, 14, 18, 19, 20, 21, 22, 15, 28, 17] and the accompanying references for further details. However, several authors have recently developed hybrid system models of networks with a single bottleneck link which employ AIMD congestion control mechanisms, most notably Hespanha [11] and Baccelli and Hong [2]. We note that the model derived in [2] is similar to the model presented here. However, whereas the model derived by Baccelli and Hong is also a random matrix model, it has an affine structure. The corresponding homogeneous (linear) part is characterized by matrices without any nonnegativity structure. In [25, 24] the same model as the one presented here is discussed. The paper [24] deals with the derivation of expected average throughputs and with the question of model validation. In [25] implications of the model for network responsiveness and network fairness are discussed, and the model validation is carried one step further in that the effects of background traffic are analyzed.

In section 2 we begin our discussion by giving an overview of AIMD congestion control and by briefly reviewing the random matrix model of AIMD network dynamic first derived in [24]. In section 3 a number of basic results are presented relating to the set of matrices used in the model. It is shown that on a jointly invariant subspace the matrices are paracontractive, which is used to show that with probability one, left products of the matrices approach the set of rank-1 column stochastic matrices. This ergodicity property plays a vital role in all the subsequent considerations. Section 4 is devoted to the analysis of the Markov chain model of the AIMD process. It is shown that the chain in question is an e-chain. Using the results of section 3 we obtain that this chain has positive and aperiodic states. From this we obtain the unique existence of an invariant probability and weak law of large number statements. Finally, the support of the invariant probability is characterized. In section 5 we collect and derive a number of results that are useful in characterizing the stochastic equilibria of various types of communication networks that employ AIMD congestion control mechanisms. In section 6 we apply these results to the study of networks employing TCP congestion control. It is shown that the model is able to predict the average behavior of TCP flows very accurately.

2. Column stochastic matrices and AIMD congestion control. A communication network consists of a number of sources and sinks connected together via links and routers. In this paper we assume that these links can be modeled as a constant propagation delay together with a queue, that the queue is operating according to a drop-tail discipline, and that all of the sources are operating an AIMD-like congestion control algorithm. In AIMD congestion control each source maintains an internal variable w_i (the window size) which tracks the number of sent unacknowledged packets that can be in transit at any time. When the window size is exhausted, the source must wait for an acknowledgment before sending a new packet. Congestion control is achieved by dynamically adapting the window size according to an additive-increase multiplicative-decrease law. Roughly speaking, the source gently probes the network for spare capacity by increasing the rate at which packets are inserted into the network, and backs off rapidly the number of packets transmitted through the network when congestion is detected through the loss of data packets. More specifically, an individual source sends packets of data through the network to a destination, and the transmission is deemed complete if an acknowledgment issued

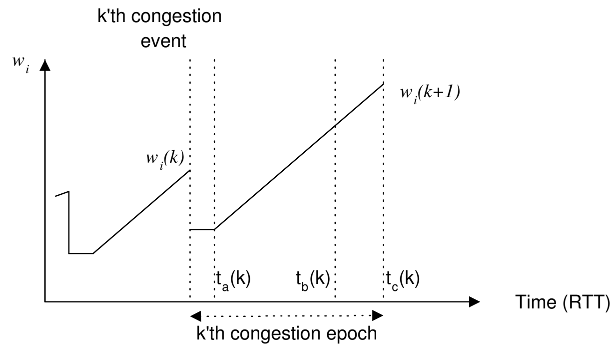

FIG. 2.1. *Evolution of window size.*

by the destination upon receipt of the packet is received by the source. As long as transmission is successful, that is, as long as all acknowledgments are received, the source increments $w_i(t)$ by a fixed amount α_i upon receipt of an acknowledgment. If an acknowledgment for a certain packet does not arrive at the sender, it is assumed that there has been a packet loss due to congestion in the network. As a consequence, the variable $w_i(t)$ is reduced in multiplicative fashion to $\beta_i w_i(t)$, where $0 < \beta_i < 1$.

2.1. A model for AIMD dynamics. In [26] a model has been presented which assumes that (i) at congestion every source experiences a packet drop; and (ii) each source has the same round-trip time (RTT).[1] In [24] this model has been extended to a random model of unsynchronized networks, where sources have different RTTs. We briefly describe the derivation of the model here. A standing assumption of the model is that all sources compete for the capacity of a single bottleneck router, and if packets are lost, this happens because the queue of that router is overflowing.

By a *congestion event* we describe the situation that more packets arrive at a router than can be serviced and the queue of the router is already full. In this case, necessarily some packets are lost. Without the assumption of synchronization, at a congestion event not all sources are necessarily informed of this congestion. For the moment uniform RTT is still assumed; we will weaken this assumption later on. Let $w_i(k)$ denote the congestion window size of source i immediately before the kth network congestion event is detected by the source.

Over the kth congestion epoch as depicted in Figure 2.1 three important events can be discerned: $t_a(k)$, $t_b(k)$, and $t_c(k)$. The time $t_a(k)$ denotes the instant at which the number of unacknowledged packets in flight equals $\beta_i w_i(k)$; $t_b(k)$ is the time at which the bottleneck queue is full; and $t_c(k)$ is the time at which packet drop is detected by some of the sources, where time is measured in units of RTT.[2] It follows from the definition of the AIMD algorithm that the window evolution is completely defined over all time instants by knowledge of the $w_i(k)$ and the event times $t_a(k)$, $t_b(k)$, and $t_c(k)$ of each congestion epoch. We therefore only need to investigate the behavior of these quantities.

[1] One RTT is the time between sending a packet and receiving the corresponding acknowledgment when there are no packet drops.

[2] Note that measuring time in units of RTT results in a linear rate of increase for each of the congestion window variables between congestion events.

We assume that sources that lose a package at congestion are informed of this loss one RTT after the queue at the bottleneck link becomes full; that is, $t_c(k) - t_b(k) = 1$. Also,

$$w_i(k) \geq 0 \quad \text{and} \quad \sum_{i=1}^{n} w_i(k) = P + \sum_{i=1}^{n} \alpha_i \ \forall k > 0, \tag{2.1}$$

where P is the maximum number of packets which can be in transit in the network at any time; P is usually equal to $q_{max} + BT_d$, where q_{max} is the maximum queue length of the congested link, B is the service rate of the congested link in packets per second, and T_d is the RTT when the queue is empty. At the $(k+1)$th congestion event

$$w_i(k+1) = \begin{cases} \beta_i^s w_i(k) + \alpha_i[t_c(k) - t_a(k)] & \text{if source } i \text{ experiences congestion,} \\ w_i(k) + \alpha_i[t_c(k) - t_a(k)] & \text{else,} \end{cases} \tag{2.2}$$

and we set

$$\beta_i(k) \in \{\beta_i^s, 1\}, \tag{2.3}$$

corresponding to whether the source experiences a packet loss or not. Then summing the equations in (2.2) and using (2.1) we obtain

$$t_c(k) - t_a(k) = \frac{1}{\sum_{i=1}^{n} \alpha_i} \left[P - \sum_{i=1}^{n} \beta_i(k) w_i(k) \right] + 1, \tag{2.4}$$

and using (2.2)–(2.4), it follows that

$$w_i(k+1) = \beta_i(k) w_i(k) + \frac{\alpha_i}{\sum_{j=1}^{n} \alpha_j} \left[\sum_{j=1}^{n} (1 - \beta_j(k)) w_j(k) \right]. \tag{2.5}$$

Thus the dynamics of an entire network of such sources is given by

$$w(k+1) = A(k) w(k), \tag{2.6}$$

where $w^T(k) = [w_1(k), \ldots, w_n(k)]$, and, writing $D(\beta(k)) = \mathrm{diag}(\beta_1(k), \ldots, \beta_n(k))$,

$$A(k) = D(\beta(k)) + \frac{1}{\sum_{j=1}^{n} \alpha_j} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \cdots \\ \alpha_n \end{bmatrix} \begin{bmatrix} 1 - \beta_1(k) & \cdots & 1 - \beta_n(k) \end{bmatrix}. \tag{2.7}$$

As the entries of $w(k)$ are nonnegative for all $k \geq 0$ the equations (2.6) define a positive linear system [4]. Using $b_i(s) \in (0, 1], i = 1, \ldots, n$, we also see that all possible matrices that appear are column stochastic. In what follows we will call column stochastic matrices of the form (2.7) *AIMD matrices.*

So far we have worked with the assumption of uniform RTT, which is quite restrictive (although it may, for example, be valid in some long-distance networks [29]). We now extend our approach to more general network conditions. As we will see, the model that we obtain shares many structural and qualitative properties of the model described above. To distinguish variables, the nominal parameters of the

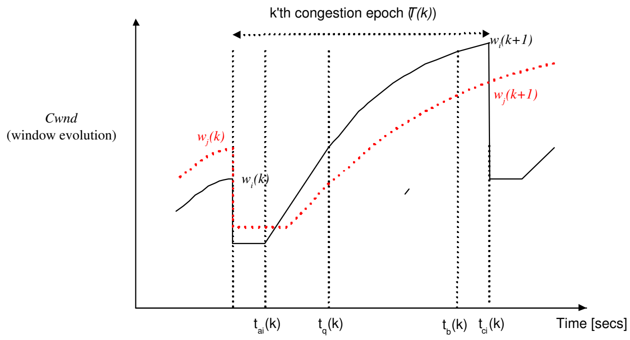

FIG. 2.2. *Evolution of window size over a congestion epoch. $T(k)$ is the length of the congestion epoch in seconds.*

sources used in the previous section are now denoted by α_i^s, β_i^s, $i = 1, \ldots, n$. Here the index s may remind the reader that these are the parameters that are chosen by each *source*.

Consider the general case of a number of sources competing for shared bandwidth in a generic dumbbell topology (where sources may have different RTTs and drops need not be synchronized). The evolution of the window size w_i of a typical source as a function of time, over the kth congestion epoch, is depicted in Figure 2.2. As before a number of important events may be discerned, where we now measure time in seconds, rather than units of RTT. Denote by $t_{ai}(k)$ the time at which the number of packets in flight belonging to source i is equal to $\beta_i^s w_i(k)$; $t_q(k)$ is the time at which the bottleneck queue begins to fill; $t_b(k)$ is the time at which the bottleneck queue is full; and $t_{ci}(k)$ is the time at which the ith source is informed of congestion. In this case the evolution of the ith congestion window variable does not evolve linearly with time after t_q seconds due to the effect of the bottleneck queue filling and the resulting variation in RTT; namely, the RTT of the ith source increases according to $RTT_i(t) = T_{d_i} + q(t)/B$ after t_q, where T_{d_i} is the RTT of source i when the bottleneck queue is empty and $0 \leq q(t) \leq q_{\max}$ denotes the number of packets in the queue. Note also that we do not assume that every source experiences a drop when congestion occurs. For example, a situation is depicted in Figure 2.2 where the ith source experiences congestion at the end of the epoch, whereas the jth source does not.

Given these general features it is clear that the modeling task is more involved than in the synchronized case. Nonetheless, it is possible to relate $w_i(k)$ and $w_i(k+1)$ using a similar approach to the synchronized case by accounting for the effect of nonuniform RTTs and unsynchronized packet drops as follows.

Due to the variation in RTT, the congestion window of a flow does not evolve linearly with time over a congestion epoch. Nevertheless, we may relate $w_i(k)$ and $w_i(k+1)$ linearly by defining an average rate $\alpha_i(k)$ depending on the kth congestion epoch:

$$\alpha_i(k) := \frac{w_i(k+1) - \beta_i(k)w(k)}{T(k)}, \tag{2.8}$$

where $T(k)$ is the duration of the kth epoch measured in seconds. Equivalently we have

$$w_i(k+1) = \beta_i(k)w_i(k) + \alpha_i(k)T(k)\,. \tag{2.9}$$

In the case when $q_{max} \ll BT_{d_i}$, $i = 1, \dots, n$, the average α_i are (almost) independent of k and given by $\alpha_i(k) \approx \alpha_i^s / T_{d_i}$ for all $k \in \mathbb{N}, i = 1, \dots, n$. The situation when

$$\alpha_i \approx \frac{\alpha_i^s}{T_{d_i}}, \quad i = 1, \dots, n, \tag{2.10}$$

is of considerable practical importance and such networks are the principal concern of this paper. See [24] for a discussion of networks where this assumption is reasonable.

In view of (2.3) and (2.9) a convenient representation of the network dynamics is obtained as follows. At congestion the bottleneck link is operating at its capacity B, i.e.,

$$\sum_{i=1}^{n} \frac{w_i(k) - \alpha_i}{RTT_{i,max}} = B, \tag{2.11}$$

where $RTT_{i,max}$ is the RTT experienced by the ith flow when the bottleneck queue is full. Note that $RTT_{i,max}$ is independent of k. Setting $\gamma_i := (RTT_{i,max})^{-1}$ we have that

$$\sum_{i=1}^{n} \gamma_i w_i(k) = B + \sum_{i=1}^{n} \gamma_i \alpha_i \,. \tag{2.12}$$

Using steps similar to the ones performed in (2.2)–(2.4) we obtain the model

$$w_i(k+1) = \beta_i(k) w_i(k) + \frac{\alpha_i}{\sum_{j=1}^{n} \gamma_j \alpha_j} \left(\sum_{j=1}^{n} \gamma_j (1 - \beta_j(k)) w_j(k) \right), \tag{2.13}$$

and the dynamics of the entire network of sources at the kth congestion event are again described by $w(k+1) = A(k) w(k)$, where

$$A(k) = D(\beta(k)) + \frac{1}{\sum_{j=1}^{n} \gamma_j \alpha_j} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \cdots \\ \alpha_n \end{bmatrix} [\gamma_1(1 - \beta_1(k)), \dots, \gamma_n(1 - \beta_n(k))], \tag{2.14}$$

and where $\beta_i(k)$ is either 1 or β_i^s. The nonnegative matrices $A_2, \dots, A_m$ are constructed by taking the matrix A_1,

$$A_1 = \begin{bmatrix} \beta_1^s & 0 & \cdots & 0 \\ 0 & \beta_2^s & 0 & 0 \\ \vdots & 0 & \ddots & 0 \\ 0 & 0 & \cdots & \beta_n^s \end{bmatrix} + \frac{1}{\sum_{j=1}^{n} \gamma_j \alpha_j} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \cdots \\ \alpha_n \end{bmatrix} \begin{bmatrix} \gamma_1(1 - \beta_1^s), & \dots, & \gamma_n(1 - \beta_n^s) \end{bmatrix},$$

and setting some, but not all, of the β_i to 1. This gives rise to $m = 2^n - 1$ matrices associated with the system (2.13) that correspond to the different combinations of source drops that are possible. These matrices are not AIMD matrices in the sense we have defined above. However, by a small transformation we come back to our original situation.

By considering the evolution of $w_\gamma^T(k) = [\gamma_1 w_1(k), \gamma_2 w_2(k), \dots, \gamma_n w_n(k)]$ we obtain the following description of the network dynamics:

$$w_\gamma(k+1) = \bar{A}(k) w_\gamma(k) \tag{2.15}$$

with $\bar{A}(k) \in \bar{\mathcal{A}} = \{\bar{A}_1, \dots, \bar{A}_m\}$, $m = 2^n - 1$, and where the $\bar{A}_i$ are obtained by the diagonal similarity transformation associated with the change of variables. As before the nonnegative matrices $\bar{A}_2, \dots, \bar{A}_m$ are constructed by taking the matrix $\bar{A}_1$ and setting some, but not all, of the β_i^s to 1. It is easy to see that all of the matrices in the set $\bar{\mathcal{A}}$ are now AIMD matrices; for convenience we use this representation of the network dynamics to prove the main mathematical results presented in this paper. Note furthermore that the similarity transformation used to bring the matrices in AIMD form depends only on the round-trip times RTT_i and not on the α_i^s, β_i^s.

2.2. Networks of flows whose parameters vary in time. Before proceeding with our analysis we note that for some applications it is convenient to allow the parameters of the matrix $A(k)$ to vary in more general a manner than that described in the previous two sections. Our model may be extended trivially to model networks whose AIMD parameters vary with time: $\alpha_i(k)$; $\beta_i(k)$. Such situations may arise in applications where the protocol adapts its parameters to reflect prevailing network conditions or in applications where variations in network delays lead to a consequent variation in the AIMD parameters (for example, due to routing changes or in wireless networks) [22]; in fact a number of AIMD networks of this type have recently been proposed by a number of authors in the context of high-speed long-distance networks [29]. We account for such behavior in this paper by defining the set $\mathcal{M}$ to be the union of a finite number of matrix sets $\bar{\mathcal{A}}_j$, each of which is defined as above but which corresponds to fixed AIMD parameters $\{\alpha_1^j, \dots, \alpha_n^j\}$ and $\{\beta_1^j, \dots, \beta_n^j\}$, $1 \le j \le h$, with $\mathcal{M} = \bigcup_{j=1}^h \bar{\mathcal{A}}_j$, where h is some fixed integer.

3. Preliminaries. The principal objective of this paper is to collect and develop analytic tools to analyze models of the form derived in section 2. We will see in section 5 that it is possible to characterize the stochastic behavior of the random variable $w(k)$ under certain assumptions. The derivation of these results is somewhat technical, and to ease exposition we introduce here a number of definitions and preliminary results.

3.1. Basic notation. The following results are based on the theory of nonnegative matrices. A matrix A or a vector x is said to be nonnegative if each of its entries is a nonnegative real number and matrices or vectors are called positive if all their entries are positive. We write $A \succ B$ or $A \succeq B$ if $A - B$ is positive, respectively, nonnegative. The set of nonnegative matrices in $\mathbb{R}^{n\times n}$ is denoted by $\mathbb{R}_+^{n\times n}$. The componentwise absolute value of $A = (a_{ij}) \in \mathbb{R}^{n\times m}$ is defined by $|A| := (|a_{ij}|) \in \mathbb{R}_+^{n\times m}$.

A special subset of $\mathbb{R}_+^{n\times n}$ are the column stochastic matrices. A matrix $A \in \mathbb{R}_+^{n\times n}$ is called column stochastic if for each of its columns the sum of the corresponding elements is equal to 1. Denoting $e := [1, 1, \dots, 1]^T$, it follows that e^T is a left eigenvector of a column stochastic matrix corresponding to the eigenvalue 1. We denote by $\mathcal{R} \subset \mathbb{R}^{n\times n}$ the set of all column stochastic matrices of rank-1 and the distance between a matrix $P \in \mathbb{R}^{n\times n}$ and the set $\mathcal{R}$ by $\operatorname{dist}(P, \mathcal{R}) = \inf\{\|P - C\| : C \in \mathcal{R}\}$, where $\|\cdot\|$ is the induced l_1-norm. Finally, the standard jth unit vector is denoted by e_j, so that $e = \sum_{j=1}^n e_j$.

3.2. Basic assumptions. Our basic objective is to model the evolution of the vector $w(k)$ for networks of AIMD flows. We consider a set of AIMD matrices $\mathcal{M} = \{M_1, \dots, M_\mu\}, \mu \ge 1$. Associated to this set we consider the deterministic system

$$x(k+1) \in \{Mx(k) \mid M \in \mathcal{M}\} \tag{3.1}$$

and a Markov chain model

$$w(k+1) = A(k)w(k)\,, \tag{3.2}$$

where for each k the $A(k)$ is a random variable with values in $\mathcal{M}$. We recall that by (2.1) the sum $\sum_i w_i(k)$ is a constant. We may thus restrict our attention to the simplex

$$\Sigma := \left\{ x \in \mathbb{R}^n_+ \mid e^T x = \sum_{i=1}^{n} x_i = 1 \right\} ,$$

and we will study the evolution of (3.2) on Σ. We assume that the random variables $A(k), k = 0, 1, \ldots,$ are independent and identically distributed (i.i.d.) and denote

$$P(A(k) = M_i) = \rho_i\,, \quad i = 1, \ldots, \mu\,.$$

As we are dealing with probabilities, necessarily, we assume $\sum_i \rho_i = 1$. With this setup the sequence $\{w(k)\}_{k\in\mathbb{N}}$ is a Markov process. The random variable of a product of length k is denoted by $\Pi(k) = A(k)A(k-1)\ldots A(0)$.

Clearly, $w(k) = \Pi(k)w(0)$, and consequently the behavior of $w(k)$, as well as the network fairness and convergence properties, are governed by the asymptotic properties of the matrix product $\Pi(k)$ as $k \to \infty$.

ASSUMPTION 3.1. *Let $\mathcal{M} = \{M_1, \ldots, M_\mu\}$ be a set of matrices of the form* (2.7)*. We assume that the probability that $A(k) = M_i$ in* (3.2) *is independent of k and equals $\rho_i > 0$.*

COMMENT 3.2. *In other words Assumption* 3.1 *says that the probability that the network dynamics are described by $w(k+1) = A(k)w(k), A(k) = M_i$ over the kth congestion epoch is ρ_i and that the random variables $A(k), k \in \mathbb{N}$ are i.i.d. Furthermore, we assume that we only have matrices in the set $\mathcal{M}$ which occur with positive probability. Without this assumption there is little insight to be gained into the dynamics of the Markov chain* (3.2) *by studying the deterministic system* (3.1)*. This assumption implies no loss of generality because we may simply remove matrices with* 0 *probability from the set $\mathcal{M}$.*

Given the probabilities ρ_i for $M_i \in \mathcal{M}$, one may then define the probability λ_j that source j experiences a backoff at the kth 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 jth source sees a drop. To put it another way, the summation is over those indices i for which the matrix M_i is defined with a value of $\beta_j \neq 1$.

ASSUMPTION 3.3. *Let $\mathcal{M} = \{M_1, \ldots, M_\mu\}$ be the set of AIMD matrices defining* (3.2) *and assume that $P(A(k) = M_i) = \rho_i, i = 1, \ldots, \mu$. We assume that $\lambda_j > 0$ for all $j \in \{1, \ldots, n\}$.*

Simply stated, by Assumption 3.3 all flows must see a drop almost surely at some time (provided that they live for a long enough time).

3.3. Column stochastic matrices. Column stochastic matrices will play a central role in the discussion in section 5. We begin by collecting some results. The following two are immediate consequences of the definition of a column stochastic matrix.

LEMMA 3.4. *A matrix $A \in \mathbb{R}_+^{n\times n}$ is column stochastic if and only if $e^T A = e^T$. Any product of a finite number of column stochastic matrices is a column stochastic matrix (i.e., the set of column stochastic matrices is a semigroup).*

It is sometimes convenient to consider the subspace orthogonal to e, which we denote by

$$S := \{z \in \mathbb{R}^n \mid e^T z = 0\}.$$

The subspace S is an invariant subspace for all column stochastic matrices. Given a column stochastic matrix A we denote by $\tilde{A} : S \to S$ the linear operator obtained by restricting A to S. Furthermore, we denote by $\|\cdot\|$ the 1-norm and the corresponding induced matrix norm.

LEMMA 3.5. *For any column stochastic matrix A it holds that $\|A\| = 1$ and $\|\tilde{A}\| \le 1$. If A is positive, then $\|\tilde{A}\| < 1$.*

Proof. The first claim is immediate from the standard characterization of the induced 1-norm as the column-sum norm. The second claim follows as $\|\tilde{A}\| \le \|A\|$ using the definition of induced norms. Finally, if A is positive, then for a vector $z \in S, \|z\| = 1$ it holds that $-A|z| \prec |Az| \prec A|z|$ as z has positive and negative entries due to $e^T z = 0$. This implies for $z \in S, \|z\| = 1$ that

$$\|\tilde{A}z\| = \|Az\| = \|\,|Az|\,\| < \|A|z|\,\| = 1.$$

This shows the assertion. ∎

A feature in the proof of our main results is the observation that products of our AIMD matrices converge to a certain compact subset of the rank-1 idempotent matrices (in the sense that the distance to this set goes to zero). We use the following lemma to estimate the distance of a matrix product from the set $\mathcal{R}$ defined at the beginning of this section.

LEMMA 3.6. *Let $A \in \mathbb{R}_+^{n\times n}$ be column stochastic; then* $\mathrm{dist}\,(A, \mathcal{R}) \le 2\|\tilde{A}\|$.

Proof. Let $A_1 = A - Aee^T/n$. Note that Aee^T/n is a rank-1 column stochastic matrix. Then $\mathrm{dist}\,(A, R) = \inf\{\|A - C\| : C \in \mathcal{R}\} \le \|A - Aee^T/n\| = \|A_1\|$. We are proving that $\|A_1\| \le 2\|\tilde{A}\|$. So let $x = z + te$, where $z \in S, t \in \mathbb{R}$ are arbitrary. Then

$$A_1 x = (A - Aee^T/n)(z + te) = Az = \tilde{A}z,$$

so

$$\|A_1 x\| \le \|\tilde{A}z\| \le \|\tilde{A}\|\|z\|.$$

To complete the proof we show that $\|z\| \le 2\|z + te\|$. Indeed, if $z_1, z_2, \dots, z_n$ are the components of z ordered such that $z_1 \ge z_2 \ge \dots \ge z_r \ge 0 > z_{r+1} \ge \dots \ge z_n$, then $\|z\| = |z_1| + |z_2| + \dots + |z_n| = 2(|z_1| + |z_2| + \dots + |z_r|)$. On the other hand for $t \ge 0$,

$$\|z + te\| = \sum_{j=1}^{n} |z_j + t| \ge \sum_{j=1}^{r} |z_j + t| \ge \sum_{j=1}^{r} |z_j| = \frac{1}{2}\|z\|,$$

thus $\|z\| \le 2\|z + te\|$. For $t < 0$ a similar argument applies. ∎

Recall that the similarity transformation described to obtain (2.15) is applied simultaneously to the matrices from (2.14). Thus each matrix $M \in \mathcal{M}$ can be written in the form

$$\mathrm{diag}(\beta_1, \beta_2, \dots, \beta_n) + \frac{1}{\sum_{j=1}^{n} \alpha_j \gamma_j}[\alpha_1\gamma_1, \dots, \alpha_n\gamma_n]^T[(1 - \beta_1), \dots, (1 - \beta_n)], \tag{3.3}$$

where $\alpha_j(M)$ are positive, and all $\beta_j(M)$ are positive and not greater than 1. The parameters γ_j are also positive and independent of $M \in \mathcal{M}$, as they are determined by the RTTs of the sources; see (2.15). Thus the matrices in $\mathcal{M}$ are column stochastic. Note that if the jth column of $M \in \mathcal{M}$ is not strictly positive, then that column is equal to e_j. Using the assumptions given in section 3.2, we now aim to prove certain convergence results for the restriction of $A(k)A(k-1)\cdots A(1)$ to S. To this end we employ the notion of paracontractivity [7, 10] from the theory of nonhomogeneous matrix products. A linear operator A on $\mathbb{R}^n$ is called *paracontractive* with respect to the norm $\|\cdot\|$ if

$$Ax \neq x \Rightarrow \|Ax\| < \|x\|. \tag{3.4}$$

We will employ the following three results to show that almost surely products of matrices from $\mathcal{M}$ converge to the set $\mathcal{R}$. The following result is proved in [7].

THEOREM 3.7. *Let $\|\cdot\|$ be a norm on $\mathbb{R}^n$ and let $\mathcal{F} \subset \mathbb{R}^{n\times n}$ be a finite set of linear operators which are paracontractive with respect to $\|\cdot\|$. Then for any sequence $\{A_k\}_{k\in\mathbb{N}} \subset \mathcal{F}^{\mathbb{N}}$, the sequence of left products $\{A_kA_{k-1}\dots A_1\}_{k\in\mathbb{N}}$ converges.*

The second result shows that all matrices from $\mathcal{M}$ are paracontractive with respect to the 1-norm on S.

LEMMA 3.8. *Let $A \in \mathcal{M}$. Then $\tilde{A}$ is paracontractive on S with respect to the 1-norm.*

Proof. As before, let $\|\cdot\|$ denote the 1-norm. For $x \in S$ we want to show (3.4). We know that any matrix from $\mathcal{M}$ can be written in the form (3.3), where $\beta_i \in (0,1], i = 1,\dots,n$, and $\beta_j < 1$ for some $j \in \{1,\dots,n\}$. Also $\alpha_i > 0$ and $\gamma_i > 0$ for $i = 1,2,\dots,n$. Without loss of generality, assume that $\beta_1 = \beta_2 = \cdots = \beta_q = 1$ for $q < n$ and $\beta_i < 1, i = q+1,\dots,n$. In this case our matrix A is of the form

$$A = \begin{bmatrix} I_q & A_{12} \\ 0 & A_{22} \end{bmatrix},$$

where I_q is the identity matrix of order q and where $A_{12}, A_{22} \succ 0$ are such that the elements of each column of A sums to 1. Pick $x \in S$. If we partition $x = \begin{bmatrix} z_1^T & z_2^T \end{bmatrix}^T$ accordingly, we have

$$Ax = \begin{bmatrix} z_1 + A_{21}z_2 \\ A_{22}z_2 \end{bmatrix}.$$

By Lemma 3.5 it follows that $\|Ax\| \leq \|x\|$. If $\|Ax\| = \|x\|$, then in each entry of Ax the summands have the same sign, because otherwise $\|Ax\| < \|A|x|\| \leq \||x|\| = \|x\|$, a contradiction. For $1 \leq j \leq q$, this implies that for $(Ax)_j = x_j + a_{jq+1}x_{q+1} + \cdots + a_{jn}x_n$ the signs of the summands coincide. Similarly for $q+1 \leq j \leq n$ the signs of the summands of $(Ax)_j = a_{jq+1}x_{q+1} + \cdots + a_{jn}x_n$ coincide. This implies $x_ix_j \geq 0$ for all $i = 1,2,\dots,n$ and all $j = q+1, q+2,\dots,n$. If we fix $j \geq q+1$, we have

$$0 = x_je^Tx = x_j(x_1 + \cdots + x_n) \geq x_j^2. \tag{3.5}$$

We conclude that $z_2 = 0$, which also means that $Ax = x$. Thus for $x \in S$, we have $\|Ax\| \leq \|x\|$ with equality if and only if $Ax = x$, as desired. ∎

Our third result is purely technical and is stated as a separate lemma to aid exposition of Theorem 3.10.

COROLLARY 3.9. *If $A \in \mathcal{M}$ is such that its not strictly positive columns are indexed by $i_1, i_2, \dots, i_q$ and $x \in S$ is such that $Ax = x$, then x lies in the subspace spanned by the vectors $e_{i_1}, e_{i_2}, \dots, e_{i_q}$.*

Proof. This follows from the previous proof, as we have seen that $Ax = x$ implies that $x_j = 0$ for $j = q+1, \dots, n$. In other words, $x \in \operatorname{span}\{e_1, \dots, e_q\}$. The general statement follows by permutation. □

Given the three previous results it is now possible to show that almost all products of matrices from $\mathcal{M}$ approach the set $\mathcal{R}$.

THEOREM 3.10. *Let $\{A_k\}_{k\in\mathbb{N}}$ be a sequence of matrices from $\mathcal{M}$. Assume that for all $i \in \{1, 2, \dots, n\}$ there is a matrix $T_i \in \mathcal{M}$ with positive ith column which occurs infinitely often in $\{A_k\}_{k\in\mathbb{N}}$. Then*

$$\lim_{k\to\infty} \{\tilde{A}_k \tilde{A}_{k-1} \cdots \tilde{A}_1\} = 0 .$$

In particular under Assumption 3.3, we have for the stochastic process $\{A(k)\}_{k\in\mathbb{N}}$ that $\lim_{k\to\infty} \tilde{A}(k)\tilde{A}(k-1)\cdots\tilde{A}(0) = 0$ almost surely.

Proof. By Lemma 3.8, the matrices $\tilde{A}_k$, $k \in \mathbb{N}$, are paracontractive with respect to $\|\cdot\|$. Using Theorem 3.7 it follows that $\{\tilde{A}_k\tilde{A}_{k-1}\cdots\tilde{A}_1\}_{k\in\mathbb{N}}$ is convergent. To prove that the limit is 0 let $s \in S$. Then there exist $y \in S$ such that $y = \lim_{k\to\infty} A_k A_{k-1}\cdots A_1 s$. We will prove that $y = 0$ from which the first assertion follows. For fixed i let $\{A_{n_k}\}_{k\in\mathbb{N}}$ be a subsequence of $\{A_k\}_{k\in\mathbb{N}}$ with $A_{n_k} = T_i$. Then

$$y = \lim_{k\to\infty} A_{n_k} A_{n_k - 1} \cdots A_1 s = T_i \lim_{k\to\infty} A_{n_k - 1} \cdots A_1 s = T_i y.$$

Thus $T_i y = y \in S$ since $s \in S$. By Corollary 3.9 the ith coordinate of y is zero. Since i is arbitrary, it follows that $y = 0$.

By Assumption 3.3 for each $j \in \{1, \dots, n\}$ the probability that matrices with positive jth column occur infinitely often in a realization of the process is equal to 1. Thus $\lim_{k\to\infty} \tilde{A}(k)\tilde{A}(k-1)\cdots\tilde{A}(0) = 0$ with probability 1. □

The next result shows that the expected distance between $A(k)A(k-1)\cdots A(1)$ and $\mathcal{R}$ decreases exponentially; a fact of independent interest.

PROPOSITION 3.11. *Let $\{A(k)\}_{k\in\mathbb{N}}$ be a sequence of random variables satisfying Assumptions 3.1 and 3.3. Let $d(k) := E(\operatorname{dist}(A(k)A(k-1)\cdots A(1), \mathcal{R}))$. Then there exist $\eta < 1$ and $C \geq 1$ such that for all k it holds that*

$$d(k) \leq C\eta^k. \tag{3.6}$$

Proof. Let $\theta = 1 - \min_{j=1,n} \lambda_j < 1$ and let l be an integer such that $1 > n\theta^l$.

At first, note that the jth column of the product of several matrices from $\mathcal{M}$ is positive if and only if one of these matrices has positive jth column, otherwise it is equal to e_j. Consider the products of length l: $\Pi(l) = A(l)A(l-1)\cdot A(1)$. The probability that the jth column of $\Pi(l)$ is not strictly positive is $o_j := (1-\lambda_j)^l \leq \theta^l$. For the probability q_l that at least one column of $\Pi(l)$ is not strictly positive, we have that $q_l \leq o_1 + o_2 + \cdots + o_n \leq n\theta^l$. Thus the probability p_l that $\Pi(l)$ is positive satisfies $p_l = 1 - q_l \geq 1 - n\theta^l > 0$. Let $k = dl + r$, where $0 \leq r < l$. We can split the product $\Pi(k) = A(k)A(k-1)\cdots A(1)$ into the product of the first r terms $D_0 = A(k)A(k-1)\cdots A(k-r+1)$ and the product of d blocks of length l: $D_i = A(il)A(il-1)\cdots A(l(i-1)+1)$ for $i = 1, 2, \dots, d$. So $\Pi(k) = D_0 D_d \cdots D_1$. Note that for all $i = 0, 1, \dots, d$, D_i, as a product of column stochastic matrices, is column stochastic, and therefore $\|D_i\| = 1$ and $\|\tilde{D}_i\| \leq 1$. With this notation we have

$$\operatorname{dist}(\Pi(k), \mathcal{R}) \leq 2\|\tilde{\Pi}(k)\| = 2\|\tilde{D}_0\tilde{D}_d\cdots\tilde{D}_1\| \leq 2\|\tilde{D}_d\cdots\tilde{D}_1\| .$$

Define

$$\delta := \max\{\|\tilde{T}\| : \ T = A_l A_{l-1} \dots A_1 > 0, A_1, A_2, \dots, A_l \in \mathcal{M}\} < 1 . \tag{3.7}$$

Since the set in (3.7) is finite, the maximum exists and is strictly less than 1 by Lemma 3.5. For any $j \in \{0, 1, 2, \dots, d\}$ the probability that exactly j of the matrices $D_1, D_2, \dots, D_d$ are positive is equal to $z_j = \binom{d}{j} p_l^j (1-p_l)^{d-j}$. We also know that if j of matrices $D_1, D_2, \dots, D_d$ are positive, then $\|(D_d D_{d-1} \cdots D_1)\,\tilde{}\,\| = \|\tilde{D}_d \cdots \tilde{D}_1\| \leq \|\tilde{D}_d\| \cdots \|\tilde{D}_1\| \leq \delta^j$. Thus we obtain

$$\begin{aligned} d(k) &\leq 2E(\|\tilde{D}_d\| \cdots \|\tilde{D}_1\|) \leq 2 \sum_{j=0}^{d} z_j \delta^j \\ &= 2 \sum_{j=0}^{d} \binom{d}{j} (p_l \delta)^j (1-p_l)^{d-j} = 2(1 + p_l\delta - p_l)^d \leq C\eta^k \,, \end{aligned}$$

where for the last inequality we choose

$$\eta := (1 - p_l + p_l \delta)^{1/l} < 1 \quad \text{and} \quad C := 2/\eta^l \,. \tag{3.8}$$

This shows the assertion. □

4. Invariant measures. In this section we study the existence of invariant measures of the Markov process $\{w(k)\}_{k\in\mathbb{N}}$. Throughout we assume that Assumptions 3.1 and 3.3 are satisfied. Our considerations are based on the results presented in [23], to which we refer the reader for further background material. We briefly present some basic properties for the Markov chain $\{w(k)\}_{k\in\mathbb{N}}$ on the simplex Σ. By $\mathcal{B}(\Sigma)$ we denote the Borel σ-algebra of Σ.

Associated with our Markov chain there is a transition kernel $P(x, X)$ for $x \in \Sigma, X \in \mathcal{B}(\Sigma)$, which gives the probability to reach the set X from the point x. This transition kernel acts on continuous functions $h : \Sigma \to \mathbb{R}$ through

$$Ph(x) = \int_\Sigma h(y)P(x, dy) = \sum_{i=1}^{\mu} \rho_i h(M_i x) \,. \tag{4.1}$$

It is obvious that Ph is continuous for continuous h, so that P is *(weak) Feller*. Furthermore we have $\|A_i\| \leq 1, i = 1, \dots, \mu$, so that $\|A_i(x - y)\| \leq \|x - y\|$. Using the uniform continuity of h it follows that for any continuous function $h : \Sigma \to \mathbb{R}$, the sequence

$$P^k h \,, \quad k \in \mathbb{N} \,,$$

defined inductively through repeated application of (4.1), is equicontinuous. Markov chains whose transition kernel have this property are called *e-chains*; see [23].

An important notion in the study of Markov chains are invariant probabilities. Recall that a probability measure π is called *invariant* for a Markov process if

$$\pi(X) = \int_\Sigma P(x, X) d\pi(x) \quad \forall X \in \mathcal{B}(\Sigma),$$

that is, intuitively, the distribution of mass on Σ given by the probability measure π is not changed if it is rearranged according to the evolution of the Markov process.

As we are considering an e-chain, we obtain from [23, Theorem 12.0.1] that an invariant probability exists in our case. We aim to show its uniqueness. To this end we first study the possible support of invariant measures. We introduce the set of sequences

$$\mathcal{L} := \{\{A_k\}_{k\in\mathbb{N}} \in \mathcal{M}^{\mathbb{N}} \mid \{A_k\}_{k\in\mathbb{N}} \text{ satisfies the conditions of Theorem 3.10}\}.$$

By Theorem 3.10 we know that the left products of a sequence $\{A_k\}_{k\in\mathbb{N}} \in \mathcal{L}$ approach the set of rank-1 column stochastic matrices. We define the set of limit points of such sequences by

$$\mathcal{R}_{\mathcal{L}} := \{R \in \mathcal{R} \mid \exists \{A_k\}_{k\in\mathbb{N}} \in \mathcal{L}, k_l \to \infty : \lim_{l\to\infty} \Pi(k_l) = R\}\,.$$

As the matrices $R \in \mathcal{R}$ are column stochastic and of rank 1 they can be represented in the form $R = ze^T$, where $z \succeq 0$ and $\|z\| = 1$. Thus the set $\mathcal{R}_{\mathcal{L}}$ naturally defines a subset of the simplex Σ by

$$\mathcal{C} := \{z \in \Sigma \mid ze^T \in \mathcal{R}_{\mathcal{L}}\}\,. \tag{4.2}$$

We note the following properties of $\mathcal{C}$.

PROPOSITION 4.1. *Consider a finite set of AIMD matrices $\mathcal{M}$ and the associated deterministic system* (3.1) *and the Markov chain* (3.2). *Let $\mathcal{C}$ be defined by* (4.2). *Then*

(i) *$\mathcal{C}$ is forward invariant under* (3.1)*;*

(ii) *for any solution $\{x(k)\}_{k\in\mathbb{N}}$, $x(0) \in \Sigma$ of* (3.1) *the distance*

$$\operatorname{dist}(x(k), \mathcal{C})$$

is nonincreasing;

(iii) *for any $z \in \mathcal{C}$ and any open neighborhood $U \subset \Sigma$ of z there is a $k_0 > 0$ such that $P^k(x, U) > \delta > 0$ for all $k \geq k_0$ and all $x \in \Sigma$;*

(iv) *for any initial condition $w_0 \in \Sigma$ we have almost surely*

$$\lim_{k\to\infty} \operatorname{dist}(w(k), \mathcal{C}) = 0\,.$$

Proof. (i) Let $x \in \mathcal{C}, B \in \mathcal{M}$. By definition there exists a sequence $\{A_k\}_{k\in\mathbb{N}} \in \mathcal{L}$ and $k_l \to \infty$ such that

$$\Pi(k_l) = A_{k_l} A_{k_l - 1} \dots A_1 \to ze^T\,.$$

We write $\Pi(k_l) = ze^T + \Delta_k$, where $\|\Delta_k\| \to 0$. Now we define a new sequence by repeating our initial sequence and inserting B, i.e., we consider the sequence

$$\{A_1, A_2, \dots, A_{k_1}, B, A_1, A_2, \dots, A_{k_2}, B, A_1, \dots, A_{k_3}, B, A_1, \dots\}\,.$$

Denoting products of length l of this sequence by $\Psi(l)$ we have

$$\Psi\left(l + \sum_{j=1}^{l} k_j\right) = B\Pi(k_l)\Psi\left((l-1) + \sum_{j=1}^{l-1} k_j\right) = B(ze^T + \Delta_k)\Psi\left((l-1) + \sum_{j=1}^{l-1} k_j\right)$$
$$= Bze^T + B\Delta_k\Psi\left((l-1) + \sum_{j=1}^{l-1} k_j\right),$$

where we have used that all matrices are column stochastic in the last step. As $\|\Delta_k\| \to 0$, this implies that $\Psi(l + \sum_{j=1}^{l} k_j) \to Bze^T$ as $l \to \infty$. The constructed sequence clearly lies in $\mathcal{L}$ so that $Bz \in \mathcal{C}$, which is what we wanted to show.

(ii) Let $x \in \Sigma$. Pick a $z \in \mathrm{cl}\,\mathcal{C}$ such that $\mathrm{dist}\,(x, \mathcal{C}) = \|x - z\|$. Then for $A \in \mathcal{M}$ it follows using (i) that

$$\mathrm{dist}\,(Ax, \mathcal{C}) \leq \|Ax - Az\| \leq \|x - z\| = \mathrm{dist}\,(x, \mathcal{C})\,.$$

This shows the assertion.

(iii) Fix $z \in \mathcal{C}$ and let $U \subset \Sigma$ be an open neighborhood of z. Then we may choose $\epsilon > 0$ such that $x \in \Sigma, \|x - z\| < \epsilon$ implies $x \in U$. By definition of $\mathcal{C}$ there exists a k_0 and a product $\Pi(k_0)$ such that $\|\Pi(k_0) - ze^T\| < \epsilon$. This implies for any $x \in \Sigma$ that

$$\|\Pi(k_0)x - z\| = \|(\Pi(k_0) - ze^T)x\| < \epsilon\,,$$

so that $\Pi(k_0)x \in U$ and, consequently, $P^{k_0}(x, U) > \delta > 0$ for all $x \in \Sigma$. As this probability is independent of x we see in particular that $P^k(z, U) > \delta > 0$ for all $k \geq k_0$ by considering the transition from $k - k_0$ to k.

(iv) This is an immediate consequence of Theorem 3.10. $\square$

In the terminology of Markov chains, we have proved in Proposition 4.1(iii) that each $z \in \mathcal{C}$ is *positive* and *aperiodic* for the Markov chain $\{w(k)\}_{k \in \mathbb{N}}$. For a general definition of positive and aperiodic states of an e-chain, see [23, pp. 456, 459]. Using the existence of positive and aperiodic states we obtain the following fundamental statement from [23, Theorem 18.0.2] and [8].

THEOREM 4.2. *Consider a finite set of AIMD matrices $\mathcal{M}$ and the associated Markov chain* (3.2). *Then*

(i) *there exists a unique invariant probability π;*

(ii) *for every $x \in \Sigma$ and every continuous function $h : \Sigma \to \mathbb{R}$ we have that if $w(0) = x$, then*

$$\lim_{k \to \infty} \frac{1}{k} \sum_{j=0}^{k-1} h(w(j)) = \int_\Sigma h(y) d\pi(y) \qquad \textit{almost surely;}$$

(iii) *for every $x \in \Sigma$ and every continuous function $h : \Sigma \to \mathbb{R}$ we have*

$$\int_\Sigma h(y) P^k(x, dy) \to \int_\Sigma h(y) d\pi(y) \qquad \textit{as } k \to \infty\,.$$

The previous result can be sharpened by using the special structure of the set of AIMD matrices $\mathcal{M}$.

THEOREM 4.3. *Consider a finite set of AIMD matrices $\mathcal{M}$ and the associated Markov chain* (3.2) *with its unique invariant probability π. Then*

$$\mathrm{supp}\,\pi = \mathrm{cl}\,\mathcal{C}\,.$$

Proof. We first show that $\mathcal{C} \subset \mathrm{supp}\,\pi$. Assume to the contrary that $x \in \mathcal{C} \setminus \mathrm{supp}\,\pi$. Then there exists an open neighborhood V of x with $V \cap \mathrm{supp}\,\pi = \emptyset$. By Proposition 4.1(iii) it follows for all $y \in \mathrm{supp}\,\pi$ that $P^k(y, V) > 0$ for some k large enough, which contradicts $x \notin \mathrm{supp}\,\pi$.

To show $\mathrm{supp}\,\pi \subset \mathrm{cl}\,\mathcal{C}$, let $\epsilon > 0$ and consider the set

$$U_\epsilon := \{x \in \Sigma \mid \mathrm{dist}\,(x, \mathcal{C}) > \epsilon\}\,.$$

As the distance of $w(k)$ to $\mathcal{C}$ is nonincreasing for every sample path by Proposition 4.1(ii), this shows that $P(x, U_\epsilon) > 0$ implies $x \in U_\epsilon$. Thus

$$\pi(U_\epsilon) = \int_\Sigma P(x, U_\epsilon) d\pi(x) = \int_{U_\epsilon} P(x, U_\epsilon) d\pi(x) .$$

If $\pi(U_\epsilon) > 0$, this shows that with probability 1 any evolution starting in U_ϵ stays in U_ϵ. This is a contradiction to $\mathrm{dist}\,(w(k), \mathcal{C}) \to 0$ with probability 1, which we know by Proposition 4.1(iv). This shows $\pi(U_\epsilon) = 0$, and as $\epsilon > 0$ was arbitrary, we obtain the assertion. ∎

The interesting point of the previous result is that the support of the invariant probability π is determined by the set of matrices $\mathcal{M}$, and only the distribution of mass on that set changes under variation of the probabilities ρ_i. In the next section we show that in some cases the expected values of the average can be elegantly expressed in terms of the data, without the knowledge of the invariant probability π.

5. Long term averages. From a practical point of view, we now present the main results of the paper. For the system defined in section 3.2 we know that the stochastic process $\{w(k)\}$ satisfies the strong law of large numbers. An important consequence of this result is that the vector of window sizes $w(k)$, averaged over time, converges in probability to a well-defined stochastic equilibrium. It is of interest to know what this equilibrium is given the data of the system.

Recall that $\Pi(k)$ is the random variable defined by $\Pi(k) = A(k-1)A(k-2)\dots A(0)$. It is prudent at this point to note that it follows from the discussion that the expectation of the random variable $A(k)$ is independent of k and is equal to

$$E(A(k)) = E(A(1)) = \sum_{i=1}^{\mu} \rho_i M_i. \tag{5.1}$$

Given Assumption 3.3, this immediately implies that matrix $E(A(1))$ is a positive column stochastic matrix and consequently has a unique Perron[3] eigenvector x_p given by $E(A(1))x_p = x_p$, $x_p^T y = 1$. Using the independence of the random variables $A(k)$, this shows the following statement.

PROPOSITION 5.1. *Consider a finite set of AIMD matrices $\mathcal{M}$ and let $\{A(k)\}_{k\in\mathbb{N}}$ be an i.i.d. stochastic process satisfying Assumptions* 3.1 *and* 3.3. *Then the expectation of $\Pi(k)$ is given by*

$$E(\Pi(k)) = \left(\sum_{i=1}^{\mu} \rho_i M_i \right)^k , \quad \textit{and we have} \quad \lim_{k\to\infty} E(\Pi(k)) = x_p e^T, \tag{5.2}$$

where the vector $x_p \succ 0$ is uniquely determined by

$$\left(\sum_{i=1}^{\mu} \rho_i M_i \right) x_p = x_p , \quad e^T x_p = 1 . \tag{5.3}$$

We are now interested in the long-term average of the window size. To this end we define the random variable $\overline{w}(k)$ by

$$\overline{w}(k) := \frac{1}{k+1} \sum_{i=0}^{k} w(i) = \left(\frac{1}{k+1} \sum_{i=0}^{k} \Pi(i) \right) w(0) = \overline{\Pi(k)} w(0) .$$

[3]Recall that for any column stochastic matrix $V \succ 0$ with Perron eigenvector x_p, it holds that $\lim_{k\to\infty} V^k = x_p e^T$ [4].

COROLLARY 5.2. *Consider a finite set of AIMD matrices $\mathcal{M}$, and let $\{A(k)\}_{k\in\mathbb{N}}$ be an i.i.d. stochastic process satisfying Assumptions* 3.1 *and* 3.3. *Then the expectation of $\overline{w}(k)$ is given by*

$$E(\overline{w}(k)) = \frac{1}{k+1}(I + E(A(1)) + E(A(1))^2 + \cdots + E(A(1))^k)w(0),$$

and with x_p defined by (5.3) *we have*

$$\lim_{k\to\infty} E(\overline{w}(k)) = x_p e^T w(0)\,.$$

Proof. This follows since $E(A(1))^k \to x_p e^T$ as $k \to \infty$. □

The following theorem shows how the average distribution of network capacities can be characterized.

THEOREM 5.3. *Consider a finite set of AIMD matrices $\mathcal{M}$ and the associated Markov chain* (3.2). *Let Assumptions* 3.1 *and* 3.3 *be satisfied. Then, almost surely,*

$$\lim_{k\to\infty} \overline{w}(k) = x_p e^T w(0), \tag{5.4}$$

where the vector x_p is defined by (5.3).

Proof. This is a consequence of Theorem 4.2 and Corollary 5.2. To be precise, by Theorem 4.2(ii) we have that if $w(0) \in \Sigma$, then

$$\overline{w}(k) \to \int_\Sigma w d\pi(w) =: E_\pi(w)$$

almost surely. (To obtain the desired result for vectors from the scalar results presented in Theorem 4.2, it suffices to consider the projections onto each coordinate.) If $w(0) \succeq 0$ is not in Σ, this equation scales by $e^T w(0)$ by linearity. Thus in particular $E(\overline{w}(k)) \to E_\pi(w)e^T w(0)$. As by Corollary 5.2 we have $E(\overline{w}(k)) \to x_p e^T w(0)$, which implies (5.4). □

To summarize, the previous result says that the average distribution of the resources of the network is given by the vector x_p, which can be simply obtained by finding the dominant eigenvalue of $\sum \rho_i M_i \succ 0$.

5.1. Stochastic equilibria of AIMD networks. Proposition 5.1 and Theorem 5.3 provide remarkable insights into the behavior of communication networks employing AIMD congestion control. In principle, they relate the asymptotic properties of such networks to the Perron eigenvector of $E(A(1))$. Since $E(A(1))$ is easily computable, it is possible not only to predict but also to control the asymptotic properties of such networks through judiciously manipulating the AIMD parameters and/or the probabilities ρ_i. In this context it is natural to ask whether the Perron eigenvector of $E(A(1))$ can be directly related to the AIMD parameters of the network. We now discuss some examples, where the calculation of $E(A(1))$ is particularly simple.

(i) *Time-invariant networks.* By this we mean that the network parameters cannot change in time and that there is a unique set of AIMD parameters $((\alpha_1, \ldots, \alpha_n), (\beta_1, \ldots, \beta_n))$ that is used in the construction of all matrices $M \in \mathcal{M}$. In this case it is readily shown that

$$E(A(1)) = \mathrm{diag}(\delta_1, \ldots, \delta_n) + \frac{1}{\sum_{i=1}^n \alpha_i\gamma_i}[\alpha_1\gamma_1, \ldots, \alpha_n\gamma_n]^T[1-\delta_1, \ldots, 1-\delta_n], \tag{5.5}$$

where $\delta_i = 1 - \lambda_i(1-\beta_i)$. Further, it follows directly by inspection that the Perron eigenvector of E(A(1)) is given by

$$x_p = \left[\frac{\alpha_1\gamma_1}{\lambda_1(1-\beta_1)} \ , \ \dots \ , \ \frac{\alpha_n\gamma_n}{\lambda_n(1-\beta_n)}\right]^T.$$

Consequently, the network convergence properties and the rates of convergence of $E(w(k))$ can be controlled directly by manipulating the network parameters $(\alpha_i, \beta_i, \rho_i)$. Clearly, such networks are of great interest since most practical wireline networks (including those employing TCP) fall into this category. A more detailed discussion of such network types can be found in [24].

(ii) *Time-varying networks.* Here we assume that there is a finite set of AIMD parameters $((\alpha_1^l, \dots, \alpha_n^l), (\beta_1^l, \dots, \beta_n^l))$, $l = 1, \dots, m$, and all matrices in $M \in \mathcal{M}$ are constructed as an AIMD matrix corresponding to one of these parameters. In this case it is convenient to consider two cases: (a) networks where the $\alpha_i = \alpha_i^l$ are independent of l and the β_i^l vary; and (b) networks where both α_i^l and β_i^l vary.

In the first case it is again readily shown that

$$E(A(1)) = \mathrm{diag}(\delta_1, \dots, \delta_n) + \frac{1}{\sum_{i=1}^n \alpha_i\gamma_i}[\alpha_1\gamma_1, \dots, \alpha_n\gamma_n]^T[1-\delta_1, \dots, 1-\delta_n], \tag{5.6}$$

where $\delta_i = E(\beta_i) < 1$. As before x_p can be found by inspection and is given by

$$x_p = \left[\frac{\alpha_1\gamma_1}{1-\delta_1} \ , \ \dots \ , \ \frac{\alpha_n\gamma_n}{1-\delta_n}\right]^T. \tag{5.7}$$

In the more general case it appears to be difficult to derive explicit formulae for x_p. One simplification occurs when the following situation prevails. The matrix $E(A(1))$ can be written as

$$E(A(1)) = \sum_{j=1}^{h} \sum_{M_i \in \bar{A}_j} \rho_i M_i = \sum_{j=1}^{h} Z_j. \tag{5.8}$$

In the case when the Z_j are positive matrices with a common Perron eigenvector x_p, it follows that x_p is also the Perron eigenvector of $E(A(1))$ and the stochastic equilibria of the corresponding communication network is defined by x_p. Hence, it follows that time-varying networks constructed by switching between networks with a common equilibrium results in a constituent network with the same equilibrium state (although the rate of convergence to this equilibrium is difficult to bound).

6. Experimental results. The mathematical results derived in section 5 are surprisingly simple when one considers the potential mathematical complexity of the unsynchronized network model (2.6). The simplicity of these results is a direct consequence of Assumptions 3.1 and 3.3. The objective of this section is therefore twofold: (i) to validate the unsynchronized model (2.6) in a general context; and (ii) to validate the analytical predictions of the model and thereby confirm that the aforementioned assumptions are appropriate in practical situations.

6.1. Networks of two unsynchronized flows: Ensemble averages. We first consider the behavior of two TCP flows in the dumbbell topology shown in

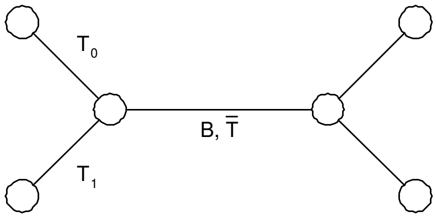

FIG. 6.1. *Dumbbell topology.*

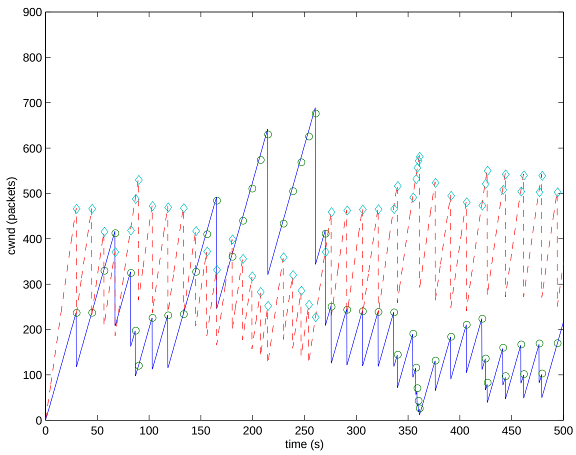

FIG. 6.2. *Evolution of window size: Predictions of the network model compared with packet-level ns2 simulation results. Key:* $\circ$ = *flow* 1 *(model),* $\Diamond$ = *flow* 2 *(model), dashed line* = *flow* 1 *(ns2), solid line* = *flow* 2 *(ns2). Network parameters:* $B = 100Mb$, $q_{max} = 80$ *packets,* $\bar{T} = 20ms$, $T_0 = 102ms$; $T_1 = 42ms$; *no background web traffic.*

Figure 6.1. Our analytic results are based upon two fundamental assumptions: (i) that the dynamics of the evolution of the source congestion windows can be accurately modeled by (2.6); and (ii) that the allocation of packet drops among the sources at congestion can be described by random variables. We consider each of these assumptions in turn.

(i) *Accuracy of dynamic model.* A comparison of the predictions made by the model (2.6) against the output of a packet-level ns2 simulation is depicted in Figure 6.2. Here, the pattern of packet drops observed in the simulation is used to select the appropriate matrix $A(k)$ from the set $\mathcal{M}$ at each congestion event when evaluating (2.6). As can be seen, the model output is very accurate. In Figure 6.3 we also plot the evolution of the linear combination $\sum_{i=1}^{n} \gamma_i w_i$, where the γ_i are defined in (2.12). It can be seen that $\sum_{i=1}^{n} \gamma_i w_i$ has the same value at each congestion event thereby validating the constraint (2.12) used in the model.

(ii) *Validity of random drop model.* It is well known that networks of TCP flows with drop-tail queues can exhibit a rich variety of deterministic drop-behaviors [9]. However, most real networks carry at least a small amount of web traffic. It is shown in [25] that already a small amount of background web traffic is enough to disrupt the coherent structure associated with phase effects and other complex phenomena previously observed in simulations of unsynchronized networks [9]. This is confirmed by

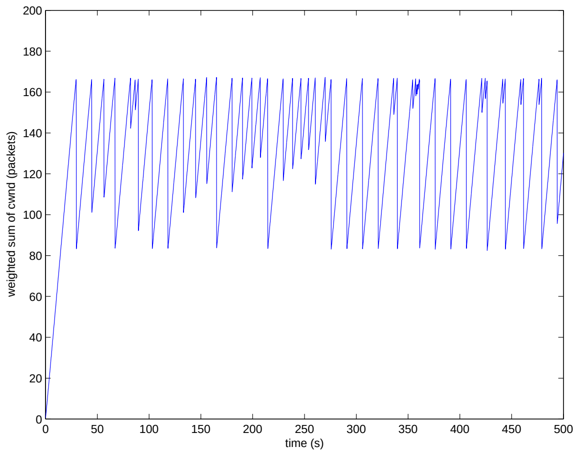

FIG. 6.3. *Evolution of* $\sum_{i=1}^{n} \gamma_i w_i$. *Network parameters:* $B = 100Mb$, $q_{max} = 80$ *packets,* $\bar{T} = 20ms$, $T_0 = 102ms$; $T_1 = 42ms$; *no background web traffic.*

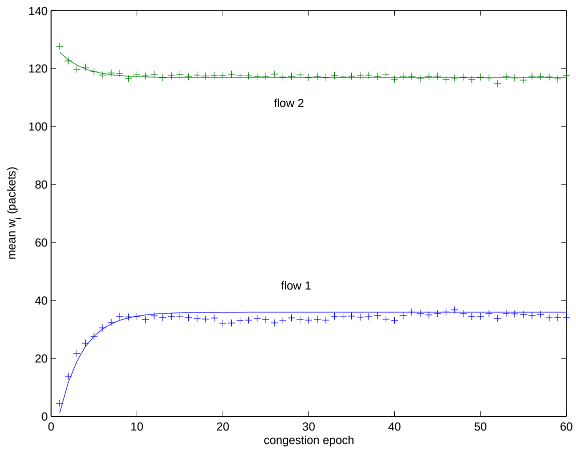

FIG. 6.4. *Variation of ensemble mean* $w_i(k)$ *with congestion epoch in dumbbell topology of Figure* 6.1. *Key:* $+ =$ ns2 *simulation result (average over* 200 *runs); solid line = Proposition* 5.1. *Network parameters:* $B = 50Mb$, $q_{max} = 50$ *packets,* $\bar{T} = 20ms$, $T_0 = 102ms$, $T_1 = 2ms$; *approximately* 0.5% *bidirectional background web traffic.*

statistical tests of this measured data, which confirm the validity of Assumptions 3.1 and 3.3.

By performing repeated packet-level simulations with different random seed values for the web traffic generator, the ensemble average congestion window can be estimated. We can also determine from the simulation results the proportion of congestion events corresponding to both flows simultaneously seeing a packet drop, flow 1 seeing a drop only, and flow 2 seeing a drop only. Using these estimates of the probabilities ρ_i, the ensemble average congestion window can also be estimated from Proposition 5.1. An example of the resulting estimates are shown in Figure 6.4. Here, we run simulations for 250 seconds with one flow started at 0 seconds and a second TCP flow started after 50 seconds (giving the first flow the opportunity to reach its

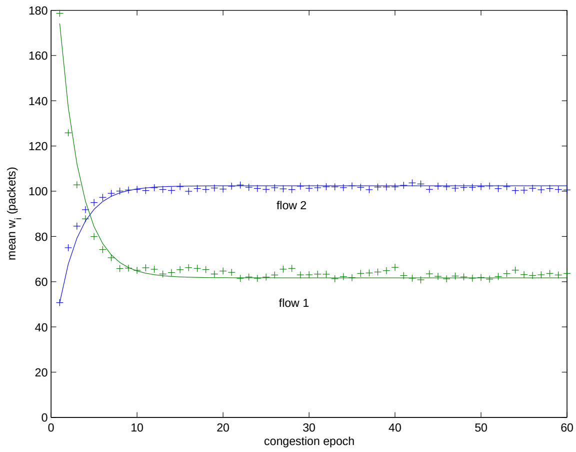

FIG. 6.5. *Variation of ensemble mean $w_i(k)$ with congestion epoch in dumbbell topology of Figure* 6.1. *Key:* $+ =$ ns2 *simulation result (average over* 200 *runs); solid line Proposition* 5.1. *Network parameters:* $B = 50Mb$, $q_{max} = 50$ *packets,* $\bar{T} = 20ms$, $T_0 = 2ms$, $T_1 = 42ms$; *approximately* 0.5% *bidirectional background web traffic.*

steady state). A small amount of bidirectional background web traffic is also included and slow-start is switched off to allow us to focus on the congestion avoidance behavior. The average congestion window evolution, estimated from 200 runs of the simulation, is plotted in Figure 6.4 together with the predictions of Proposition 5.1. It can be seen that the agreement is remarkably good. Not only is the long-term average accurately captured, but so is the manner in which the flows converge to this long-term average. That is, the model accurately describes the dynamic evolution over time, on average, of the TCP flows and thereby is useful for the analysis of both short and long-lived flows. The results shown in Figure 6.4 are for a single choice of network conditions, but the model remains accurate for other conditions; see, for example, Figure 6.5. As can be seen from the figures, the predictions of Proposition 5.1 and the ns2 simulations are consistently in close agreement.

Acknowledgment. The authors wish to thank John Foy for useful discussions.

REFERENCES

[1] E. Altman, T. Jiménez, and R. Núñez-Queija, *Analysis of two competing TCP/IP connections*, Perform. Evaluation, 49 (2002), pp. 43–55.

[2] F. Baccelli and D. Hong, *AIMD, fairness and fractal scaling of TCP traffic*, in Proceedings of the IEEE INFOCOM 2002, New York, 2002, pp. 229–238.

[3] F. Baccelli and D. Hong, *Interaction of TCP flows as billiards*, in Proceedings of the IEEE INFOCOM 2003, San Francisco, 2003, pp. 841–853.

[4] A. Berman and R. J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, SIAM, Philadelphia, 1994.

[5] P. Brown, *Resource sharing of TCP connections with different round trip times*, in Proceedings of the IEEE INFOCOM 2000, Tel Aviv, Israel, 2000, pp. 1734–1741.

[6] Y. Chait, C. V. Hollot, V. Misra, H. Han, and Y. Halevi, *Dynamic analysis of congested TCP networks*, in Proceedings of American Control Conference, San Diego, 1999, pp. 2430–2435.

[7] L. Elsner, I. Koltracht, and M. Neumann, *On the convergence of asynchronous paracontractions with applications to tomographic reconstruction from incomplete data*, Linear Algebra Appl., 130 (1990), pp. 65–82.

[8] J. H. ELTON, *An ergodic theorem for iterated maps*, Ergodic Theory Dynam. Systems, 7 (1987), pp. 481–488.

[9] S. FLOYD AND V. JACOBSON, *Traffic phase effects in packet-switched gateways*, J. Internetworking: Practice and Experience, 3 (1992), pp. 115–156.

[10] D. HARTFIEL, *Nonhomogeneous Matrix Products*, World Scientific, Singapore, 2002.

[11] J. HESPANHA, *Stochastic hybrid systems: Application to communication networks*, in Hybrid Systems, Computation, and Control, R. Alur and G. J. Pappas, eds., Proceedings of HSCC 2004, Lecture Notes in Comput. Sci. 2993, Springer, Heidelberg, 2004, pp. 387–401.

[12] J. HESPANHA, S. BOHACEK, K. OBRACZKA, AND J. LEE, *Hybrid modeling of TCP congestion control*, in Hybrid Systems, Computation, and Control, M. D. Di Benedetto and A. L. Sangiovanni-Vicentelli, eds., Proceedings of the HSCC 2001, Lecture Notes in Comput. Sci. 2034, Springer, Heidelberg, 2001, pp. 291–304.

[13] C. V. HOLLOT AND Y. CHAIT, *Nonlinear stability analysis of a class of TCP/AQM networks*, in Proceedings of the 40th IEEE Conference on Decision and Control, Orlando, FL, 2001, pp. 2309–2314.

[14] C. V. HOLLOT, V. MISRA, D. TOWSLEY, AND W. GONG, *A control theoretic analysis of RED*, in Proceedings of the IEEE INFOCOM, Anchorage, AK, 2001, pp. 1510–1519.

[15] C. V. HOLLOT, V. MISRA, D. TOWSLEY, AND W. GONG, *Analysis and design of controllers for AQM routers supporting TCP flows*, IEEE Trans. Automat. Control, 47 (2002), pp. 945–959.

[16] D. HONG AND D. LEBEDEV, *Many TCP User Asymptotic Analysis of the AIMD Model*, Tech. report INRIA 4229, INRIA Rocquencourt, 2001.

[17] R. JOHARI AND D. TAN, *End-to-end congestion control for the internet: Delays and stability*, IEEE/ACM Trans. Networking, 9 (2001), pp. 818–832.

[18] F. P. KELLY, *Mathematical modelling of the internet*, in Proceedings of the ICIAM 99 (Edinburgh), 4th International Congress of Industrial Applied Mathematics, Oxford University Press, Oxford, UK, 2000, pp. 105–116.

[19] S. S. KUNNIYUR AND R. SRIKANT, *Stable, scalable, fair congestion control, and AQM schemes that achieve high utilization in the internet*, IEEE Trans. Automat. Control, 48 (2003), pp. 2024–2029.

[20] S. LOW, F. PAGANINI, AND J. DOYLE, *Internet congestion control*, IEEE Control Systems Mag., 32 (2002), pp. 28–43.

[21] S. MASCOLO, *Congestion control in high speed communication networks using the Smith principle*, Automatica, 35 (1999), pp. 1921–1935.

[22] L. MASSOULIE, *Stability of distributed congestion control with heterogeneous feedback delays*, IEEE Trans. Automat. Control, 47 (2002), pp. 895–902.

[23] S. P. MEYN AND R. L. TWEEDIE, *Markov Chains and Stochastic Stability*, Comm. Control Engrg. Ser., Springer, London, 1993.

[24] R. SHORTEN, F. WIRTH, AND D. LEITH, *A positive systems model of TCP-like congestion control: Asymptotic results*, IEEE/ACM Trans. Networking, 14 (2006), pp. 616–629.

[25] R. SHORTEN, C. KING, F. WIRTH, AND D. LEITH, *Modelling TCP congestion control dynamics in droptail environments*, Automatica, to appear.

[26] R. N. SHORTEN, D. J. LEITH, J. FOY, AND R. KILDUFF, *Analysis and design of AIMD congestion control algorithms in communication networks*, Automatica, 41 (2005), pp. 725–730.

[27] R. SRIKANT, *Internet Congestion Control*, Control Theory 14, Birkhäuser Boston, Boston, MA, 2004.

[28] G. VINNICOMBE, *On the stability of networks operating TCP-like congestion control*, in Proceedings of the 15th IFAC World Congress on Automatic Control, Barcelona, Spain, 2002.

[29] L. XU, K. HARFOUSH, AND I. RHEE, *Binary increase congestion control (BIC) for fast long-distance networks*, in Proceedings of IEEE INFOCOM 2004, Hong Kong, 2004, pp. 2514–2524.