

Minimal Evacuation Times and Stability

Leonidas Georgiadis, *Senior Member, IEEE*, Georgios S. Paschos, Lavy Libman, *Senior Member, IEEE*, and Leandros Tassioulas, *Fellow, IEEE*

Abstract—We consider a system where packets (jobs) arrive for processing using one of the policies in a given class. We study the connection between the *minimal evacuation time* and the *stability region* of the system and show that evacuation time optimal policies can be used for stabilizing the system (and for characterizing its stability region) under broad assumptions. Conversely, we show that while a stabilizing policy can be suboptimal in terms of evacuation time, one can always design a randomized version of any stabilizing policy that achieves an optimal evacuation time in the asymptotic regime when the number of evacuated packets scales to infinity.

Index Terms—Evacuation time, stability, throughput.

I. INTRODUCTION

IN THIS paper, we consider a time-slotted system where packets arrive to n different input queues (there may be other system queues to which packets are placed during their processing; see Fig. 1). The packets are processed by a policy from an admissible class. We are interested in the stability region of such a system. A related problem is the following. For the same system, a number of packets is placed in the input queues, no arrivals may occur in the future, and it is required that the time to process all these packets (evacuation time) is minimal. Our purpose is to investigate the relation between system stability and minimum evacuation time. Under certain general assumptions on admissible policies and system statistics, we show that the stability region of the system is completely characterized by the asymptotic growth rate of minimal evacuation time. We make very few assumptions on the system structure, and hence the result is applicable to a large number of communication systems as well as more general control systems.

Manuscript received April 08, 2013; revised January 26, 2014; accepted March 04, 2014; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor I. Kesslasy. The work of L. Georgiadis and L. Tassioulas was supported by the European Union (European Social Fund ESF) and Greek national funds through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF)—Research Funding Program: Thales. Investing in knowledge society through the European Social Fund. The work of G. Paschos was supported by the WiNC project of the Action: Supporting Postdoctoral Researchers, funded by national and Community funds (European Social Fund). NICTA is funded by the Australian Department of Communications and the Australian Research Council through the ICT Centre of Excellence program.

L. Georgiadis is with the Department of Electrical and Computer Engineering, Aristotle University of Thessaloniki, Thessaloniki 54124, Greece.

G. S. Paschos is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 01960 USA (e-mail: gpaschos@gmail.com).

L. Libman is with the School of Computer Science and Engineering, UNSW Australia and NICTA, Sydney, NSW 2052, Australia.

L. Tassioulas is with the Department of Computer Engineering and Telecommunications, University of Thessaly, Volos 38221, Greece.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2014.2312271

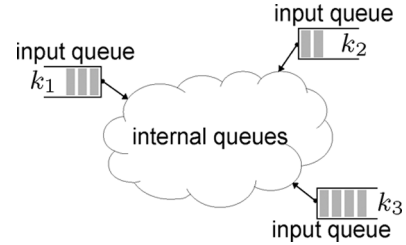


Fig. 1. Packetized system with $\mathbf{k} = (k_1, k_2, k_3)$ packets at the inputs at the beginning of time; this paper deals with policies that evacuate such systems in minimum time.

However, we point out that the result, while intuitive, has to be applied with caution since there are systems for which its application leads to wrong conclusions. We then look at the reverse problem: Given a throughput optimal policy, is this policy also optimal in evacuating packets? We demonstrate that this is not usually true, but randomized versions of throughput optimal policies exist that achieve evacuation time optimality in an asymptotic sense.

Concepts akin to evacuation time and their relation to stability have been investigated in earlier works. Baccelli and Foss [1] consider a system fed by a marked point process and operating under a given policy. The concept of *dater* is used to describe the time of last activity in the system, if the system is fed only by the m th to n th, $m \leq n$ of the points of the marked process. Assuming that the dater is a deterministic function of the arrival times and the marks of the point process, and under additional assumption on dater sample paths, they show that stability under the specified policy is characterized by the asymptotic behavior of daters. These results are extended to continuous-time input processes by Altman [2]. In our setup, the system evolution may depend on random factors as well as the characteristics of the arrival process. Moreover, we do not make sample path assumptions on specific policies. Rather, we specify features that admissible policies may have, and based on these, we characterize the stability region of the class of admissible policies by the asymptotic growth rate of minimal (over all admissible policies) evacuation times.

A different yet related concept, *workload*, was introduced by Harrison [3]; under a specific policy, the workload $w(t)$ is defined as the time the server must work to clear all of the inventory of the system at time t . This basic concept is used in the analysis of fluid limits to derive significant results and provide intuition for good control policies in specific complex networks [4]. In this paper, we concentrate on the minimal evacuation time over a whole class of policies and relate its asymptotic rate of growth to system stability. There is also similarity between the manner we obtain the asymptotic growth rate of the minimal evacuation time, and the scaling that is done in order to obtain fluid limit approximation of a system when it is operated under a specific policy [4]. However, the limits obtained

by the two approaches concern different quantities. In our case, we work directly with the original system, and construction of the fluid limit is not needed. We also consider the stability region of a system under a class of policies rather than specific policies, and we do not impose specific requirements on system structure.

A number of works have recently used evacuation times as a tool for proving important results in communications networks. In [5], the stability proof of dynamic index coding policies is based on minimal evacuation times. Angelakis *et al.* [6] study the complexity of evacuation-optimal scheduling policies for wireless networks. Evacuation times are used in [7] for performance analysis of the two-receiver Broadcast Erasure channel with feedback. In our current work, we formalize the general system model and prove the key necessity and sufficiency results in order to facilitate its future use in applications.

II. SYSTEM MODEL AND ADMISSIBLE POLICIES

In the following, we use the vector $\mathbf{x} = [x_1, x_2, \dots, x_n]^\top$. Also, $\mathbf{x} \geq \mathbf{y}$ means $x_i \geq y_i, i = 1, 2, \dots, n$ and

$$[\mathbf{x}] \doteq [[x_1], \dots, [x_n]]$$

where $[x]$ is the least integer larger than or equal to x . With \mathbf{m}, \mathbf{k} we denote vectors of nonnegative integers and with \mathbf{r}, \mathbf{s} vectors of nonnegative reals.

We consider a time-slotted system where slot $t = 0, 1, \dots$ corresponds to the time interval $[t, t + 1)$. The system has n input queues of infinite length where packets¹ arrive; these packets may have certain properties, e.g., service times, routing options, etc. There may be additional queues in the system, where packets may be placed during its operation; we call these queues *internal*. At the beginning of time-slot t , $A_i(t)$ packets arrive at input i . (In particular, we use $A_i(0)$ to denote the number of packets in the queue of input i when the system commences operation at $t = 0$.) We assume that the arrival processes satisfy the ergodicity condition

$$\lim_{t \rightarrow \infty} \frac{\sum_{\tau=0}^t A_i(\tau)}{t} = \lambda_i, \quad i = 1, 2, \dots, n \quad (1)$$

as well as

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{\tau=1}^t A_i(t) \right]}{t} = \lambda_i, \quad i = 1, 2, \dots, n. \quad (2)$$

The operation of the system is characterized by a finite set of *system states* \mathcal{S} , and control sets \mathcal{G}_s for each $s \in \mathcal{S}$: If at the beginning of a slot the system state is $s \in \mathcal{S}$, one of the available controls $g \in \mathcal{G}_s$ is applied. There may be randomness in the behavior of the system, that is, given s and g at the beginning of a slot, the system state and the results at the end of a slot (e.g., packet erasures) may be random (e.g., due to ambient noise in wireless networks); the assumptions about the permissible random distributions will be made precise later.

Arriving packets are processed by the system following a policy π , belonging to a class of admissible policies Π . At time t , when the system state is s , an admissible policy specifies: 1) the control $g \in \mathcal{G}_s$ to be chosen, and 2) an action α among a set of

available actions \mathcal{A}_g when control g is chosen. An action specifies how packets are handled within the system. The choice of controls and actions depends on the *system history* up to t , denoted by \mathcal{H}_t . The history \mathcal{H}_t includes all information about packet arrival instants, packet departure instants, system states, controls, actions taken, and results, up to and including time t .

Remark: In the mathematical analysis of systems, the “state” of the mathematical model may include part of \mathcal{H}_t , and actions are usually not distinguished from controls. For the purposes of this work, the terms *system states* and *controls* are explicitly used to refer to the operational characteristics of the system and are distinct from the actions taken once the system characteristics are set. For example, the sizes of the queues at time t are part of the information captured by \mathcal{H}_t , rather than the system state. Also, we emphasize that the choice of one action or another within a given control (for example, which particular packet is transmitted from a given queue) does not affect the system state or slot outcome. This distinction is needed in order to define well the statistical assumptions used in the development that follows. We next present several examples to clarify these notions.

Example 1: Controls Versus Actions. Assume a wireless transmitter that can transmit to a destination over one of two channels, I or II (e.g., over two different carriers). Data arriving at the transmitter are classified in two types A, B. Packets from each of the classes are placed in distinct infinite-size queues.

The channels can be in one of four states, $(s_1, s_2) \in \{(l, h), (h, l), (l, l), (h, h)\}$. The controls \mathcal{G}_s available when in state $s = (s_1, s_2)$ determine: 1) whether a channel will be used for transmission, and 2) the transmission power p over the channel(s) selected for transmission. Suppose at a given slot we choose $(1, p)$, then the rate of transmission on channel I will be $r(p, s_1)$, while channel II cannot be used in this slot. In this case, the action set \mathcal{A}_g consists of two elements, a_A and a_B , indicating the type of packet to be transmitted over channel I. If, say, the action specifies that type-A packets are selected but there are no packets of type A, then a dummy (non-information-bearing) packet is used. The choice of action does not make a difference to the dynamics of the system state. \square

Example 2: Consider a communication system consisting of two nodes, a, b . Arriving packets are stored in an infinite queue at node a and must be delivered to node b . The two nodes are connected with two links, ℓ_1, ℓ_2 , at *most one* of which may be activated at a time. If link ℓ_1 is activated, a packet can be successfully transmitted in one slot, but both links cannot be activated for the next nine slots. If link ℓ_2 is activated, a transmitted packet is erased with probability 0.5 (and received successfully with probability 0.5) and both links can be activated in the next slot.

The states for this system can be described by the elements $\{0, 1, 2, \dots, 9\}$, where state 0 means that both links can be activated and state $i \geq 1$ means that no link can be activated for the next i slots.

The control set for state 0 is $\mathcal{G}_0 = \{g_0, g_1, g_2\}$, where g_0 means no link activation, g_1 means activation of link ℓ_1 , and g_2 means activation of link ℓ_2 . The control set for the rest of the states consists only of g_0 . From state 0, if control g_0 or g_2 is taken, the state returns to 0 in the next slot, while if g_1 is taken, the state becomes 9. From state $i \geq 1$ the system moves to state $i - 1$ in the next slot.

¹In this paper, we use the term *packet*, which describes an arriving unit in a communication network. However, our work applies to any general service system with arrival processes and queues, e.g., manufacturing systems, road networks, network switches, etc. Therefore, the subsequent discussion and results should be understood generically.

At state 0, control g_0 results in “inactive” channels. If control g_2 is taken, the result is either “unsuccessful” or “successful” transmission on channel ℓ_2 —a random event—and if control g_1 is taken, the result is “successful transmission” on channel ℓ_1 . Here, a “successful” transmission should be taken to mean that a packet will be successfully delivered to node b if transmitted in the slot (in other words, a “good” underlying transmission link); it does not preclude the respective control to include a possible action that does not make a transmission in the slot at all.

The controls under which one of the links is activated are associated with two actions: 1) the action of transmitting a packet on the corresponding link, and 2) the action of not transmitting a packet (“null” action). For the control that does not activate any link, the associated action set is only the “null” action. \square

Departures: There are well-defined times when each arriving packet is considered to depart from the system. For example, in a store-and-forward communication network where a packet arrives at node i and must be delivered to a single node j , it is natural to consider the departure time as the time at which this packet is delivered to node j . Similarly, if the packet must be multicast to a subset \mathcal{K} of the nodes, the departure time of the packet can be defined as the first time at which all nodes in \mathcal{K} receive the packet. However, in some systems, several definitions of departure times may make sense, and the particular choice depends on the performance measures of interest. As an example, consider the case where network coding is used to transmit encoded packets. In this case, a packet p arriving at a single-destination node j may be considered as departed when the destination node j can decode the packet based on the packets already received by that node. On the other hand, if the decoded packet is still needed for decoding of other packets, it may be of interest to define the departure time of p as the first time the packet is not needed for further decoding. At any time between the arrival and departure times of a packet p , we say that p is “in the system.”

Features of Admissible Policies: We assume that the class of admissible policies Π must possess the following features.

- F1) At any time t , the history of the system up to t , \mathcal{H}_t is fully known.
- F2) At any time t , if the system state is $s(t)$, a policy may choose any control $g(t) \in \mathcal{G}_{s(t)}$ and action $a \in \mathcal{A}_{g(t)}$. No further constraints are imposed on the sequence of controls and actions that may be chosen by the policy, namely $(g(1), g(2), \dots)$ and $(a(1), a(2), \dots)$.²
- F3) If at time t there are \mathbf{k} packets at the input queues, it is permissible to pick any $\mathbf{m} \leq \mathbf{k}$ packets and continue processing the \mathbf{m} packets, along with other packets that may be in the system, *without taking into account* the remaining $\mathbf{k} - \mathbf{m}$ packets.

We emphasize that the above features refer to the *class* of admissible policies, rather than any policy in particular. Thus, any given policy may not necessarily have all of the features; we

²In particular, we emphasize that the set of controls and actions available to a policy at time t may depend on the system state at that time, but may not be constrained by the history of the system, including the history of packet arrivals and queue sizes. Therefore, if necessary, the outcomes of actions involving packets from one or more queues should be meaningfully defined even when some of the queues involved are empty, e.g., if an action selects a packet from an empty queue for transmission, “dummy” (non-information-bearing) packet may be transmitted.

only require that policies that have the above features are not excluded from the admissible class (e.g., due to the system structure or operational constraints). Moreover, these features are generic and easy to verify, they hold naturally in many systems, and apply to classes of policies as widely as possible while still maintaining the desired connection between evacuation times and stability region. However, in certain systems, the above features may not be available; in such systems, the results of this work, however intuitive, may not hold. The following examples illustrate the importance of the above features and clarify their definitions.

Example 3: Time-Average Constraints Invalidate F2. To emphasize the importance of F2, consider a single transmitter where the transmission power $p(t)$ is a control chosen at every time-slot from a set $\{0, 1, 2\}$. The number of packets served at every slot depends on the chosen transmitted power; for simplicity, assume exactly $p(t)$ packets are served. Consider the throughput region achieved by the class of policies with a time-average power consumption below a given bound \bar{p} , i.e., $\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{t=0}^{n-1} p(t) \leq \bar{p}$, where $\bar{p} < 2$. Note that the bound on the time-average power consumption does not impose a constraint on the power consumed at any particular time, and thus at any time-slot t , a policy may choose any $p(t) \in \{0, 1, 2\}$. However, this class of policies violates F2 since some control sequences $(g(1), g(2), \dots)$ where $g(t) \in \{0, 1, 2\}$ are not admissible due to the above constraint. Consequently, our results will not apply to this class of policies. \square

Example 4: Two-Transmitter Aloha-Type System F3. Consider a system consisting of two transmitters attempting to transmit arriving packets to a single destination. Each transmitter has its own queue. Activation of both transmitters in the same slot results in the loss of any transmitted packet. We can model this system by considering that it has a single state, and that the control set consists of pairs (g_1, g_2) where $g_i = 1$ ($g_i = 0$) indicates that transmitter $i \in \{1, 2\}$ becomes active (inactive). If both queues are nonempty, then the transmitters are activated with probability $q(t)$, $0 \leq q(t) \leq 1$, $q(t)$ being the *same* for both transmitters. Hence, the controls $(0, 0), (1, 0), (0, 1), (1, 1)$ are chosen randomly with probabilities of $(1 - q_t)^2, q_t(1 - q_t), q_t(1 - q_t), q_t^2$, respectively. When one of the queues is empty, any of the controls can be chosen either randomly or deterministically.

In this example, it is not permissible for a policy to serve one of the transmitters first; e.g., if $\mathbf{k} = (1, 1)$, a policy is not allowed to select (with probability 1) to transmit first the vector of packets $(1, 0)$ and next the vector $(0, 1)$. Thus, the class of admissible policies does not satisfy Feature F3, and our results on the relationship between evacuation time and stability region do not cover this Aloha-type system. \square

Evacuation Times: At the beginning of slot 0, let the system state be s , and let there be $k_i \geq 0, i = 1, \dots, n$ packets at input i and no arrivals afterwards, i.e., $A_i(0) = k_i, A_i(t) = 0, t = 1, 2, \dots$. Let $T_s^\pi(\mathbf{k}) \geq 0, \mathbf{k} \neq \mathbf{0}$ be the time it takes until all of these packets depart from the system under policy π . We call $T_s^\pi(\mathbf{k})$ the *evacuation time* under policy π when the system starts in state s with \mathbf{k} packets at the inputs, and denote its average value, $\bar{T}_s^\pi(\mathbf{k}) = \mathbb{E}[T_s^\pi(\mathbf{k})], \mathbf{k} \neq \mathbf{0}$. It will also be convenient to define $\bar{T}_s^\pi(\mathbf{0}) = 1$, a convention that has the meaning of advancing one slot whenever the system is empty.

Let

$$\bar{T}_s^*(\mathbf{k}) \triangleq \inf_{\pi \in \Pi} \bar{T}_s^\pi(\mathbf{k}) \quad (3)$$

and

$$\bar{T}^*(\mathbf{k}) \triangleq \max_{s \in \mathcal{S}} \bar{T}_s^*(\mathbf{k}).$$

We call $\bar{T}^*(\mathbf{k})$ the *critical evacuation time function*. It will be seen that under certain statistical assumptions, this function determines the stability region of the considered policies.

Note that according to the definition of $\bar{T}_s^*(\mathbf{k})$, for any $\epsilon > 0$ we can always find an admissible policy π such that

$$\bar{T}_s^\pi(\mathbf{k}) \leq \bar{T}_s^*(\mathbf{k}) + \epsilon. \quad (4)$$

This fact will be used repeatedly in the development that follows.

Statistical Assumptions: Next, we present statistical assumptions regarding the system under consideration.

SA1) For all \mathbf{k} and $s \in \mathcal{S}$, it holds $\bar{T}_s^*(\mathbf{k}) < \infty$.

SA2) Individual packet characteristics (such as packet length, service time, etc.) are independent and statistically identical for all packets arriving at a given input and independent across inputs.

SA3) If at the beginning of a slot t the system state is $s_t \in \mathcal{S}$ and control $g_t \in \mathcal{G}_{s_t}$ is taken, the results at time $t+1$ are independent of the system history before t . However, the system state s_{t+1} and the results at time $t+1$ may depend on both s_t and g_t . Hence, the system states may be affected by the controls (but not actions) taken by a policy. Formally, if W_t is the (random) outcome at the end of a slot, we have for all t

$$\Pr(W_{t+1}, S_{t+1} \mid s_t, g_t, \mathcal{H}_t) = \Pr(W_{t+1}, S_{t+1} \mid s_t, g_t).$$

SA4) At time $t = 0, 1, 2, \dots$, let there be \mathbf{k} packets in the system (where k_i is the number of packets still in the system from those that originally arrived at input i ; they may or may not still be at the input queues). There is a policy π_h that can process all these packets until they all depart from the system by time $t + F^{\pi_h}(\mathbf{k})$ ($F^{\pi_h}(\mathbf{k})$ may be random), such that

$$\mathbb{E}[F^{\pi_h}(\mathbf{k})] \leq C \sum_{i=1}^n k_i \quad (5)$$

where C is a finite constant (which may depend on system statistics but not on \mathbf{k}).

SA5) Let \mathbf{e}_i be the unit n -dimensional vector with 1 at the i th coordinate and 0 elsewhere. It holds for all $i = 1, \dots, n$, and all \mathbf{k} and s

$$\bar{T}_s^*(\mathbf{k}) - \bar{T}_s^*(\mathbf{k} + \mathbf{e}_i) \leq D_0 < \infty. \quad (6)$$

Statistical Assumption SA4 is easy to verify in several systems. For example, in a communication network a policy that usually satisfies this assumption is the one that picks one of the \mathbf{k} packets, transmits it to its destination, then picks another packet and so on, until all the packets are delivered to their destinations. Note that assumption SA4 implies SA1; we keep assumption SA1 separate because, as will be seen shortly, only this assumption is needed to establish the key property (namely, subadditivity) of $\bar{T}^*(\mathbf{k})$.

Statistical Assumption SA5 is needed to justify a technical condition in the development that follows. This assumption may also be easy to verify for several systems. It says that if the number of packets at the system inputs at time 0 is *increased* by one, then the minimal average evacuation time under any initial state cannot be *decreased* by more than a fixed amount. For example, this assumption is always satisfied if $\bar{T}_s^*(\mathbf{k})$ is nondecreasing in \mathbf{k} , i.e.,

$$\bar{T}_s^*(\mathbf{k}) \leq \bar{T}_s^*(\mathbf{k} + \mathbf{e}_i). \quad (7)$$

In particular, it can be easily shown that condition (7) holds if policies have the ability to generate “dummy” packets (i.e., packets that bear no information and are used just for policy implementation), a feature that is available in many communication networks. Indeed, assume that at time $t = 0$ the system is in state s and there are \mathbf{k} packets at the system inputs. Pick $\epsilon > 0$ and a policy π such that

$$\bar{T}_s^\pi(\mathbf{k} + \mathbf{e}_i) \leq \bar{T}_s^*(\mathbf{k} + \mathbf{e}_i) + \epsilon.$$

Consider the following policy π_0 for evacuating \mathbf{k} packets: Generate a “dummy” packet for input i , place the $\mathbf{k} + \mathbf{e}_i$ packets at the inputs, and use policy π to evacuate the system. By construction, $\bar{T}_s^{\pi_0}(\mathbf{k}) \leq \bar{T}_s^\pi(\mathbf{k} + \mathbf{e}_i)$ (the inequality may be strict if the departure time of the dummy packet turns out to be strictly larger than the departure times of the rest of the packets). Hence

$$\bar{T}_s^*(\mathbf{k}) \leq \bar{T}_s^{\pi_0}(\mathbf{k}) \leq \bar{T}_s^\pi(\mathbf{k} + \mathbf{e}_i) \leq \bar{T}_s^*(\mathbf{k} + \mathbf{e}_i) + \epsilon.$$

Since ϵ is arbitrary, (7) follows.

III. PROPERTIES OF $\bar{T}^*(\mathbf{k})$

The following property of the critical evacuation time function will play a key role in the subsequent analysis. We provide the proofs to our claims in the appendixes.

Lemma 5: The Critical Evacuation Time Function is subadditive, i.e., the following holds for $\mathbf{m} \geq \mathbf{0}$, $\mathbf{k} \geq \mathbf{0}$:

$$\bar{T}^*(\mathbf{k} + \mathbf{m}) \leq \bar{T}^*(\mathbf{k}) + \bar{T}^*(\mathbf{m}). \quad (8)$$

The proof of Lemma 5 is given in Appendix A.

Let \mathbb{N}_0 and \mathbb{R}_0 be respectively the set of nonnegative integers and nonnegative real numbers. We extend the domain of $\bar{T}^*(\mathbf{k})$ from \mathbb{N}_0^n to \mathbb{R}_0^n as follows. For $\mathbf{r} \in \mathbb{R}_0^n$, let

$$\bar{T}^*(\mathbf{r}) = \bar{T}^*(\lceil \mathbf{r} \rceil). \quad (9)$$

The function $\bar{T}^*(\mathbf{r})$ is not necessarily subadditive in \mathbb{R}_0^n since, in general, subadditivity at integer points does not imply subadditivity over \mathbb{R}_0 . For example, the function $f(2l) = al$ and $f(2l+1) = al + A$, $l = 0, 1, \dots$, with $a < A$, is subadditive in \mathbb{N}_0 , while for $r_1 = r_2 = 1.5$, $f(\lceil r_1 + r_2 \rceil) = f(3) = a + A$ and $f(\lceil r_1 \rceil) + f(\lceil r_2 \rceil) = 2a < f(\lceil r_1 + r_2 \rceil)$. However, as we show next, $\bar{T}^*(\mathbf{r})$ has asymptotically linear *rate of growth*, a fundamental property of subadditive functions.

Theorem 6: For any $\mathbf{r} \in \mathbb{R}_0^n$, the limit function

$$\hat{T}(\mathbf{r}) = \lim_{t \rightarrow \infty} \frac{\bar{T}^*(t\mathbf{r})}{t} \quad (10)$$

exists and is finite, positively homogeneous, convex and Lipschitz continuous, i.e., for a positive constant D it holds

$$|\hat{T}(\mathbf{r}) - \hat{T}(\mathbf{s})| \leq D \sum_{i=1}^n |r_i - s_i|.$$

Moreover, for any converging sequence $\mathbf{r}_t \in \mathbb{R}_0^n$ such that $\lim_{t \rightarrow \infty} \mathbf{r}_t = \boldsymbol{\lambda} < \infty$, it holds

$$\lim_{t \rightarrow \infty} \frac{\bar{T}^*(t\mathbf{r}_t)}{t} = \hat{T}(\boldsymbol{\lambda}). \quad (11)$$

Here, “positively homogeneous” means that for any $\rho \geq 0$,

$$\hat{T}(\rho\mathbf{r}) = \rho\hat{T}(\mathbf{r}). \quad (12)$$

The proof of Theorem 6 is given in Appendix B.

IV. STABILITY—NECESSITY

Let $D_{s,i}^\pi(t), t \geq 1$, be the number of packet arrivals at input i that have departed from the system during time-slot t under policy $\pi \in \Pi$ when the system starts in state s . Define also $D_{s,i}^\pi(0) = 0$. In the following, we will use the notation

$$\tilde{A}_i(t) = \sum_{\tau=0}^t A_i(\tau) \quad \tilde{D}_{s,i}^\pi(t) = \sum_{\tau=0}^t D_{s,i}^\pi(\tau)$$

to denote the cumulative number of arrivals and departures respectively up to time t . Hence, the number of packet arrivals at input i that are still in the system at time t is $Q_{s,i}^\pi(t) = \tilde{A}_i(t) - \tilde{D}_{s,i}^\pi(t)$ (these packets may at time t be scattered among internal system queues as well as the original input queue). We define the vector $\mathbf{Q}_s^\pi(t) = (Q_{s,i}^\pi(t))_{i=1}^n$ and the total system occupancy

$$Q_s^\pi(t) = \sum_{i=1}^n Q_{s,i}^\pi(t).$$

Let \mathcal{M} be a probability measure over the space of permissible arrival processes; in other words, \mathcal{M} captures the statistical assumptions about the arrival processes, such as the distribution of the arrival sizes, whether or not the arrivals are independent over time and between queues, etc. Let \mathcal{M}_λ be a probability measure over arrival processes that satisfy ergodicity conditions (1) and (2) with a rate vector $\boldsymbol{\lambda}$.

Definition 7: System Stability. A policy $\pi \in \Pi$ is called stable for an arrival rate vector $\boldsymbol{\lambda} \geq \mathbf{0}$, if under any initial system state s , the following holds:

$$\lim_{q \rightarrow \infty} \limsup_{t \rightarrow \infty} \Pr(Q_s^\pi(t) > q) = 0 \quad (13)$$

(where the probability in (13) is taken with respect to the arrival process statistics \mathcal{M}_λ and the system state transitions). A stochastic process $Q_s^\pi(t)$ satisfying (13) is called *substable* in [8].

The stability region \mathcal{R}^π of a policy π is the closure of the set of the arrival rate vectors for which the policy is stable. The stability region \mathcal{R} of the system is the closure of the union of $\mathcal{R}^\pi, \pi \in \Pi$.³ A policy whose stability region is \mathcal{R} is called *stabilizing*.

We show in Theorem 8 below that under (1) and (2), it holds $\mathcal{R} \subseteq \{\mathbf{r} \geq \mathbf{0} : \hat{T}(\mathbf{r}) \leq 1\}$. Furthermore, in Section V, we show that under the assumption that the packet arrival vectors

are independent and identically distributed (i.i.d.) over time, we also have $\{\mathbf{r} \geq \mathbf{0} : \hat{T}(\mathbf{r}) \leq 1\} \subseteq \mathcal{R}$, hence $\mathcal{R} = \{\mathbf{r} \geq \mathbf{0} : \hat{T}(\mathbf{r}) \leq 1\}$, and we present an explicit policy called *Epoch-based* that is stabilizing.

Theorem 8: (Necessity). Let (1) and (2) hold. If $\mathbf{r} \in \mathcal{R}$, then

$$\hat{T}(\mathbf{r}) \leq 1.$$

The proof is given in Appendix C.

We note that there are classes of policies for which the limit $\hat{T}(\boldsymbol{\lambda})$ can be formally defined, but Theorem 8 does not hold in all its generality since some of the features of admissible policies in Section II are not satisfied. Consider the following examples.

Example 9: F3 Not Satisfied. Consider the following system. There are two input queues. If only one of the queues is nonempty, a single packet from that input is processed in one time-slot. If both queues are nonempty, then a pair of packets from both queues must be processed in three time-slots. This system is a simplified version of the system in Example 4, and the specified policies do not satisfy Feature F3. It can be easily seen that $\hat{T}^*(k_1, k_2) = 3 \min\{k_1, k_2\} + |k_1 - k_2|$, hence formally

$$\hat{T}(r_1, r_2) = 3 \min\{r_1, r_2\} + |r_1 - r_2|.$$

The region $\hat{T}(r_1, r_2) \leq 1$ is described by

$$\mathbf{r} \geq \mathbf{0} : r_1 + 2r_2 \leq 1, r_1 \geq r_2 \text{ or } 2r_1 + r_2 \leq 1, r_2 \geq r_1. \quad (14)$$

Clearly, the vector $(1/2, 1/2)$ does not belong in this region. Consider, however, that single packets arrive in alternating time-slots to inputs 1 and 2, hence the arrival rate vector is $(1/2, 1/2)$. Then, simply processing immediately the arriving packets results in a stable policy.

Notice also that the region in (14) is not convex, while the region in Theorem 8 is convex since $\hat{T}(r_1, r_2)$ is convex. \square

The arrival processes in the previous example are not stationary, hence one may wonder whether imposing slightly stronger assumptions on the arrival processes would render the claim of Theorem 8 valid. An example is presented below, where the arrival processes are i.i.d., but Theorem 8 still does not hold since admissible policies do not satisfy Feature F3.

Example 10: i.i.d. Arrivals and F3. Let $M > 1$ and consider a system with a single input and the following restriction on the policies. If the number of packets in the inputs is

$$k = lM + v \quad 0 \leq v \leq M - 1$$

then a policy may either decide to idle in a slot or to transmit m packets, $1 \leq m \leq M + v$, in which case it takes l slots to process all m packets. Under this restriction, we have

$$\bar{T}^*(k) = \sum_{i=1}^l i = \frac{l(l+1)}{2}$$

hence

$$\begin{aligned} \hat{T}(r) &= \lim_{t \rightarrow \infty} \frac{\bar{T}^*([tr])}{t} \\ &= \lim_{t \rightarrow \infty} \frac{1}{t} \left(\frac{([tr] - v_t)([tr] - v_t + 1)}{2M^2} \right) = \infty. \end{aligned}$$

Applying formally Theorem 8, we deduce that the system is unstable for any positive arrival rate. Consider, however, that the arrival process is i.i.d. but bounded, such that at most $2M - 1$ packets may arrive at the beginning of each slot (including slot 0, i.e., to be in the system when it commences operation). Then,

³We emphasize that the stability region of a policy may in general depend on the permitted statistical assumptions about the arrival processes; for example, a policy may be unstable for a certain rate vector $\boldsymbol{\lambda}$ if general stationary arrival processes are allowed, but become stable if the individual queue arrivals are required to be independent. The above definition of stability is generic and captures a number of common definitions of stability in the literature, and the subsequent discussion in this section is orthogonal to any specific assumptions imposed on the arrival process, beyond the basic ergodicity condition of (1) and (2).

the policy that transmits all the packets immediately is stable, i.e., under the stated conditions on arrival process statistics, the system is stable for any arrival rate $\lambda \leq 2M - 1$. \square

For the systems described in the last two examples, there were rates outside the region obtained by formally using $\hat{T}(\mathbf{r})$, for which the systems were stabilizable. The next example shows an opposite case, namely where the system is unstable for rates inside the formally obtained region (again, due to not satisfying Feature F3).

Example 11: System With Priorities and Switchover Times. Consider a single server with two inputs, where arrivals at input 1 have priority over arrivals at input 2: If there are packets from input 1 in the system, one of these packets must be served, while packets from input 2 can be delayed. All packets are served within one slot, however there is a preparatory time of one slot to switch the system to serve packets from another input (i.e., an idle slot occurs whenever a switch from one input to the other is required).

This system has two states, s_1, s_2 , where state s_i means that the server is set to serve packets of input i . For this system

$$\begin{aligned} \bar{T}_{s_1}^*(k_1, k_2) &= \begin{cases} k_1 + 1 + k_2, & \text{if } k_2 \neq 0 \\ k_1, & \text{if } k_2 = 0 \end{cases} \\ \bar{T}_{s_2}^*(k_1, k_2) &= \begin{cases} 1 + k_1 + 1 + k_2, & \text{if } k_1 \neq 0, k_2 \neq 0 \\ 1 + k_1, & \text{if } k_2 = 0 \\ k_2, & \text{if } k_1 = 0. \end{cases} \end{aligned}$$

Hence, $\hat{T}(r_1, r_2) = r_1 + r_2$ and the region formally is

$$\{\mathbf{r} \geq \mathbf{0} : r_1 + r_2 \leq 1\}.$$

Consider, however, an arrival pattern where the system starts at state s_1 , and a single packet arrives at input 1 at every $t = 4j, j = 0, 1, \dots$; hence, $\lambda_1 = 0.25$. Packets at input 2 arrive according to an i.i.d. process of rate $\lambda_2 > 0.5$. It can be easily seen that in any interval $[4j, 4j + 8)$, the number of packets served from input 2 cannot be larger than 4, hence the departure rate for packets at input 2 cannot be more than 0.5 and the system is unstable, even though $\lambda_1 + \lambda_2 < 1$. \square

One may wonder whether if the initial state of the system at time $t = 0$ is fixed, say $s(0) = s_0$, then stability is determined by $\bar{T}_{s_0}^*(\mathbf{k})$ only. The following final example illustrates that this is not always the case, i.e., the condition of Theorem 8 applies to the critical (worst-case) evacuation time function, and not just the evacuation time function corresponding to s_0 .

Example 12: Importance of Maximizing Over States. Consider a system with two servers, where server 1 takes l slots to serve a packet, and server 2 takes $L > l$ slots. The system can be in one of three states, $(0, 0), (1, 0), (0, 1)$, where 0 denotes an inactive and 1 denotes an active server. Suppose that there are no (or null) controls, and that state transitions are random with the following transition probabilities. For $(0, 0) = \mathbf{0}$

$$\Pr\{(1, 0) \mid \mathbf{0}\} = \Pr\{(0, 1) \mid \mathbf{0}\} = \Pr\{(0, 0) \mid \mathbf{0}\} = \frac{1}{3}$$

$$\Pr\{(1, 0) \mid (1, 0)\} = \Pr\{(0, 1) \mid (0, 1)\} = 1.$$

If the system starts at state $(0, 0)$, it takes on average 1.5 slots to move to one of the other states, and the transition to either state occurs with equal probability. Then, since no further change of states occurs afterwards, it will take either lk or Lk slots to evacuate k packets. Hence

$$\bar{T}_{(0,0)}^*(k) = \frac{3}{2} + \frac{l+L}{2}k.$$

It can also be easily verified that $\bar{T}_{(1,0)}^*(k) = lk$ and $\bar{T}_{(0,1)}^*(k) = Lk$; thus

$$\bar{T}^*(k) = \max \left\{ \frac{3}{2} + \frac{l+L}{2}k, lk, Lk \right\}$$

and $\hat{T}(r) = Lr$, which leads to the stability condition $\lambda \leq \frac{1}{L}$.

Assume now that the system starts in state $s = (0, 0)$ (an initial state that may be “natural” in some sense), and formally use $\bar{T}_{(0,0)}^*(k)$ in place of $\bar{T}^*(k)$. Then, we would conclude that $\hat{T}(r) = \frac{l+L}{2}r$, and hence that the system is stable when

$$\lambda \leq \frac{2}{l+L}.$$

This, however, is wrong since for $\frac{2}{l+L} > \lambda > \frac{1}{L}$, under state transition $(0, 0) \rightarrow (0, 1)$, an event of positive probability, the input rate will be larger than the output rate. \square

V. EPOCH-BASED POLICY—SUFFICIENCY

In this section, we consider a specific policy that we henceforth refer to as an *Epoch-Based* policy. The idea of the policy (which is defined formally below) is to divide the time into *epochs* and focus on the efficient evacuation of packets present in the system at the start of an epoch, while new packets that arrive during the epoch are excluded from processing. A similar idea was used to derive a stabilizing policy in [7] for a two-user broadcast erasure channel with feedback and in [9] for the input-queued switch. The main result of this section is that, for the special case of i.i.d. arrival processes, the epoch-based policy is stabilizing, provided that the underlying evacuation policy within each epoch is efficient (i.e., informally, minimizes the expected evacuation time for the packets present at the start of the epoch). More precisely, in this section we make the assumption that the arrival process vectors $\mathbf{A}(t)$ are i.i.d. with respect to time for $t = 1, 2, \dots$ (for a given time-slot t , the components of the vector $\mathbf{A}(t)$ may be dependent; also, the initial number of packets in the system at $t = 0$, namely $\mathbf{A}(0)$, can be arbitrary and is not required to have the same distribution as for $t \geq 1$). We then show that the epoch-based policy is stabilizing for any such arrival processes if the arrival rate λ satisfies $\hat{T}(\lambda) < 1$.

Consider the set

$$\mathcal{R}_l = \{\lambda \geq \mathbf{0} : \hat{T}(\lambda) < 1\}.$$

This set is nonempty since

$$\hat{T}(\mathbf{0}) = \lim_{t \rightarrow \infty} \frac{\bar{T}^*(t \cdot \mathbf{0})}{t} = 0 \quad (15)$$

hence $\mathbf{0} \in \mathcal{R}_l$. We will construct a policy that is stable for any $\lambda \in \mathcal{R}_l$. The continuity, convexity of $\hat{T}(\lambda)$ and (15) imply that the closure of \mathcal{R}_l is the set $\{\lambda \geq \mathbf{0} : \hat{T}(\lambda) \leq 1\}$ and hence

$$\{\lambda \geq \mathbf{0} : \hat{T}(\lambda) \leq 1\} \subseteq \mathcal{R}.$$

Combined with the necessity result of Section IV, we then conclude that

$$\mathcal{R} = \{\lambda \geq \mathbf{0} : \hat{T}(\lambda) \leq 1\}.$$

We now present a policy that stabilizes the system for any $\lambda \in \mathcal{R}_l$, that is

$$\hat{T}(\lambda) < 1. \quad (16)$$

Definition 13: Epoch-Based Policy π_ϵ . Pick $\epsilon > 0$ such that

$$0 < \epsilon < 1 - \hat{T}(\lambda)$$

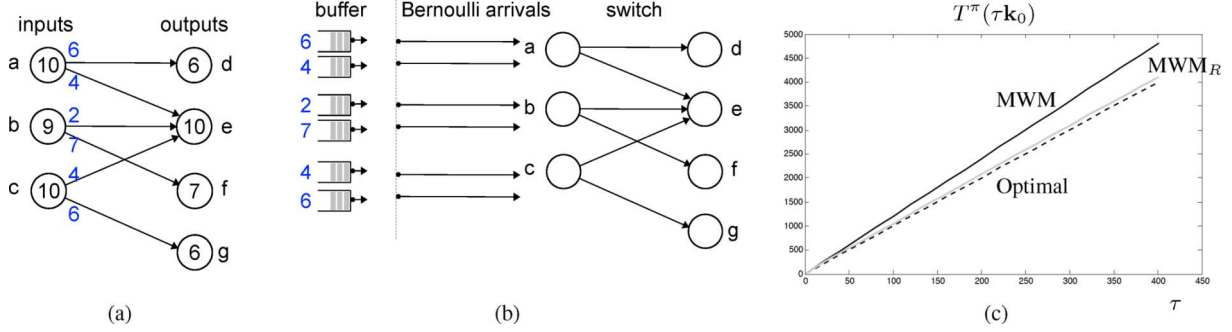


Fig. 2. (a) An evacuation example where MWM leads to suboptimal evacuation schedule; (b) The external buffer-based randomization technique of stabilizing policies; (c) Simulation results: a randomized version of MWM (called MWM_R) is asymptotically optimal.

and for each \mathbf{k} and s , pick an evacuation policy $\pi_{\mathbf{k},s}$ such that⁴

$$\bar{T}_s^{\pi_{\mathbf{k},s}}(\mathbf{k}) \leq \bar{T}_s^*(\mathbf{k}) + \epsilon. \quad (17)$$

Policy π_ϵ operates recursively in (random) time intervals $[t_{m-1}, t_m]$, $m = 1, 2, \dots$, called “epochs,” as follows. Epoch 1 starts at time $t_0 = 0$ at state $S_0 = s_0$ with $\hat{\mathbf{A}}(0) = \mathbf{A}(0) = \mathbf{k}$ packets at the inputs; the evacuation policy $\pi_{\mathbf{k},s}$ is used to evacuate the \mathbf{k} packets by time $t_1 = T_s^{\pi_{\mathbf{k},s}}(\mathbf{k})$, while any new packet arrivals during the epoch are kept at the inputs, but excluded from processing. Let S_m be the state of the system at time t_m . Epoch $m + 1$, $m \geq 1$ starts at time t_m with $\mathbf{k}_m = \mathbf{A}(t_m) - \hat{\mathbf{A}}(t_{m-1})$ packets at the inputs and policy $\pi_{\mathbf{k}_m, S_m}$ is used to evacuate the \mathbf{k}_m packets by time t_{m+1} . Note that due to F1–F3, policy π_ϵ is admissible.

Let $T_m = t_m - t_{m-1}$, $m = 1, 2, \dots$ be the length of the m th epoch. Since the arrival process vector is i.i.d, due to F2 and the Statistical Assumptions of Section II, the process $\{(T_m, S_m)\}_{m=1}^\infty$ constitutes a (homogeneous) Markov chain with stationary transition probabilities. Note that with this formulation, the initial state of the Markov chain, (T_1, S_1) , is a random variable whose distribution depends on $\mathbf{A}(0)$ and s_0 .

The main result of this section is the following.

Theorem 14: (Sufficiency). For any $\lambda \geq 0$ such that

$$\hat{T}(\lambda) < 1 \quad (18)$$

policy π_ϵ stabilizes the system. Hence, the stability region of Π is \mathcal{R} .

The proof of this theorem is given in Appendix D. We conjecture that this result can be extended to hold under more general arrival process statistics.

VI. FROM STABILIZING POLICIES TO POLICIES WITH ASYMPTOTICALLY OPTIMAL EVACUATION TIMES

In the preceding sections, we characterized the stability region of a class of policies through the asymptotic rates of minimal evacuation times and constructed a stabilizing policy based on policies having minimal evacuation times. In this section, we address the reverse problem: *If a stabilizing policy π for the given class of policies is known, is it possible to construct, based*

⁴The definition of the epoch-based policy involves an arbitrary small $\epsilon > 0$ since $\bar{T}_s^*(\mathbf{k})$ is defined as the *infimum* of evacuation times by all admissible policies; consequently a policy that actually achieves $\bar{T}_s^*(\mathbf{k})$ may not necessarily exist (or may not be admissible). Of course, if such a policy is admissible, then it can be used in the subsequent derivations.

on π , a minimum-evacuation time policy, at least in an asymptotic sense? To address this question, we first need the following definition.

Definition 15: Asymptotic Optimality. Policy π is *asymptotically optimal* (with regard to evacuation times) if for any $\mathbf{r} \geq 0$ it holds

$$\lim_{\tau \rightarrow \infty} \frac{\max_{s \in \mathcal{S}} \{\bar{T}_s^{\pi}(\lceil \mathbf{r}\tau \rceil)\}}{\tau} = \hat{T}(\mathbf{r}).$$

Note that the definition implies that

$$\lim_{\tau \rightarrow \infty} \frac{\max_{s \in \mathcal{S}} \{\bar{T}_s^{\pi}(\lceil \mathbf{r}\tau \rceil)\}}{\bar{T}^*(\lceil \mathbf{r}\tau \rceil)} = 1$$

hence the justification of asymptotic optimality. Also note that since \mathbf{r} represents the direction of growth, one may restrict attention to the case $\sum_{i=1}^n r_i = 1$.

Consider now that π is stabilizing. The next example shows that it is possible that $\bar{T}_s^{\pi}(\mathbf{k})$ may be neither optimal nor asymptotically optimal.

Example 16: Input Queued Switch. Consider the switch fabric shown in Fig. 2(a) with input set \mathcal{I} and output set \mathcal{O} , where the numbers on each link $i \rightarrow j$ denote the number of packets queued at input i for node j . Since this system only has a single state, we drop the state index in the notation. The policy that, in each slot, activates the matching that maximizes the sum of queue backlogs (weights) is called *maximum weight matching* policy (MWM). It is well known that MWM is stabilizing for this system [9]. Let $\mathbf{k} = (k_{ij})_{i \in \mathcal{I}, j \in \mathcal{O}}$ where k_{ij} is the number of packets at input i destined to output j . Now, consider the case where we are given the vector of packets \mathbf{k}_0 shown in Fig. 2(a) and we want to minimize the evacuation time for this \mathbf{k}_0 . From Hall’s theorem [10], it follows that the minimum evacuation time of \mathbf{k} is simply

$$T^*(\mathbf{k}) = \max \left\{ \max_j \sum_{i \in \mathcal{I}} k_{ij}, \max_i \sum_{j \in \mathcal{O}} k_{ij} \right\} \quad (19)$$

and in the present example $T^*(\mathbf{k}_0) = 10$. This optimal evacuation time can be achieved by finding a *critical matching* at each slot that reduces the maximum row/column sum of matrix \mathbf{k} by one [11]. Incidentally, a stabilizing policy based on the minimum evacuation time for the input switch is implemented [9], incurring a delay bound that is logarithmic in the switch size [12].

Now, check the operation of MWM on the particular evacuation example. The maximum matching at the first slot will be the one containing the links (a, d) , (b, f) , (c, g) with a weight

of $6 + 7 + 6 = 19$. However, note that this matching does not serve any of the packets destined to node e , thus, after the first slot, we still have $\sum_{i \in \mathcal{I}} x_{ie} = 10$, and by (19), we conclude that $T^{\text{MWM}}(\mathbf{k}_0) \geq 1 + T^*(\mathbf{k}_0)$. Thus, MWM is not an evacuation time optimal policy.

Next, assume that the vector of packets to be evacuated is scaled to $\tau \mathbf{k}_0$. We immediately get $T^*(\tau \mathbf{k}_0) = 10\tau$. Also, note that MWM will select the same matching as before (where no packet reaches e) for at least the first τ slots before the queue sizes drop for any other matching to have comparable weight. Applying the same reasoning, we conclude that after τ slots the packets at the input queues require at least another 10τ slots, i.e., $T^{\text{MWM}}(\tau \mathbf{k}_0) \geq \tau + T^*(\tau \mathbf{k}_0)$. Thus, by the definition above, MWM is not an asymptotically optimal policy in the evacuation time sense. \square

Notwithstanding the above example showing that a stabilizing policy is not in general asymptotically optimal for evacuation times, our main result of this section shows that an asymptotically optimal evacuation time can be achieved based on any given stabilizing policy, using a randomization technique that is presented next.

A. Randomized Version of Stable Policies

Consider a stabilizing policy π , whose stability region is \mathcal{R} . In addition to F1–F3 and SA1–SA5, we further assume that it is permissible to generate dummy packets in the system.

Let there be $\lceil \mathbf{r}\tau \rceil$ packets at the inputs of the system where $\mathbf{r} \neq \mathbf{0}$. We construct a randomized version of π , denoted by π_R^ϵ , which will be used in Theorem 17. Policy π_R^ϵ depends on a (very small) parameter ϵ , $0 < \epsilon < 1$.

- 1) If $\hat{T}(\mathbf{r}) > 0$, let $\alpha = 1/\hat{T}(\mathbf{r})$ so that, by positive homogeneity of $\hat{T}(\mathbf{r})$,

$$\hat{T}(\alpha \mathbf{r}) = \alpha \hat{T}(\mathbf{r}) = 1. \quad (20)$$

If $\hat{T}(\mathbf{r}) = 0$, then let $\alpha = 1/\epsilon$. In both cases, $\alpha \mathbf{r} \in \mathcal{R}$.

- 2) Since policy π is stabilizing and $\alpha \mathbf{r} \in \mathcal{R}$, there is an arrival rate vector λ_ϵ with $\sum_{i=1}^n |\alpha r_i - \lambda_{\epsilon,i}| \leq \epsilon$, such that the system is stable under π . Fix λ_ϵ and construct independent sequences of random variables, $\{A_i(t)\}_{t=0}^\infty$, $i = 1, \dots, n$, where for each i , $A_i(t)$ consists of i.i.d. random variables with

$$\mathbb{E}[A_i(t)] = \lambda_{\epsilon,i}. \quad (21)$$

- 3) Consider that the $\lceil \mathbf{r}\tau \rceil$ packets are placed in “buffers” outside the system [see Fig. 2(b) for the example of the input switch]. Mimic the actions of policy π , assuming that the packet (virtual) arrival process at time t is $\mathbf{A}(t)$, $t = 0, \dots, (\lceil \tau/\alpha \rceil)$. That is, at the beginning of time-slot t , pick $A_i(t)$ packets from the $\lceil r_i \tau \rceil$ that were originally in the buffers and consider them as “arrivals” to the system input queues. The processing of packets then follows policy π . If it happens that at some time $\tau < \lceil \tau/\alpha \rceil$ all $\lceil r_i \tau \rceil$ packets have departed, generate dummy packets to implement the policy. After time $\lceil \tau/\alpha \rceil$, if there are still \mathbf{m} (where $m_i > 0$ for some i) packets in the system, then pick policy π_h of Statistical Assumption SA4 to empty the system from the \mathbf{m} packets in $F^{\pi_h}(\mathbf{m})$ slots, where

$$\mathbb{E}[F^{\pi_h}(\mathbf{m}) \mid \mathcal{H}_t] \leq C \sum_{i=1}^n m_i. \quad (22)$$

Remark: If for policy π it holds for the time to empty the system from the \mathbf{m} packets, $F^\pi(\mathbf{m})$

$$\mathbb{E}[F^\pi(\mathbf{m}) \mid \mathcal{H}_t] \leq C^\pi \sum_{i=1}^n m_i \quad (23)$$

then we can use π to empty the system. The MWM policy for the input switch example has this property.

Theorem 17: Let π be a stabilizing policy. Then, for any $\mathbf{r} \geq \mathbf{0}$, it holds

$$\limsup_{t \rightarrow \infty} \max_s \left\{ \frac{\bar{T}_s^{\pi_R}(\lceil \mathbf{r}\tau \rceil)}{\tau} \right\} \leq (1 + C\epsilon)\hat{T}(\mathbf{r}) \quad \text{if } \hat{T}(\mathbf{r}) > 0$$

$$\limsup_{t \rightarrow \infty} \max_s \left\{ \frac{\bar{T}_s^{\pi_R}(\lceil \mathbf{r}\tau \rceil)}{\tau} \right\} \leq \epsilon + C\epsilon^2 \quad \text{if } \hat{T}(\mathbf{r}) = 0.$$

The proof is in Appendix E.

B. Simulations

We run a simulation experiment on the switch of Fig. 2(a) with $\mathbf{r} = \mathbf{k}_0$. In Fig. 2(c), the x -axis is τ and the y -axis is $T^\pi(\tau \mathbf{k}_0)$ for three different policies: 1) the optimal evacuation time policy described above; 2) MWM; and 3) the randomized version of MWM described in Theorem 17. Since MWM satisfies (23) in this case, we use MWM at both phases (the phase where virtual arrivals move packets from the buffer to the switch queues and the one where all the remaining packets are emptied). In this case

$$\hat{T}(\mathbf{k}) = \max \left\{ \max_j \sum_{i \in \mathcal{I}} k_{ij}, \max_i \sum_{j \in \mathcal{O}} k_{ij} \right\}$$

so we generate random variables

$$\mathbb{E}[A_{xy}(t)] \triangleq (1 - \epsilon)\alpha k_{xy} = \frac{(1 - \epsilon)k_{xy}}{\max \left\{ \sum_i k_{ij}, \sum_j k_{ij} \right\}} = p_{xy}.$$

Hence, it suffices to generate Bernoulli random variables with $\Pr(A_{xy}(t) = 1) = p_{xy}$.

From the simulations, it is evident that MWM can be arbitrarily suboptimal in terms of evacuation times (the gap $T^{\text{MWM}}(\tau \mathbf{k}_0) - T^*(\tau \mathbf{k}_0)$ diverges), but the randomized version of MWM (called MWM_R in the figure) asymptotically matches the rate of the optimal. Intuitively, the randomization of the arrivals provide a crucial regularity that is fundamentally important for the MWM policy to be efficient with respect to evacuation.

VII. CONCLUSION

This paper studied the connection between evacuation times and the stable throughput region under a general system model. Using this connection, if an optimal evacuation time policy is known, one immediately obtains a stabilizing policy for the system. Furthermore, if the optimal evacuation time can be computed in a tractable fashion, one obtains a closed-form expression of the system stability region. An application of this connection in the context of wireless networks with network coding and opportunistic routing can be found in [13] and [14], respectively. In addition, even in systems where both evacuation times and the system evolution are intractable, the connection provides a useful tool for deriving basic structural properties of the stability region. For example, this tool was used in [15] to prove the equivalence between the stability region and the

information-theoretic capacity in a broadcast channel with erasures and feedback.

APPENDIX A

PROOF OF SUBADDITIVITY (LEMMA 5)

Proof: Let $\epsilon > 0$ and let the system be in state s at time 0. An admissible policy π that evacuates $\mathbf{k} + \mathbf{m}$ packets is the following.

a) Pick an admissible policy $\pi_{\mathbf{k}}$ such that

$$\bar{T}_s^{\pi_{\mathbf{k}}}(\mathbf{k}) \leq \bar{T}_s^*(\mathbf{k}) + \epsilon/2.$$

b) Evacuate the \mathbf{k} packets following policy $\pi_{\mathbf{k}}$. According to Feature F3, this is permissible. From Statistical Assumption SA3, we conclude that the average evacuation time in this case is $\bar{T}_s^{\pi_{\mathbf{k}}}(\mathbf{k})$. Let s_1 be the state of the system by time $\bar{T}_s^{\pi_{\mathbf{k}}}(\mathbf{k})$. Both s_1 and $\bar{T}_s^{\pi_{\mathbf{k}}}(\mathbf{k})$ are known to $\pi_{\mathbf{k}}$ (hence to π) due to F1. Note that s_1 is a random variable that depends on s .

c) Again, pick an admissible policy $\pi_{\mathbf{m}}$ such that

$$\bar{T}_{s_1}^{\pi_{\mathbf{m}}}(\mathbf{m}) \leq \bar{T}_{s_1}^*(\mathbf{m}) + \epsilon/2.$$

d) Evacuate the \mathbf{m} packets following policy $\pi_{\mathbf{m}}$. Due to F2 and Statistical Assumptions SA2 and SA3, the average evacuation time (given s_1) in this case is $\bar{T}_{s_1}^{\pi_{\mathbf{m}}}(\mathbf{m})$.

The average evacuation time of π is

$$\begin{aligned} \bar{T}_s^{\pi}(\mathbf{k} + \mathbf{m}) &= \bar{T}_s^{\pi_{\mathbf{k}}}(\mathbf{k}) + \mathbb{E}[\bar{T}_{s_1}^{\pi_{\mathbf{m}}}(\mathbf{m})] \\ &\leq \bar{T}_s^*(\mathbf{k}) + \mathbb{E}[\bar{T}_{s_1}^*(\mathbf{m})] + \epsilon \end{aligned} \quad (24)$$

where the expectation in (24) is with respect to s_1 . Hence

$$\begin{aligned} \bar{T}^*(\mathbf{k} + \mathbf{m}) &= \max_{s \in S} \bar{T}_s^*(\mathbf{k} + \mathbf{m}) \stackrel{(3)}{\leq} \max_{s \in S} \bar{T}_s^{\pi}(\mathbf{k} + \mathbf{m}) \\ &\stackrel{(24)}{\leq} \max_{s \in S} \{ \bar{T}_s^*(\mathbf{k}) + \mathbb{E}[\bar{T}_{s_1}^*(\mathbf{m})] \} + \epsilon \\ &\leq \max_{s \in S} \bar{T}_s^*(\mathbf{k}) + \max_{s \in S} \bar{T}_s^*(\mathbf{m}) + \epsilon \end{aligned}$$

where the last inequality follows by applying expectations with respect to s_1 on $\bar{T}_{s_1}^* \leq \max_{s \in S} \bar{T}_s^*$ and plugging the result into the expression. Since ϵ is arbitrary, the lemma follows. ■

APPENDIX B

PROOF OF THEOREM 6

An analogy to Theorem 6 has been derived in [16] for subadditive functions defined on \mathbb{R}^n . The extension of Critical Evacuation Time Function to \mathbb{R}_0^n given in (9) is not necessarily subadditive, and hence we need different arguments to show the result.

Let $f(\mathbf{k}) : \mathbb{N}_0^n \rightarrow \mathbb{R}_0$ be a subadditive function. Let \mathcal{U} be the set of n -dimensional vectors whose coordinates are either zero or one, and define, $U \triangleq \max_{\mathbf{u} \in \mathcal{U}} f(\mathbf{u})$. We will need the following lemma.

Lemma 18: For any $\mathbf{k} \in \mathbb{N}_0^n - \{\mathbf{0}\}$, it holds

$$f(\mathbf{k}) \leq U \max_i k_i.$$

Proof: Assume without loss of generality that for some $c \leq n$, $0 < k_1 \leq k_2 \leq \dots \leq k_c$ and, in case $c < n$, then $k_{c+1} = \dots = k_n = 0$. Write

$$\mathbf{k} = [k_1 \dots k_n]^\top = \sum_{i=1}^c (k_i - k_{i-1}) \mathbf{u}_i$$

where $k_0 = 0$, and the j th element of \mathbf{u}_i is given by

$$u_{i,j} = \begin{cases} 0, & \text{if } i > 1 \quad \text{and} \quad j = 1, \dots, i-1 \\ 1, & \text{if } j = i, \dots, c \\ 0, & \text{if } j > c. \end{cases}$$

By subadditivity, we have

$$f(\mathbf{k}) \leq \sum_{i=1}^c (k_i - k_{i-1}) f(\mathbf{u}_i) \leq U k_c.$$

Next we extend the definition of $f(\mathbf{k})$ to \mathbb{R}_0^n by defining

$$f(\mathbf{r}) \triangleq f(\lceil \mathbf{r} \rceil), \quad \mathbf{r} \in \mathbb{R}_0^n.$$

We then have the following theorem.

Theorem 19: For any $\mathbf{r} \in \mathbb{R}_0^n$, the limit function

$$\hat{f}(\mathbf{r}) = \lim_{t \rightarrow \infty} \frac{f(t\mathbf{r})}{t} \quad (25)$$

exists, is finite, and positively homogeneous.

Proof: Assume without loss of generality that $r_1 \geq r_2 \geq \dots \geq r_n$. For consistency define $r_{n+1} = 0$. If $r_1 = 0$, then $\mathbf{r} = \mathbf{0}$ and (25) is obvious. Assume next that for some c , $1 \leq c \leq n$, $r_c > 0$ and $r_{c+1} = 0$.

Let $\epsilon > 0$ and $\beta = \liminf_{t \rightarrow \infty} f(t\mathbf{r})/t \geq 0$. Using Lemma 18, we have

$$\begin{aligned} \frac{f(t\mathbf{r})}{t} &= \frac{f(\lceil t\mathbf{r} \rceil)}{t} \leq U \frac{\max_i \lceil tr_i \rceil}{t} \\ &< U \frac{\max_i \{tr_i\} + 1}{t} = U \left(\max_i \{r_i\} + \frac{1}{t} \right). \end{aligned}$$

Hence, $\beta < \infty$.

To show existence of the limit in (25), it suffices to show that

$$\limsup_{t \rightarrow \infty} \frac{f(t\mathbf{r})}{t} \leq \beta + \delta(\epsilon) \quad (26)$$

where $\lim_{\epsilon \rightarrow 0} \delta(\epsilon) = 0$.

By definition of β , there are infinitely many t , such that $f(t\mathbf{r})/t \leq \beta + \epsilon$. Since we also have

$$r_i \leq \frac{\lceil tr_i \rceil}{t} < r_i + \frac{1}{t} \quad (27)$$

we can pick t_0 large enough so that the following inequalities hold.

$$\frac{f(t_0\mathbf{r})}{t_0} \leq \beta + \epsilon \quad (28)$$

$$r_i \leq \frac{\lceil t_0 r_i \rceil}{t_0} < r_i + \epsilon, \quad i = 1, \dots, c. \quad (29)$$

Using Euclidean division, write for $i = 1, \dots, c$

$$\lceil tr_i \rceil = l_{t,i} \lceil t_0 r_i \rceil + v_{t,i}, \quad 0 \leq v_{t,i} \leq \lceil t_0 r_i \rceil - 1. \quad (30)$$

If $c < n$, define also

$$l_{t,i} = v_{t,i} = 0, \quad i = c+1, \dots, n. \quad (31)$$

We then have

$$\begin{aligned} f(t\mathbf{r}) &= f(\lceil t\mathbf{r} \rceil) \\ &= f(l_{t,1} \lceil t_0 r_1 \rceil + v_{t,1}, \dots, l_{t,n} \lceil t_0 r_n \rceil + v_{t,n}) \\ &\leq f(l_{t,1} \lceil t_0 r_1 \rceil, \dots, l_{t,n} \lceil t_0 r_n \rceil) + f(\mathbf{v}_t). \end{aligned} \quad (32)$$

Next, write

$$[l_{t,1} \lceil t_0 r_1 \rceil \dots l_{t,n} \lceil t_0 r_n \rceil]^\top = \sum_{j=1}^c (l_{t,j} - l_{t,j-1}) \mathbf{v}_j$$

where $l_{t,0} = 0$ and the i th coordinate of \mathbf{v}_j , $v_{j,i}$, is defined for $1 \leq j \leq c$ as

$$v_{j,i} = \begin{cases} 0, & \text{if } j \neq 1 \text{ and } i = 1, \dots, j-1 \\ \lceil t_0 r_i \rceil, & \text{if } i = j, \dots, n. \end{cases} \quad (33)$$

Notice that since $r_j \geq r_{j+1}$, it holds $l_{t,j-1} \leq l_{t,j}$, $1 \leq j \leq c$. Using subadditivity, we then have from (32)

$$f(t\mathbf{r}) \leq \sum_{j=1}^c (l_{t,j} - l_{t,j-1})f(\mathbf{v}_j) + f(\mathbf{v}_t).$$

Hence

$$\begin{aligned} \frac{f(t\mathbf{r})}{t} &\leq \sum_{j=1}^c \frac{(l_{t,j} - l_{t,j-1})t_0}{t} \frac{f(\mathbf{v}_j)}{t_0} + \frac{f(\mathbf{v}_t)}{t} \\ &= \frac{l_{t,1}t_0}{t} \frac{f(t_0\mathbf{r})}{t_0} + \sum_{j=2}^c \frac{(l_{t,j} - l_{t,j-1})t_0}{t} \frac{f(\mathbf{v}_j)}{t_0} + \frac{f(\mathbf{v}_t)}{t}. \end{aligned} \quad (34)$$

By (30) and (31), \mathbf{v}_t takes a finite number of values, hence $f(\mathbf{v}_t)$ is a bounded sequence, and

$$\lim_{t \rightarrow \infty} \frac{f(\mathbf{v}_t)}{t} = 0.$$

Also, from (27), (29), and (30), we have for $1 \leq i \leq c$

$$\begin{aligned} r_i &\leq \frac{\lceil tr_i \rceil}{t} = \frac{l_{t,i}t_0}{t} \frac{\lceil t_0 r_i \rceil}{t_0} + \frac{v_{t,i}}{t} < \frac{l_{t,i}t_0}{t} (r_i + \epsilon) + \frac{v_{t,i}}{t} \\ r_i + \frac{1}{t} &> \frac{\lceil tr_i \rceil}{t} = \frac{l_{t,i}t_0}{t} \frac{\lceil t_0 r_i \rceil}{t_0} + \frac{v_{t,i}}{t} \geq \frac{l_{t,i}t_0}{t} r_i + \frac{v_{t,i}}{t} \end{aligned}$$

hence, using the fact that \mathbf{v}_t is a bounded sequence, we conclude

$$1 - \frac{\epsilon}{r_c} \leq 1 - \frac{\epsilon}{r_i} \leq \frac{r_i}{r_i + \epsilon} \leq \liminf_{t \rightarrow \infty} \frac{l_{t,i}t_0}{t} \leq \limsup_{t \rightarrow \infty} \frac{l_{t,i}t_0}{t} \leq 1.$$

Taking into account the latter inequalities and (28), we have from (34)

$$\begin{aligned} \limsup_{t \rightarrow \infty} \frac{f(t\mathbf{r})}{t} &\leq (\beta + \epsilon) \limsup_{t \rightarrow \infty} \frac{l_{t,1}t_0}{t} \\ &+ \sum_{j=2}^c \left(\limsup_{t \rightarrow \infty} \frac{l_{t,j}t_0}{t} - \liminf_{t \rightarrow \infty} \frac{l_{t,j-1}t_0}{t} \right) \frac{f(\mathbf{v}_j)}{t_0} \\ &\leq \beta + \epsilon + \frac{\epsilon}{r_c} \sum_{j=2}^c \frac{f(\mathbf{v}_j)}{t_0} \\ &\leq \beta + \epsilon + \frac{\epsilon}{r_c} U_c \frac{\max_i \lceil t_0 r_i \rceil}{t_0} \quad \text{by Lemma 18 and (33)} \\ &\leq \beta + \epsilon + \frac{\epsilon}{r_c} U_n (r_1 + \epsilon) \quad \text{by (29).} \end{aligned}$$

Hence, (26) holds with $\delta(\epsilon) = \epsilon + \frac{\epsilon}{r_c} U_n (r_1 + \epsilon)$.

Positive homogeneity follows immediately since for $\alpha \geq 0$

$$\hat{f}(\alpha\mathbf{r}) = \lim_{t \rightarrow \infty} \frac{f(t\alpha\mathbf{r})}{t} = \alpha \lim_{t \rightarrow \infty} \frac{f(t\alpha\mathbf{r})}{\alpha t} = \alpha \hat{f}(\mathbf{r}).$$

The next lemma is needed to establish further properties of $f(\mathbf{k})$ in Theorem 21.

Lemma 20: Let a subadditive function $f(\mathbf{k})$, $\mathbf{k} \in \mathbb{N}_0^n$ satisfy

$$f(\mathbf{k}) - f(\mathbf{k} + \mathbf{e}_i) \leq D_0 \quad \forall i = 1, \dots, n. \quad (35)$$

Then, for $D = \max\{f(\mathbf{e}_1), \dots, f(\mathbf{e}_n), D_0\}$, the following holds:

$$|f(\mathbf{k}) - f(\mathbf{k} + \mathbf{e}_i)| \leq D, \quad \text{for all } i = 1, \dots, n. \quad (36)$$

$$|f(\mathbf{k}) - f(\mathbf{m})| \leq D \sum_{i=1}^n |k_i - m_i| \quad (37)$$

$$|f(\mathbf{r}) - f(\mathbf{s})| < D \sum_{i=1}^n |r_i - s_i| + nD. \quad (38)$$

Proof: By subadditivity

$$f(\mathbf{k} + \mathbf{e}_i) \leq f(\mathbf{k}) + f(\mathbf{e}_i)$$

hence

$$f(\mathbf{k} + \mathbf{e}_i) - f(\mathbf{k}) \leq \max_i f(\mathbf{e}_i) \doteq D_1.$$

Taking into account (35), we conclude

$$|f(\mathbf{k} + \mathbf{e}_i) - f(\mathbf{k})| \leq \max\{D_1, D_0\} \doteq D$$

which shows (36).

To show (37), we use backward induction on the number c of coordinates of \mathbf{k}, \mathbf{m} that are equal. If $c = n$, then clearly (37) holds. Let (37) hold for $c \leq n$ and assume without loss of generality that $k_i = m_i, i = 1, \dots, c-1$ and $k_i \neq m_i, i \geq c, k_c > m_c$. We then have

$$\begin{aligned} |f(\mathbf{k}) - f(\mathbf{m})| &= |f(\mathbf{k}) - f(k_1, \dots, k_{c-1}, m_c, k_{c+1}, \dots, k_n) \\ &\quad + f(k_1, \dots, k_{c-1}, m_c, k_{c+1}, \dots, k_n) - f(\mathbf{m})| \\ &\leq |f(\mathbf{k}) - f(k_1, \dots, k_{c-1}, m_c, k_{c+1}, \dots, k_n)| \\ &\quad + |f(m_1, \dots, m_{c-1}, m_c, k_{c+1}, \dots, k_n) - f(\mathbf{m})| \\ &\leq |f(\mathbf{k}) - f(k_1, \dots, k_{c-1}, m_c, k_{c+1}, \dots, k_n)| \\ &\quad + D \sum_{i=c+1}^n |k_i - m_i| \text{ by the ind. hypothesis.} \end{aligned}$$

Now, write

$$\begin{aligned} &|f(\mathbf{k}) - f(k_1, \dots, k_{c-1}, m_c, k_{c+1}, \dots, k_n)| \\ &= \left| \sum_{i=0}^{k_c - m_c - 1} f(k_1, \dots, k_{c-1}, m_c + i + 1, k_{c+1}, \dots, k_n) \right. \\ &\quad \left. - f(k_1, \dots, k_{c-1}, m_c, k_{c+1}, \dots, k_n) \right| \\ &\leq \sum_{i=0}^{k_c - m_c - 1} |f(k_1, \dots, k_{c-1}, m_c + i + 1, k_{c+1}, \dots, k_n) \\ &\quad - f(k_1, \dots, k_{c-1}, m_c, k_{c+1}, \dots, k_n)| \\ &\stackrel{\text{by (36)}}{\leq} \sum_{i=0}^{k_c - m_c - 1} D = D|k_c - m_c| \end{aligned}$$

and since $k_i = m_i, i = 1, \dots, c-1$, we have

$$|f(\mathbf{k}) - f(\mathbf{m})| \leq D \sum_{i=c}^n |k_i - m_i| = D \sum_{i=1}^n |k_i - m_i|$$

i.e., the inductive hypothesis holds for $c-1$ as well.

Finally, for (38), write

$$\begin{aligned} |f(\mathbf{r}) - f(\mathbf{s})| &= |f(\lceil \mathbf{r} \rceil) - f(\lceil \mathbf{s} \rceil)| \\ &\leq D \sum_{i=1}^n |\lceil r_i \rceil - \lceil s_i \rceil| \text{ by (37)} \\ &< D \sum_{i=1}^n |r_i - s_i| + Dn \end{aligned}$$

where the last inequality follows from $||\lceil r_i \rceil - \lceil s_i \rceil| < |r_i - s_i| + 1$. ■

The next theorem provides further useful properties of $\hat{f}(\mathbf{r})$ under condition (35).

Theorem 21: If a subadditive function $f(\mathbf{k})$, $\mathbf{k} \in \mathbb{N}_0^n$ satisfies (35), then the limit function

$$\hat{f}(\mathbf{r}) = \lim_{t \rightarrow \infty} \frac{f(t\mathbf{r})}{t}$$

is subadditive, convex, Lipschitz continuous, i.e., it holds

$$|\hat{f}(\mathbf{r}) - \hat{f}(\mathbf{s})| \leq D \sum_{i=1}^n |r_i - s_i|$$

and for any sequence $\mathbf{r}_t \in \mathbb{R}_0^n$ such that

$$\lim_{t \rightarrow \infty} \mathbf{r}_t = \boldsymbol{\lambda} < \infty$$

it holds

$$\lim_{t \rightarrow \infty} \frac{f(t\mathbf{r}_t)}{t} = \hat{f}(\boldsymbol{\lambda}). \quad (39)$$

Proof: To show subadditivity, we proceed as follows. Since for any a, b it holds

$$\lceil a + b \rceil + x = \lceil a \rceil + \lceil b \rceil \quad \text{for some } x = 0, 1, 2$$

we write

$$\lceil t(\mathbf{r}_1 + \mathbf{r}_2) \rceil + \mathbf{x} = \lceil t\mathbf{r}_1 \rceil + \lceil t\mathbf{r}_2 \rceil.$$

Also, by (37)

$$f(\lceil t(\mathbf{r}_1 + \mathbf{r}_2) \rceil) - f(\lceil t(\mathbf{r}_1 + \mathbf{r}_2) \rceil + \mathbf{x}) \leq D \sum_{i=1}^n x_i \leq 2nD.$$

Hence

$$\begin{aligned} f(t(\mathbf{r}_1 + \mathbf{r}_2)) - 2nD &\leq f(\lceil t(\mathbf{r}_1 + \mathbf{r}_2) \rceil + \mathbf{x}) \\ &= f(\lceil t\mathbf{r}_1 \rceil + \lceil t\mathbf{r}_2 \rceil) + f(\mathbf{x}) \\ &\leq f(t\mathbf{r}_1) + f(t\mathbf{r}_2) + f(\mathbf{x}). \end{aligned}$$

Dividing the last inequality by t and taking limits shows that $\hat{f}(\mathbf{r}_1 + \mathbf{r}_2) \leq \hat{f}(\mathbf{r}_1) + \hat{f}(\mathbf{r}_2)$.

Convexity follows easily from positive homogeneity and subadditivity

$$\begin{aligned} \hat{f}(p\mathbf{r}_1 + (1-p)\mathbf{r}_2) &\leq \hat{f}(p\mathbf{r}_1) + \hat{f}((1-p)\mathbf{r}_2) \\ &= p\hat{f}(\mathbf{r}_1) + (1-p)\hat{f}(\mathbf{r}_2). \end{aligned}$$

Lipschitz continuity follows easily as well from (38) by replacing \mathbf{r}, \mathbf{s} with $t\mathbf{r}, t\mathbf{s}$, dividing by t and taking limits.

Finally, let

$$\lim_{t \rightarrow \infty} \mathbf{r}_t = \boldsymbol{\lambda} < \infty.$$

Using (38), write

$$\begin{aligned} \left| \frac{f(t\mathbf{r}_t)}{t} - \hat{f}(\boldsymbol{\lambda}) \right| &= \left| \frac{f(t\mathbf{r}_t)}{t} - \frac{f(t\boldsymbol{\lambda})}{t} + \frac{f(t\boldsymbol{\lambda})}{t} - \hat{f}(\boldsymbol{\lambda}) \right| \\ &\leq \left| \frac{f(t\mathbf{r}_t) - f(t\boldsymbol{\lambda})}{t} \right| + \left| \frac{f(t\boldsymbol{\lambda})}{t} - \hat{f}(\boldsymbol{\lambda}) \right| \\ &\leq D \sum_{i=1}^n |r_{t,i} - \lambda_i| + \frac{nD}{t} + \left| \frac{f(t\boldsymbol{\lambda})}{t} - \hat{f}(\boldsymbol{\lambda}) \right|. \end{aligned}$$

Taking limits in the last inequality shows (39). ■

Theorem 6 will follow directly from Theorems 19 and 21 if we verify that the critical evacuation time function satisfies (35). However, this follows easily from (6) since

$$\begin{aligned} \bar{T}^*(\mathbf{k}) - \bar{T}^*(\mathbf{k} + \mathbf{e}_i) &= \max_s \bar{T}_s^*(\mathbf{k}) - \max_i \bar{T}_s^*(\mathbf{k} + \mathbf{e}_i) \\ &\leq \max_s \{ \bar{T}_s^*(\mathbf{k}) - \bar{T}_s^*(\mathbf{k} + \mathbf{e}_i) \} \\ &\leq D_0, \quad \text{by (6).} \end{aligned}$$

APPENDIX C

PROOF OF NECESSITY (THEOREM 8)

We need the following lemma.

Lemma 22: (System Stability Implies Mean Rate Stability). If (13), (1), and (2) hold, then

$$\lim_{t \rightarrow \infty} \frac{\mathbb{E}[Q_s^\pi(t)]}{t} = 0. \quad (40)$$

Proof of Lemma 22: It follows from (1), (2) and the corollary to [17, Theorem 16.14] that the sequences $\{\tilde{A}_i(t)/t\}$ $i = 1, \dots, n$ are uniformly integrable, hence the sequence $\{\sum_{i=1}^n \tilde{A}_i(t)/t\}$ is also uniformly integrable. Since

$$0 \leq \frac{Q_s^\pi(t)}{t} \leq \frac{\sum_{i=1}^n \tilde{A}_i(t)}{t}$$

we conclude that the sequence $\{Q_s^\pi(t)/t\}$ is also uniformly integrable. We will show in the next paragraph that $\{Q_s^\pi(t)/t\}$ converges in probability to 0. Equality (40) will then follow from [17, Theorem 25.12].

Pick any $\epsilon > 0$ (arbitrarily small) and a $q \geq 0$ large enough so that according to (13) it holds

$$\limsup_{t \rightarrow \infty} \Pr\{Q_s^\pi(t) > q\} \leq \epsilon.$$

Since we can pick t_0 large enough so that $\epsilon t > q$, $t \geq t_0$, we have

$$\limsup_{t \rightarrow \infty} \Pr\left\{\frac{Q_s^\pi(t)}{t} > \epsilon\right\} \leq \limsup_{t \rightarrow \infty} \Pr\{Q_s^\pi(t) > q\} \leq \epsilon$$

i.e., $\{Q_s^\pi(t)/t\}$ converges in probability to 0. ■

Proof of Theorem 8: Pick $\mathbf{r} \in \mathcal{R}$. Since \mathbf{r} belongs to the closure of the rates for which the system is stabilizable, for any $\delta > 0$ we can find a $\boldsymbol{\lambda} \geq \mathbf{0}$, $\|\boldsymbol{\lambda} - \mathbf{r}\| \leq \delta$, for which the system is stable under some policy $\pi_0 \in \Pi$. By continuity of $\hat{T}(\mathbf{r})$ it suffices to show that for any such $\boldsymbol{\lambda}$

$$\hat{T}(\boldsymbol{\lambda}) \leq 1. \quad (41)$$

Let the initial system state be $s \in \mathcal{S}$. Fix an arbitrary time index t and generate random number of packets $\mathbf{A}(0), \dots, \mathbf{A}(t)$ according to the distribution of the arrival processes. Consider that all $\tilde{\mathbf{A}}(t) = \sum_{\tau=0}^t \mathbf{A}(\tau)$ packets are in the system at the beginning of time and construct the following evacuation policy π .

1) Mimic the actions of policy π_0 for up to t time-slots, assuming that the packet arrival process at time τ is $\mathbf{A}(\tau)$, $\tau = 1, \dots, t$. Due to Features F1, F3 this mimicking is permissible.

2) If all $\tilde{\mathbf{A}}(t)$ packets are transmitted by time t , then the evacuation time of π is at most t . Else, after t time-slots, there will be $Q_s^{\pi_0}(t) > 0$ packets in the system. According to Statistical Assumption SA4, pick a policy π_h to evacuate the $Q_s^{\pi_0}(t)$ packets in $F^{\pi_h}(Q_s^{\pi_0}(t))$ slots, where using (5)

$$\mathbb{E}\left[F^{\pi_h}(Q_s^{\pi_0}(t)) \mid \tilde{\mathbf{A}}(t), \tilde{D}_s^{\pi_0}(t)\right] \leq C Q_s^{\pi_0}(t). \quad (42)$$

The evacuation time of π given $\tilde{\mathbf{A}}(t)$ is at most $t + F^{\pi_h}(\mathbf{Q}_s^{\pi_0}(t))$ —“at most” because all $\tilde{\mathbf{A}}(t)$ packets may have left before time t —and hence, taking the conditional average, we have

$$\begin{aligned}\bar{T}_s^*(\tilde{\mathbf{A}}(t)) &\leq t + \mathbb{E} \left[F^{\pi_h}(\mathbf{Q}_s^{\pi_0}(t)) \mid \tilde{\mathbf{A}}(t) \right] \\ &= t + \mathbb{E} \left[\mathbb{E} \left[F^{\pi_h}(\mathbf{Q}_s^{\pi_0}(t)) \mid \tilde{\mathbf{A}}(t), \tilde{\mathbf{D}}_s^{\pi_0}(t) \right] \mid \tilde{\mathbf{A}}(t) \right] \\ &\leq t + C \mathbb{E} \left[Q_s^{\pi_0}(t) \mid \tilde{\mathbf{A}}(t) \right] \quad \text{by (42)}.\end{aligned}$$

Next, using the last inequality

$$\begin{aligned}\bar{T}^*(\tilde{\mathbf{A}}(t)) &= \max_{s \in \mathcal{S}} \left\{ \bar{T}_s^*(\tilde{\mathbf{A}}(t)) \right\} \\ &\leq t + C \max_{s \in \mathcal{S}} \mathbb{E} \left[Q_s^{\pi_0}(t) \mid \tilde{\mathbf{A}}(t) \right] \\ &\leq t + C \sum_{s \in \mathcal{S}} \mathbb{E} \left[Q_s^{\pi_0}(t) \mid \tilde{\mathbf{A}}(t) \right].\end{aligned}$$

Taking expectations with respect to $\tilde{\mathbf{A}}(t)$ and dividing by t , we have

$$\mathbb{E} \left[\frac{\bar{T}^* \left(\frac{\tilde{\mathbf{A}}(t)}{t} \right)}{t} \right] \leq 1 + C \sum_{s \in \mathcal{S}} \frac{\mathbb{E} [Q_s^{\pi_0}(t)]}{t}. \quad (43)$$

Since $\lim_{t \rightarrow \infty} \frac{\tilde{\mathbf{A}}(t)}{t} = \lambda$, by (1), using (11) from Theorem 6, we then obtain,

$$\lim_{t \rightarrow \infty} \frac{\bar{T}^* \left(\frac{\tilde{\mathbf{A}}(t)}{t} \right)}{t} = \hat{T}(\lambda).$$

Hence

$$\begin{aligned}\hat{T}(\lambda) &= \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{\bar{T}^* \left(\frac{\tilde{\mathbf{A}}(t)}{t} \right)}{t} \right] \\ &\leq \liminf_{t \rightarrow \infty} \mathbb{E} \left[\frac{\bar{T}^* \left(\frac{\tilde{\mathbf{A}}(t)}{t} \right)}{t} \right] \quad \text{by Fatou's lemma} \\ &\leq \liminf_{t \rightarrow \infty} \left(1 + C \sum_{s \in \mathcal{S}} \frac{\mathbb{E} [Q_s^{\pi_0}(t)]}{t} \right) \quad \text{by (43)} \\ &= 1 + C \sum_{s \in \mathcal{S}} \lim_{t \rightarrow \infty} \frac{\mathbb{E} [Q_s^{\pi_0}(t)]}{t} = 1 \quad \text{by (40)}.\end{aligned}$$

APPENDIX D PROOF OF SUFFICIENCY (THEOREM 14)

In the discussion that follows, we use the terminology and related results in [18]. For $x, y \in \mathcal{G}$, if x leads to y , we write $x \rightsquigarrow y$ and if x communicates with y , $x \longleftrightarrow y$. A Markov chain with countable state space \mathcal{G} is called irreducible if all states in \mathcal{G} belong to the same *essential* class, i.e., all states communicate with each other.

The proof of stability of the Epoch-Based Policy is based on the following theorem (see [19] and [20]).

Theorem 23: Let $\{X_m\}_{m=1}^{\infty}$ be a homogeneous, irreducible, and aperiodic Markov chain with countable state space \mathcal{G} . Let $v(x)$ be a nonnegative real function defined on the state space (Lyapunov function). If there exists a finite set $\mathcal{A} \subseteq \mathcal{G}$ such that $v(x) \geq \epsilon > 0$, $x \in \mathcal{A}^c = \mathcal{G} - \mathcal{A}$

$$\mathbb{E}[v(X_2) \mid X_1 = x] < \infty, \quad x \in \mathcal{A} \quad (44)$$

and for some $\delta, 1 \geq \delta > 0$

$$\mathbb{E}[v(X_2) \mid X_1 = x] \leq (1 - \delta) v(x), \quad x \in \mathcal{A}^c \quad (45)$$

then the Markov Chain is geometrically ergodic (positive recurrent) and $\mathbb{E}[v(\hat{X})] < \infty$, where \hat{X} has the steady-state distribution of $\{X_m\}_{m=1}^{\infty}$.

For the general model under consideration in the current work, irreducibility and aperiodicity may not hold. Hence, we need some preparatory work to use Theorem 23. The following lemma will be useful.

Lemma 24: Let $\{X_m\}_{m=1}^{\infty}$ be a homogeneous Markov Chain, not necessarily irreducible and/or aperiodic.

a) With the notation of Theorem 23, conditions (44) and (45) imply

$$\mathbb{E}[v(X_2) \mid X_1 = x] \leq U + (1 - \delta)v(x) \quad \text{for all } x \in \mathcal{G} \quad (46)$$

where $U = \max_{x \in \mathcal{A}} \mathbb{E}[v(X_2) \mid X_1 = x]$.

b) Conversely, if $v(x) \geq 0$ and there are constants $U > 0$, $\delta, \delta_1, 0 < \delta_1 < \delta \leq 1$, and a finite set \mathcal{B} such that (46) holds and

$$\frac{U}{v(x)} \leq \delta_1 \quad \text{for all } x \in \mathcal{B}^c \quad (47)$$

then (44) and (45) hold with $\mathcal{A} \leftarrow \mathcal{B}$ and $\delta \leftarrow \delta - \delta_1$.

c) If (46) holds, then for $m \geq 2$

$$\mathbb{E}[v(X_m) \mid X_1 = x] \leq \frac{U}{\delta} + (1 - \delta)^{m-1} v(x). \quad (48)$$

Proof: It is clear that (44) and (45) imply (46). Assume now that (46) and (47) hold. Then, clearly (44) is satisfied for all $x \in \mathcal{B}$. Also, since the following holds for $x \in \mathcal{B}^c$

$$\begin{aligned}\mathbb{E}[v(X_2) \mid X_1 = x] &\leq \left(1 - \left(\delta - \frac{U}{v(x)} \right) \right) v(x) \\ &\leq (1 - (\delta - \delta_1)) v(x)\end{aligned}$$

it follows that (45) is satisfied for $x \in \mathcal{B}^c$ with $\delta \leftarrow \delta - \delta_1$.

To prove (48), use the Markov property and (46) to write

$$\begin{aligned}\mathbb{E}[v(X_m) \mid X_1 = x] &= \mathbb{E}[\mathbb{E}[v(X_m) \mid X_{m-1}, X_1 = x]] \\ &\leq U + (1 - \delta) \mathbb{E}[v(X_{m-1}) \mid X_1 = x]\end{aligned}$$

and hence by induction

$$\begin{aligned}\mathbb{E}[v(X_m) \mid X_1 = x] &\leq U \sum_{i=0}^{m-2} (1 - \delta)^i + (1 - \delta)^{m-1} v(x) \\ &\leq \frac{U}{\delta} + (1 - \delta)^{m-1} v(x).\end{aligned}$$

The next lemma states that when (16) holds, the Markov process described in Section V, namely $\{(T_m, S_m)\}_{m=1}^{\infty}$ (where T_m is the duration of the m th epoch and S_m is the system state at the end of the m th epoch) has the drift property described in Lemma 24.

Lemma 25: For the Markov process $\{(T_m, S_m)\}_{m=1}^\infty$, define $v((\tau, s)) = \tau$. If $\hat{T}(\lambda) < 1$, then there are $U > 0$ and $\delta > 0$ such that

$$\mathbb{E}[v((T_2, S_2)) \mid (T_1, S_1) = (\tau, s)] \leq U + (1 - \delta) v((\tau, s))$$

for all $(\tau, s) \in \mathcal{G}$ and (47) is also satisfied.

Proof: Using the definition of v , and the fact that given \mathbf{k}_1 and S_1 , T_2 is independent of T_1 , write

$$\begin{aligned} \mathbb{E}[v((T_2, S_2)) \mid (T_1, S_1) = (\tau, s)] &= \mathbb{E}[T_2 \mid T_1 = \tau, s_1 = s] \\ &= \mathbb{E}[\mathbb{E}[T_2 \mid T_1 = \tau, s_1 = s, \mathbf{k}_1(\tau)] \mid (T_1, S_1) = (\tau, s)] \\ &= \mathbb{E}[\bar{T}_s^{\pi_{\mathbf{k}_1, s}}(\mathbf{k}_1(\tau))] \end{aligned} \quad (49)$$

We have by construction of π_ϵ

$$\begin{aligned} \mathbb{E}[\bar{T}_s^{\pi_{\mathbf{k}_1, s}}(\mathbf{k}_1(\tau))] &\leq \mathbb{E}[\bar{T}_s^*(\mathbf{k}_1(\tau))] + \epsilon \\ &\leq \mathbb{E}[\bar{T}^*(\mathbf{k}_1(\tau))] + \epsilon. \end{aligned} \quad (50)$$

Since the arrival process vectors are i.i.d, it holds with probability 1

$$\lim_{\tau \rightarrow \infty} \frac{\mathbf{k}_1(\tau)}{\tau} = \lambda$$

and

$$\begin{aligned} \lim_{\tau \rightarrow \infty} \frac{\bar{T}^*(\mathbf{k}_1(\tau))}{\tau} &= \lim_{\tau \rightarrow \infty} \frac{\bar{T}^*\left(\frac{\mathbf{k}_1(\tau)}{\tau}\tau\right)}{\tau} \\ &= \hat{T}(\lambda) \quad \text{by (11)}. \end{aligned} \quad (51)$$

We will show at the end of the proof that the sequence $\bar{T}^*(\mathbf{k}_1(\tau))/\tau, \tau = 1, \dots$ is uniformly integrable, which will imply that

$$\begin{aligned} \limsup_{\tau \rightarrow \infty} \frac{\mathbb{E}[T_2 \mid T_1 = \tau, s_n = s]}{\tau} &\leq \lim_{\tau \rightarrow \infty} \mathbb{E}\left[\frac{\bar{T}^*(\mathbf{k}_1(\tau))}{\tau}\right] + \epsilon \quad \text{by (49) and (50)} \\ &= \mathbb{E}\left[\lim_{\tau \rightarrow \infty} \frac{\bar{T}^*(\mathbf{k}_1(\tau))}{\tau}\right] + \epsilon \quad \text{by uniform integrability} \\ &= \hat{T}(\lambda) + \epsilon \quad \text{by (51)}. \end{aligned}$$

Therefore, for δ such that $0 < \delta < 1 - \hat{T}(\lambda) - \epsilon$, there exists τ_δ such that for all pairs (τ, s) with $\tau > \tau_\delta$, it holds

$$\mathbb{E}[T_2 \mid T_1 = \tau, S_1 = s] \leq (1 - \delta)\tau$$

hence

$$\mathbb{E}[T_2 \mid T_1 = \tau, S_1 = s] \leq U + (1 - \delta)\tau$$

where

$$U = \max_{(\tau, s) \in \mathcal{G}: \tau \leq \tau_\delta} \mathbb{E}[T_2 \mid T_1 = \tau, S_1 = s].$$

Also, (47) is satisfied since $\lim_{\tau \rightarrow \infty} v(\tau) = \tau = \infty$.

It remains to show that $\bar{T}^*(\mathbf{k}_1(\tau))/\tau, \tau = 1, 2, \dots$ is uniformly integrable. Using (5), we have

$$0 \leq \frac{\bar{T}^*(\mathbf{k}_1(\tau))}{\tau} \leq C \sum_{i=1}^n \frac{k_{1,i}(\tau)}{\tau}. \quad (52)$$

Now, we have with probability one

$$\lim_{\tau \rightarrow \infty} \frac{k_{1,i}(\tau)}{\tau} = \lambda_i.$$

On the other hand, since the length of an epoch T_1 is independent of the arrivals during this epoch, we have

$$\frac{\mathbb{E}[k_{1,i}(\tau)]}{\tau} = \frac{\lambda_i \tau}{\tau} = \lambda_i.$$

Since the nonnegative sequences $\frac{k_{1,i}(\tau)}{\tau}, i = 1, 2, \dots, n, \tau = 1, 2, \dots$ converge both with probability one and in expectation, they are uniformly integrable (see [17, Theorem 16.4]). Using this fact, uniform integrability of $\bar{T}^*(\mathbf{k}_1(\tau))/\tau$ follows from (52). ■

We next present the main theorem of this section, which shows the stability of policy π_ϵ .

Theorem 26: For any $\lambda \geq 0$ such that

$$\hat{T}(\lambda) < 1 \quad (53)$$

policy π_ϵ stabilizes the system.

Proof: The idea of the proof is the following. Assume that the system starts at time $t = 0$ in system state s , with $\mathbf{A}(0) = \mathbf{k}$ packets at the inputs. We use the queue occupancy notation of $Q_s^\pi(t), Q_s^\pi(t)$ from Section IV, but we henceforth omit the indices s and π to simplify the notation. Under π_ϵ , it will be shown through Theorem 23 that (53) implies that we can identify a state (τ_a, s_a) to which the chain $(T_m, S_m)_{m=1}^\infty$ returns infinitely often. Define $m_l, l = 1, \dots$ to be the sequence of epoch indices when the Markov chain is in state (τ_a, s_a) . Then, due to the Markov property, the process consisting of the successive intervals between the times at which the process $(T_m, S_m)_{m=1}^\infty$ returns to state (τ_a, s_a) , i.e.,

$$L_l = \sum_{j=m_l+1}^{m_{l+1}} T_j, \quad l = 1, 2, \dots \quad (54)$$

consists of i.i.d. random variables and, as will be seen

$$\mathbb{E}[L_l] < \infty. \quad (55)$$

Hence, the process

$$Z_0 = \sum_{j=1}^{m_1} T_j, \quad Z_l = Z_{l-1} + L_l, \quad l \geq 1$$

constitutes a (delayed) renewal process.

Observe next that by the operation of π_ϵ , $\{Q(Z_l)\}_{l=0}^\infty$, the number of packets in system at times Z_l , is statistically the same as the number of arrivals in a interval of length τ_a . Since packet arrivals are i.i.d. and the operations of the process during the interval τ_a do not depend on these arrivals, $\{Q(Z_l)\}_{l=0}^\infty$, consists of i.i.d. random variables with $\mathbb{E}[Q(Z_l)] = \lambda \tau_a < \infty$. This, and the operation of π_ϵ imply that the process $\{Q(t)\}_{t=0}^\infty$, is regenerative with respect to $\{Z_l\}_{l=0}^\infty$. Let g be the period of the distribution of the cycle length L_l . It then follows from [21, p. 128, Corollary 1.5] and (55) that

$$\begin{aligned} \lim_{\alpha \rightarrow \infty} \frac{1}{g} \sum_{\beta=0}^{g-1} \Pr(Q(\alpha g + \beta) > q) &= \lim_{\alpha \rightarrow \infty} \frac{1}{g} \sum_{\beta=0}^{g-1} \mathbb{E}[1_{\{Q(\alpha g + \beta) > q\}}] = \frac{\mathbb{E}\left[\sum_{j=0}^{L_1-1} 1_{\{Q(Z_0+j) > q\}}\right]}{\mathbb{E}[L_1]}. \end{aligned} \quad (56)$$

Observe next that the random variables $Y_j(q) = 1_{\{Q(Z_0+j) > q\}}$ are decreasing in q , and since $Q(Z_0 + j)$ are finite, $\lim_{q \rightarrow \infty} 1_{\{Q(Z_0+j) > q\}} = 0$. Using the monotone convergence theorem, we then have

$$\begin{aligned} \lim_{q \rightarrow \infty} \mathbb{E}\left[\sum_{j=0}^{L_1-1} 1_{\{Q(Z_0+j) > q\}}\right] &= \mathbb{E}\left[\lim_{q \rightarrow \infty} \sum_{j=0}^{L_1-1} 1_{\{Q(Z_0+j) > q\}}\right] \\ &= 0. \end{aligned} \quad (57)$$

From (56) and (57), we conclude that

$$\lim_{q \rightarrow \infty} \lim_{\alpha \rightarrow \infty} \sum_{\beta=0}^{g-1} \Pr(Q(\alpha g + \beta) > q) = 0. \quad (58)$$

Since for $t = \alpha_t g + \beta_t$ it holds

$$\Pr(Q(t) > q) \leq \sum_{\beta=0}^{g-1} \Pr(Q(\alpha_t g + \beta) > q)$$

we conclude from (58) that

$$\lim_{q \rightarrow \infty} \limsup_{t \rightarrow \infty} \Pr(Q(t) > q) = 0$$

i.e., policy π_ϵ is stable.

To implement the plan outlined above, we must show the existence of a state to which the Markov chain returns infinitely often, as well as (58). For this, we will use Theorem 23, but because of the generality of the model under consideration, we cannot *a priori* claim irreducibility and aperiodicity in order to apply it directly. Instead, we rely first on Lemma 24 using the result of Lemma 25.

Let $C((\tau, s))$ be the communicating class to which a state (τ, s) belongs. We consider two cases as follows.

a) If $C((\tau, s))$ is essential, and $(T_1, S_1) = (\tau, s)$, we have $(T_m, S_m) \in C((\tau, s))$ $m = 1, 2, \dots$ and the evolution of the process with initial condition (τ, s) constitutes an irreducible Markov chain. If this chain is periodic with period d , then the process $\{(T_{dk+1}, S_{dk+1})\}_{k=0}^\infty$ is an aperiodic Markov chain [18, p. 14]. For this chain, we can apply Theorem 23 to show positive recurrence, as follows. Since by Lemma 25 the process $\{(T_m, S_m)\}_{m=1}^\infty$ satisfies (46), it also satisfies (48). Hence, the process $\{(T_{dk+1}, S_{dk+1})\}_{k=0}^\infty$ satisfies (46), and since $\lim_{\tau \rightarrow \infty} v((\tau, s)) = \infty$, it also satisfies (47). Therefore, by Lemma 24, we can apply Theorem 23 to $\{(T_{dk+1}, S_{dk+1})\}_{k=0}^\infty$ to deduce that it is geometrically ergodic with

$$\mathbb{E}[v(\hat{X})] = \mathbb{E}[\hat{T}] < \infty. \quad (59)$$

From the above discussion, we conclude that the initial state (τ, s) is visited infinitely often, and the successive visit indices are of the form $m_l = dV_l + 1$, $l = 0, 1, 2, \dots$, $V_1 = 0$, where $V_l, l = 1, \dots$ are integer-valued i.i.d. random variables with

$$\mathbb{E}[V_l] < \infty. \quad (60)$$

Let now

$$\tilde{L}_k = \sum_{j=1}^d T_{dk+j} \quad k = 0, 1, \dots \quad (61)$$

The nonnegative process $\{\tilde{L}_k\}_{k=0}^\infty$ is regenerative with respect to $\{V_l\}_{l=1}^\infty$, and by the regenerative theorem and (60), it holds

$$\lim_{k \rightarrow \infty} \frac{\mathbb{E}[\sum_{m=0}^k \tilde{L}_m]}{k} = \frac{\mathbb{E}[\sum_{m=0}^{V_1-1} \tilde{L}_m]}{\mathbb{E}[V_1]}. \quad (62)$$

Observe next that by (54) and (61)

$$\mathbb{E}\left[\sum_{m=0}^{V_1-1} \tilde{L}_m\right] = \mathbb{E}[L_1]. \quad (63)$$

Hence in order to show (55), it suffices to show

$$\mathbb{E}\left[\sum_{m=0}^{V_1-1} \tilde{L}_m\right] < \infty \quad (64)$$

or, by (62)

$$\lim_{k \rightarrow \infty} \frac{\mathbb{E}[\sum_{m=0}^k \tilde{L}_m]}{k} < \infty. \quad (65)$$

Notice that by (61), we have

$$\begin{aligned} \mathbb{E}\left[\sum_{m=0}^k \tilde{L}_m\right] &= \mathbb{E}\left[\sum_{m=1}^{d(k+1)} T_m\right] \\ &= \sum_{m=1}^{d(k+1)} \mathbb{E}[T_m] \\ &\leq \sum_{m=1}^{d(k+1)} \left(\frac{U}{\delta} + (1-\delta)^{m-1}\tau\right) \quad \text{by (48)} \\ &\leq \frac{U}{\delta}(dk + d + 1) + \frac{\tau}{\delta} \end{aligned}$$

from which (65) follows.

b) Consider next the case where $C((\tau, s))$ is inessential, i.e., there is at least one state $y \in \mathcal{G} - C((\tau, s))$ such that for $x \in C((\tau, s))$, $x \rightsquigarrow y$ but $y \not\rightsquigarrow x$; here, with x, y we denote pairs of the form (τ, s) . Hence, there is at least one other communicating class reachable from $C((\tau, s))$. The communicating classes reachable from $C((\tau, s))$ will be either essential or inessential. We argue that the process $\{(T_m, S_m)\}_{m=1}^\infty$ will enter an essential class in finite time. Assume the contrary, that is, there is a set of sample paths Ω_I with $\Pr\{\Omega_I\} > 0$, for which the process remains always in some inessential class. Since inessential states are nonrecurrent (see [18, p. 18, Theorem 4]), the process visits each inessential state only a finite number of times. This implies that on Ω_I , $\lim_{m \rightarrow \infty} T_m = \infty$, and since $\Pr\{\Omega_I\} > 0$, we conclude that

$$\mathbb{E}\left[\lim_{m \rightarrow \infty} T_m \mid (T_1, S_1) = (\tau, s)\right] = \infty.$$

Applying next Fatou's Lemma, we have

$$\begin{aligned} \liminf_{m \rightarrow \infty} \mathbb{E}[T_m \mid (T_1, S_1) = (\tau, s)] \\ \geq \mathbb{E}\left[\lim_{m \rightarrow \infty} T_m \mid (T_1, S_1) = (\tau, s)\right] = \infty \end{aligned}$$

which contradicts (48).

Since the process enters again an essential class in finite time, we can apply the arguments of case a) to complete the proof. ■

APPENDIX E PROOF OF THEOREM 17

Proof: Assume first that $\hat{T}(\mathbf{r}) > 0$. Let the initial system state be $s \in \mathcal{S}$. At time $\lceil \tau/\alpha \rceil$, the following packets will be in the system: 1) the $\mathbf{Q}_s^\pi(\lceil \tau/\alpha \rceil)$ packets at various queues due to processing according to π ; and 2) the $\mathbf{G}_s^\pi(\lceil \tau/\alpha \rceil)$ packets that were originally in the input queues but were not processed by π up to time $\lceil \tau/\alpha \rceil$, i.e.,

$$\mathbf{G}_{s,i}^\pi(\lceil \tau/\alpha \rceil) = \max \left\{ \lceil r_i \tau \rceil - \sum_{t=0}^{\lceil \tau/\alpha \rceil - 1} A_i(t), 0 \right\}.$$

The evacuation time of π_R^ϵ is at most

$$\lceil \tau/\alpha \rceil + F^{\pi_h} (Q_s^\pi(\lceil \tau/\alpha \rceil) + G_s^\pi(\lceil \tau/\alpha \rceil)).$$

Taking expectations, we have

$$\begin{aligned} \bar{T}_s^{\pi_R^\epsilon}(\lceil \tau/\alpha \rceil) &\leq \lceil \tau/\alpha \rceil + \mathbb{E} [F^{\pi_h} (Q_s^\pi(\lceil \tau/\alpha \rceil) + G_s^\pi(\lceil \tau/\alpha \rceil))] \\ &\leq \lceil \tau/\alpha \rceil + C \sum_{i=1}^n \mathbb{E} [Q_{s,i}^\pi(\lceil \tau/\alpha \rceil) + G_{s,i}^\pi(\lceil \tau/\alpha \rceil)]. \end{aligned} \quad (66)$$

Observe now that by the strong law of large numbers

$$\begin{aligned} \lim_{\tau \rightarrow \infty} \frac{G_{s,i}^\pi(\lceil \tau/\alpha \rceil)}{\tau} &= \lim_{\tau \rightarrow \infty} \max \left\{ \frac{\lceil \tau/\alpha \rceil}{\tau} - \frac{1}{\alpha} \sum_{t=0}^{\lceil \tau/\alpha \rceil - 1} \frac{A_i(t)}{\tau/\alpha}, 0 \right\} \\ &\stackrel{(21)}{=} \max \left\{ r_i - \frac{1}{\alpha} \lambda_{\epsilon,i}, 0 \right\} \leq \frac{|\alpha r_i - \lambda_{\epsilon,i}|}{\alpha}. \end{aligned}$$

By uniform integrability of $\sum_{t=0}^{\tau} A_i(t)/\tau$, we conclude from the last inequality that

$$\lim_{\tau \rightarrow \infty} \mathbb{E} \left[\frac{G_{s,i}^\pi(\tau/\alpha)}{\tau} \right] \leq \frac{|\alpha r_i - \lambda_{\epsilon,i}|}{\alpha}.$$

Dividing by τ and taking limits in (66) and using the fact that by the stability of π and Lemma 22 it holds

$$\lim_{\tau \rightarrow \infty} \sum_{s \in S} \frac{\mathbb{E} [Q_{s,i}^\pi(\lceil \tau/\alpha \rceil)]}{\tau/\alpha} = 0$$

we then have

$$\limsup_{\tau \rightarrow \infty} \frac{\bar{T}_s^{\pi_R^\epsilon}(\lceil \tau/\alpha \rceil)}{\tau} \leq \frac{1}{\alpha} + \frac{C}{\alpha} \sum_{i=1}^n |\alpha r_i - \lambda_{\epsilon,i}| \leq \hat{T}(\mathbf{r}) + \frac{C}{\alpha} \epsilon.$$

Hence

$$\begin{aligned} \limsup_{\tau \rightarrow \infty} \max_{s \in S} \left\{ \frac{\bar{T}_s^{\pi_R^\epsilon}(\tau\mathbf{r})}{\tau} \right\} &= \max_{s \in S} \left\{ \limsup_{\tau \rightarrow \infty} \frac{\bar{T}_s^{\pi_R}(\tau\mathbf{r})}{\tau} \right\} \\ &\leq \hat{T}(\mathbf{r}) + \frac{C}{\alpha} \epsilon \end{aligned}$$

and the theorem follows. Similar arguments for the case $\hat{T}(\mathbf{r}) = 0$ lead to the corresponding bound. ■

REFERENCES

- [1] F. Baccelli and S. Foss, "On the saturation rule for the stability of queues," *J. Appl. Probab.*, vol. 32, pp. 494–507, Jun. 1995.
- [2] E. Altman, "The Baccelli-Foss saturation rule for stability for continuous time input processes," in *Proc. Allerton Conf. Commun., Control, Comput.*, Urbana, IL, USA, Sep. 1994, pp. 396–403.
- [3] J. Harrison and J. Van Mieghem, "Dynamic control of Brownian networks: State space collapse and equivalent workload formulations," *Ann. Appl. Probab.*, vol. 7, no. 3, pp. 747–771, Aug. 1997.
- [4] S. Meyn, *Control Techniques for Complex Networks*. Cambridge, U.K.: Cambridge Univ. Press, Jan. 2008.
- [5] M. Neely, A. Tehrani, and Z. Zhang, "Dynamic index coding for wireless broadcast networks," in *Proc. IEEE INFOCOM*, Orlando, FL, USA, Mar. 2012, pp. 316–324.
- [6] V. Angelakis, A. Ephremides, Q. He, and D. Yuan, "On emptying a wireless network in minimum time," in *Proc. IEEE ISIT*, Cambridge, MA, USA, Jul. 2012, pp. 2671–2675.
- [7] Y. Sagduyu, L. Georgiadis, L. Tassioulas, and A. Ephremides, "Capacity and stable throughput regions for broadcast erasure channel with feedback: An unusual union," *IEEE Trans. Inf. Theory*, vol. 59, no. 5, pp. 2841–2862, May 2013.

- [8] R. Loynes, "The stability of a queue with non-independent inter-arrival and service times," *Math. Proc. Cambridge Philos. Soc.*, vol. 58, no. 3, pp. 497–520, Jul. 1962.
- [9] S. Iyer and N. McKeown, "Maximum size matchings and input queued switches," in *Proc. Allerton Conf. Commun., Control, Comput.*, Urbana, IL, USA, Oct. 2002.
- [10] M. Hall, Jr., *Combinatorial Theory*. Waltham, MA, USA: Blaisdell, 1967.
- [11] T. Weller and B. Hajek, "Scheduling nonuniform traffic in a packet-switching system with small propagation delay," *IEEE/ACM Trans. Netw.*, vol. 5, no. 6, pp. 813–823, Dec. 1997.
- [12] M. Neely, E. Modiano, and Y.-S. Cheng, "Logarithmic delay for $N \times N$ packet switches under the crossbar constraint," *IEEE/ACM Trans. Netw.*, vol. 15, no. 3, pp. 657–668, Jun. 2007.
- [13] G. Paschos, C. Fragiadakis, L. Georgiadis, and L. Tassioulas, "Wireless network coding with partial overhearing information," in *Proc. IEEE INFOCOM*, Turin, Italy, Apr. 2013, pp. 2337–2345.
- [14] L. Libman, G. Paschos, and L. Georgiadis, "Throughput regions and optimal policies in wireless networks with opportunistic routing," in *Proc. WiOpt*, Tsukuba, Japan, May 2013, pp. 516–523.
- [15] L. Georgiadis, G. Paschos, L. Tassioulas, and L. Libman, "Stability and capacity through evacuation times," in *Proc. Inf. Theory Workshop (ITW)*, Lausanne, Switzerland, Sep. 2012, pp. 347–351.
- [16] N. Bingham and A. Ostaszewski, "Generic subadditive functions," *Proc. Amer. Math. Soc.*, vol. 136, no. 12, pp. 4257–4266, Jul. 2008.
- [17] P. Billingsley, *Probability and Measure*. New York, NY, USA: Wiley, 1995.
- [18] K. Chung, *Markov Chains With Stationary Transition Probabilities*. New York, NY, USA: Springer-Verlag, 1960.
- [19] N. Popov, "Conditions for geometric ergodicity of countable Markov chain," *Sov. Math. Dokl.*, vol. 18, no. 3, pp. 676–679, 1977.
- [20] D. Isaacson, "A characterization of geometric ergodicity," *Probability Theory and Related Fields*, vol. 49, no. 3, pp. 267–273, 1979.
- [21] S. Asmussen, *Applied Probability and Queues*. New York, NY, USA: Wiley, 1987.

Leonidas Georgiadis (S'76–M'78–SM'96) received the Diploma degree in electrical engineering from Aristotle University of Thessaloniki, Thessaloniki, Greece, in 1979, and the M.S. and Ph.D. degrees in electrical engineering from the University of Connecticut, Storrs, CT, in 1981 and 1986, respectively.

Since 1995, he has been with the Telecommunications Department, Aristotle University, Thessaloniki, Greece. His interests are in the areas of wireless networks, high-speed networks, distributed systems, routing, scheduling, congestion control, modeling, and performance analysis.

Georgios S. Paschos received the diploma in electrical and computer engineering from Aristotle University of Thessaloniki, Thessaloniki, Greece, in 2002, and the Ph.D. degree from the University of Patras, Patras, Greece, in 2006, both in electrical and computer engineering.

Since 2012, he has been a Postdoctoral Associate with the Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology (MIT), Cambridge, MA, USA. His main interests are in the areas of wireless communications, networks, and stochastic modeling.

Lavy Libman (S'97–M'04–SM'08) received the B.Sc. degrees in electrical engineering and computer engineering and the M.Sc. and Ph.D. degrees in electrical engineering from the Technion—Israel Institute of Technology, Haifa, Israel, in 1992, 1997, and 2003, respectively.

He is a Senior Lecturer with the School of Computer Science and Engineering, UNSW, Sydney, Australia. His research interests are in cooperative communication, network coding, game-theoretic modeling and protocol design in wireless networks, and their applications in intelligent transport systems.

Leandros Tassioulas (S'89–M'91–SM'05–F'07) received the Diploma from the Aristotle University of Thessaloniki, Thessaloniki, Greece, in 1987, and the Ph.D. degree from the University of Maryland, College Park, MD, USA, in 1991, both in electrical engineering.

He has been a Professor of telecommunication networks with the Department of Computer and Communication Engineering, University of Thessaly, Volos, Greece, since 2002 and Associate Director of the Informatics and Telematics Institute of the Center for Research and Technology Hellas (CERTH). His research interests are in the field of computer and communication networks with emphasis on fundamental mathematical models and algorithms, architectures and protocols of wireless systems, sensor networks, novel internet architectures, and network testbed experimentation.