

MIT Open Access Articles

*Distributed Random Access Algorithm:
Scheduling and Congestion Control*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Walrand, J., with Libin Jiang and Shah, D., Jinwoo Shin. "Distributed Random Access Algorithm: Scheduling and Congestion Control." *Information Theory, IEEE Transactions On* 56.12 (2010) : 6182-6207. Copyright © 2010, IEEE

As Published: <http://dx.doi.org/10.1109/TIT.2010.2081490>

Publisher: Institute of Electrical and Electronics Engineers

Persistent URL: <http://hdl.handle.net/1721.1/61979>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Distributed Random Access Algorithm: Scheduling and Congestion Control

Libin Jiang, *Student Member, IEEE*, Devavrat Shah, *Member, IEEE*, Jinwoo Shin, and Jean Walrand, *Fellow, IEEE*

Abstract—This paper provides proofs of the rate stability, Harris recurrence, and ε -optimality of carrier sense multiple access (CSMA) algorithms where the random access (or backoff) parameter of each node is adjusted dynamically. These algorithms require only local information and they are easy to implement. The setup is a network of wireless nodes with a fixed conflict graph that identifies pairs of nodes whose simultaneous transmissions conflict. The paper studies two algorithms. The first algorithm schedules transmissions to keep up with given arrival rates of packets. The second algorithm controls the arrivals in addition to the scheduling and attempts to maximize the sum of the utilities, in terms of the rates, of the packet flows at different nodes. For the first algorithm, the paper proves rate stability for strictly feasible arrival rates and also Harris recurrence of the queues. For the second algorithm, the paper proves the ε -optimality in terms of the utilities of the allocated rates. Both algorithms are iterative and we study two versions of each of them. In the first version, both operate with strictly local information but have relatively weaker performance guarantees; under the second version, both provide stronger performance guarantees by utilizing the additional information of the number of nodes in the network.

Index Terms—Congestion control, distributed medium access, positive recurrent, random access, scheduling.

I. INTRODUCTION

THE problem of scheduling and controlling congestion in networks with conflicting nodes has received considerable attention over the last few years for communication networks such as the Internet (cf., [34] and [53]), stochastic processing networks (cf., [25] and [24]), and switched networks (cf., [59]). Chronologically, the major steps have been random access algorithms, the stability of maximum weight (MW) scheduling, randomized versions of MW, greedy algorithms

with good throughput properties, and optimal¹ local algorithms. We provide a brief overview of literature in that order followed by contributions of this paper.

A number of random access algorithms for scheduling transmissions of nodes were proposed, starting with the classical ALOHA protocol [1], [43]. Hajek and van Loon [23] first showed that an adaptive version of ALOHA achieves the maximum throughput possible for that particular type of network. Works by Kelly and McPhee [33], [32], [40], Mosely and Humblet [48], Tsybakov and Likhanov [65], Aldous [2], Hastad, Leighton, and Rogoff [26], and Goldberg *et al.* [20] establish various negative and positive results about the setup when time is slotted, packets are unit size, and packets may be queued or not queued. These papers assume that the nodes do not sense the transmission of other nodes. For an online survey (until October 2002) of contention resolution without carrier sense, see [21]. More recently, Gupta and Stolyar [22] and Stolyar [62] proposed algorithms that can achieve the capacity of slotted ALOHA by dynamically adjusting the access probabilities. However, this adaptation requires information exchange between nodes. Another class of random access algorithms are based on carrier sense multiple access (CSMA) where nodes can “sense” whether any conflicting transmission is active. Eryilmaz, Marbach, and Ozdaglar [41] showed that with a particular interference model, by properly choosing the access probabilities with CSMA, the maximum throughput can be achieved in the asymptotic regime of small sensing delay and large network size. A related work by Bordenave, McDonald, and Proutière [4] analyzes the “capacity” of large network (or mean field limit) for a given set of access probabilities.

The MW algorithm was proposed by Tassioulas and Ephremides [64]. This algorithm schedules the independent set (nonconflicting nodes) with the maximum sum of queue lengths. These authors show that the sum of the squares of the queue lengths is a Lyapunov function, thus proving stability. Variants of this algorithm have good delay properties (cf., [58] and [59]). Unfortunately, finding the MW independent set is NP-complete and requires global information of the queue lengths, making such algorithms difficult to implement. The central idea of considering the maximization of the sum of the user utilities is due to [34]. See also [39] and [45]. Combining this objective with the scheduling appears in [49], [50], [14], [15], [61], and [36]. For a related survey, see [8] and [60].

¹In this paper, optimal performance means rate stability or positive Harris recurrence of queueing Markov process in the context of scheduling, and utility maximizing rate allocation in the context of congestion control. We will also use throughput optimality to refer to optimality in the context of scheduling problem.

Manuscript received July 04, 2009; revised May 10, 2010. Date of current version November 19, 2010. This work was supported in part by the National Science Foundation (NSF) under Projects CNS 0546590, TF 0728554, by the DARPA ITMANET project, by the AFOSR Complex Networks project, and by the Multi-disciplinary University Research Initiative (MURI) under Grant BAA 07-036.18.

L. Jiang and J. Walrand are with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720 USA (e-mail: ljiang@eecs.berkeley.edu; wlr@eecs.berkeley.edu).

D. Shah is with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: devavrat@mit.edu).

J. Shin is with the Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: jinwoos@mit.edu).

Communicated by R. D. Yates, Associate Editor for Communication Networks.

Digital Object Identifier 10.1109/TIT.2010.2081490

Randomized versions of MW algorithm by Tassiulas [63] and its variant by Giaccone, Prabhakar, and Shah [19] provide a simpler (centralized) implementation of MW for input-queued switches while retaining the throughput property. A distributed implementation of this algorithm based on distributed sampling and distributed (a la gossip; cf., [55]) summation procedure was proposed by Modiano, Shah, and Zussman [46]. This algorithm, though simple and distributed, requires network-wide information exchange for each new scheduling decision. To overcome this limitation, Rajagopalan, Shah, and Shin [51], [52] proposed a random access CSMA algorithm in which the access probability of a node is a function of its own queue size and the estimation of the maximum of queue sizes in the network. The maximum of queue sizes in the network is a global property. However, as shown in [52], a useful estimation of it can be maintained at each node through exactly *one* message/number (through broadcast transmission) exchange with its neighbor. This algorithm can be viewed as a continuous-time-reversible Markov chain on the space of schedules (independent sets of conflict graph) with time-varying transition probabilities that are function of the queue sizes. In [51] and [52], authors show that the average weight of schedule with respect to the invariant distribution of this Markov chain (also known as the Glauber dynamics in statistical physics literature) based on the instantaneous queue sizes is close to the maximum weight. Therefore, if this reversible Markov chain is always close to its invariant distribution, then the algorithm effectively simulates the MW. However, queue sizes change and hence the transition probabilities as well as stationary distribution of this Markov chain changes. In [52], through a novel *network adiabatic theorem* for reversible Markov chains with time-varying transition probabilities, authors establish that by choosing access probabilities as slowly varying function [like $\log \log(\cdot + e)$] of queue sizes, the above mentioned Markov chain remains always close to its invariant distribution. This subsequently establishes the throughput optimality of the algorithm. In summary, the choice of an appropriate function of the queue size plays key role in establishing the throughput optimality.

In [52], authors also propose an algorithm in which each node chooses its access probability as the slowly varying function of its own queue size only (i.e., ignoring the estimation of maximum of queue sizes). This algorithm is totally distributed (no exchange of information or messages between neighbors). They conjecture it to be throughput optimal (in a stronger sense of Harris recurrence). The conjecture, as this paper is written, remains unresolved. Further extensions of the algorithm of [51] and [52] to circuit-switched networks is provided in recent work by Shah and Shin [56]. An interested reader can find a summary of design and analysis of MW-based scheduling algorithms (until 2007) for switched networks in a book chapter by Shah [54].

Greedy algorithms are simpler than MW. Parallel iterative matching [3] and iSLIP [42] were shown to be 50% throughput optimal [9]. Subsequently, Dimakis and Walrand [13] identified sufficient conditions on the network topology for throughput optimality of greedy algorithms. Those conditions were further weakened to obtain fractional throughput results about a class of wireless networks by Joo, Lin, and Shroff [31] and

Lecante, Ni, and Srikant [35]. These algorithms are generally not throughput optimal and require multiple rounds of message exchanges among nodes.

Another class of local algorithms was proposed by Jiang and Walrand [29]. The algorithms adjust access probabilities in CSMA for both scheduling and congestion control by means of a novel optimization problem and its relation to certain reversible networks. The result is a totally distributed algorithm. They conjecture it to be throughput optimal and utility maximizing for scheduling and congestion control respectively. In [30], the authors use a suggestion by Shah to adjust the access probabilities over increasing intervals, and they adapt techniques from stochastic approximation to prove the convergence, rate stability, and optimality of the algorithms in [29]. Independently, Liu *et al.* [37] showed that, under some technical assumptions, the algorithm in [29] converges to an approximate utility maximizing solution. However, their result does not establish the throughput optimality (i.e., stability of queue size in some form). Further, implicitly their algorithm requires some knowledge about the entire system.

The key idea of [29] is that, instead of using the MW schedule, the algorithm attempts to improve the schedule to match the arrival rates into the queues. The scheduling algorithm is parameterized by the access probabilities or aggressiveness or backoff time, which decide the rates at which nodes attempt transmission using carrier sense information. One then defines a distance between the distribution (over the independent sets) realized by the current parameters of the algorithm and desired distribution. The key point of the algorithm is to minimize that distance by adjusting the parameters. The gradient of that distance with respect to the aggressiveness of one node turns out to be the difference between the average service and arrival rates at that node. Accordingly, the algorithm follows the gradient and adjusts the parameters based on the empirical service and arrival rates at the nodes and thus is local. Indeed, the queue length reflects this difference between service and arrival rate. In that sense, this algorithm can be interpreted as utilizing queue lengths.

The technical challenges in proving the convergence and the optimality of the algorithm in [29] is as follows. First, given a set of CSMA parameters (the aggressiveness of the nodes), a node can only measure the empirical arrival and service rates, but not the average values of those quantities required by the algorithm in principle. Second, the algorithm keeps changing the CSMA parameters. The intuition to overcome the challenges is that if the parameters remain constant for long enough (i.e., they are changed slowly), then the distribution of the underlying Markov chain approaches its invariant distribution. Consequently, the algorithm approaches the desired gradient algorithm. The general idea of using a random version of the desired gradient is at the heart of stochastic approximation (see [5], [6], and [37]). Usually, in such scenarios, when the controlled variables (in this case the CSMA parameters) are confined in a “compact” set, then generic results from stochastic approximation (cf., [5] and [6]) will provide desired convergence result. However, it is not the case here and therefore additional steps are required to show that the Markov chain approaches its invariant distribution fast enough. The needed technical tool is a bound on the mixing time of the Markov chain. Here, as in [52] and [30], we use

a uniformized version of the continuous-time Markov chain to exploit a bound available for the mixing time of discrete-time Markov chains.

The current paper provides an alternate proof of the rate stability in the scheduling algorithm. Moreover, it proves the Harris recurrence of the queue lengths when using a variant of the algorithm that, in addition, requires that each node knows the total number of nodes in the network. (Under that assumption, we also show that for any given $\varepsilon > 0$, there is a congestion control algorithm that is ε -optimal.) Both the proof in [30] and the current proof of rate stability first establish the convergence of the CSMA parameters to some desired values, and then use this to prove rate stability. One major difference is that to establish the convergence, the proof in [30] uses a quadratic Lyapunov function, whereas the current proof uses a Lyapunov function related to KL-divergence. The proof of the Harris recurrence involves constructing a “petite set” that is positive recurrent. Intuitively, such a set is a generalization of a recurrent state for a countable Markov chain. (The state space is not countable in our problem.) Once the Markov chain hits this set, it starts afresh with at least some measure, thus leading to the ergodicity of this Markov chain.

Finally, it is worth taking note of similarities and difference between the algorithm of Rajagopalan, Shah, and Shin [52] and the algorithm of Jiang and Walrand [29]. Both of these algorithms are random access CSMA and adjust the access (or backoff) parameters adaptively as function of local information. Through this adjustment, both algorithms induce a time-varying Markov chain on the space of schedules (independent sets of conflict graph) that is reversible. However, both algorithms differ in the way they adjust or adapt the access (or backoff) parameters. Specifically, to adjust the access parameters, the algorithm of [52] uses a slowly varying function of instantaneous queue sizes while the algorithm of [29] (considered in this paper) uses empirical arrival and service rates. While the difference of these rates has relation to queue sizes induced over the adjustment interval, it is not instantaneous queue sizes considered in [52], if we use time-varying step sizes and update intervals (as in Algorithm 1). On the other hand, if we use constant step sizes and update intervals (as in the variant in [28]), then the access parameters become proportional to the queue sizes (although extra care needs to be taken to keep the parameters bounded, slightly decreasing the capacity region).

The paper is organized as follows. Section II defines the network model. The main results are stated in Section III. Some preliminaries about Markov chains as well as a relevant (CSMA) Markov chain are introduced in Section IV. The throughput properties of scheduling algorithms are proved in Section V. Specifically, rate stability and Harris recurrence properties are proved in Sections V-A and V-B, respectively. Section VI analyzes the congestion control problem. Section VII concludes the paper.

II. MODEL AND PROBLEM STATEMENT

Our network graph is a collection of n queues. Time is indexed by $t \in \mathbb{R}_+$. Let $Q_i(t) \in \mathbb{R}_+$ denote the amount of work in the i th queue at time t and let $\mathbf{Q}(t) = [Q_i(t)]_{1 \leq i \leq n}$. Initially,

$t = 0$ and $\mathbf{Q}(0) = \mathbf{0}$, i.e., the system starts empty.² Work arrives to each queue either as per an exogenous arrival process or is controlled by each queue as per a certain algorithm. Each queue can potentially be serviced at unit rate resulting in the departure of work from it. Throughout this paper, we will assume single-hop network. That is, once work departs from a queue, it leaves the network. In this paper, we will not consider multihop network. We strongly believe that the results of this paper can be extended to multihop network in a straightforward manner (cf., see [60]).

The queues are offered service as per the constraint imposed by interference. To define this constraint, let $G = (V, E)$ denote the inference graph between queues. Here vertices $V = \{1, \dots, n\}$ represent the n queues and edges $E \subset V \times V$ represent interfering queues: $(i, j) \in E$ iff transmissions of queues i and j interfere with each other (that is, the interference relationship is symmetric). Let $\mathcal{N}(i) = \{j \in V : (i, j) \in E\}$ denote the neighbors of node i . Let $\sigma_i(t) \in \{0, 1\}$ denote whether queue i is transmitting at time t , with notation that $\sigma_i(t) = 1$ represents transmission. Let $\boldsymbol{\sigma}(t) = [\sigma_i(t)]$. Then, interference imposes the constraint that for all $t \in \mathbb{R}_+$

$$\boldsymbol{\sigma}(t) \in \mathcal{I}(G)$$

$$\mathcal{I}(G) \triangleq \{\boldsymbol{\rho} \in \{0, 1\}^n : \rho_i + \rho_j \leq 1, \forall (i, j) \in E\}. \quad (1)$$

The resulting queueing dynamics are described as follows. For $0 \leq s < t$ and $1 \leq i \leq n$

$$Q_i(t) = Q_i(s) - \int_s^t \sigma_i(r) \mathbf{1}_{\{Q_i(r) > 0\}} dr + A_i(s, t)$$

where $A_i(s, t)$ denotes the cumulative arrival to queue i in the time interval $(s, t]$ and $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function. Finally, define the cumulative departure process $\mathbf{D}(t) = [D_i(t)]$, where

$$D_i(t) = \int_0^t \sigma_i(r) \mathbf{1}_{\{Q_i(r) > 0\}} dr.$$

We define the capacity region of such a network. The capacity region $\mathcal{C} \subset [0, 1]^n$ is the convex hull of the feasible scheduling set $\mathcal{I}(G)$, i.e.,

$$\mathcal{C} = \left\{ \sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \alpha_{\boldsymbol{\rho}} \boldsymbol{\rho} : \sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \alpha_{\boldsymbol{\rho}} = 1, \alpha_{\boldsymbol{\rho}} \geq 0, \forall \boldsymbol{\rho} \in \mathcal{I}(G) \right\}.$$

The intuition behind this definition of capacity region comes from the fact that any algorithm has to choose a schedule from $\mathcal{I}(G)$ at each time and hence the time average of the “service rate” induced by any algorithm must belong to \mathcal{C} .

Scheduling Problem. In this setup, we assume that the arrival process at each queue is exogenous. Recall that $A_i(s, t)$ denotes the work that has arrived to queue i in the time interval $(s, t]$; $A_i(t) \triangleq A_i(0, t)$ represents cumulative arrival process.

²The assumption of system starting empty is not crucial and it is chosen for the ease of exposition. Specifically, most results stated in this paper remain unchanged even if the initial queues take any finite values. The only exception is that, in Theorem 4, the bound on the queue size is increased by the initial queue size.

We assume that the increments in the arrival process over integral times, i.e., $A_i(k, k+1)$ for $k \in \mathbb{Z}_+$, are independent and identically distributed (i.i.d.) with bounded support. Moreover, we assume that $A_i(1) \in [0, K]$ and $\Pr(A_i(1) = 0) > 0$ for all i . Note that this setup naturally allows for A_i and A_j to be very different processes for $i \neq j$. Finally, we define $\lambda_i = \mathbb{E}[A_i(1)]$. Under our setup the strong law of large numbers implies that

$$\lim_{t \rightarrow \infty} \frac{A_i(t)}{t} = \lambda_i, \quad \text{with probability 1.} \quad (2)$$

Let $\lambda = [\lambda_i]$. We assume that $\lambda_{\min} \triangleq \min_{1 \leq i \leq n} \lambda_i > 0$ without loss of generality.³ In this setup, we need a scheduling algorithm that decides $\sigma(t)$ each instant $t \in \mathbb{R}_+$. Intuitively, we would expect that a good algorithm will keep the queues as small as possible. To make this notion formal, first note that if $\lambda \notin \Lambda$, then no algorithm can keep the queues finite, where

$$\Lambda = \{\lambda \in \mathbb{R}_+^n : \lambda \leq \gamma \text{ componentwise, for some } \gamma \in \mathcal{C}\}.$$

Motivated by this observation, we call λ *strictly admissible* if $\lambda \in \Lambda^\circ$, where

$$\Lambda^\circ = \{\lambda \in \mathbb{R}_+^n : \lambda < \gamma \text{ componentwise, for some } \gamma \in \mathcal{C}\}.$$

We call a scheduling algorithm *rate stable* if for any $\lambda \in \Lambda^\circ$, the following holds with probability 1:

$$\lim_{t \rightarrow \infty} \frac{1}{t} D_i(t) = \lambda_i, \quad \forall 1 \leq i \leq n.$$

Given (2), this is equivalent to

$$\lim_{t \rightarrow \infty} \frac{Q_i(t)}{t} = 0, \quad \forall 1 \leq i \leq n.$$

Rate stability is a weaker notion of *throughput optimality* or *stability* of the network. A stronger notion requires that for any $\lambda \in \Lambda^\circ$ the underlying network Markov process is *positive recurrent* or more generally *positive Harris recurrent*. This is a stronger property compared to the rate stability as it implies existence of unique stationary distribution and in our setup ergodicity of the network Markov process. Subsequently, it leads to the finiteness of queue sizes with probability 1 in stationarity. See Section IV-C for definition and further implications of Harris recurrence.

In summary, the problem of scheduling requires designing an algorithm that makes the network-wide decisions $\sigma(t) \in \mathcal{I}(G)$ for all t so that the network is throughput optimal (rate stable or positive recurrent). The algorithm should utilize only local information, i.e., $\sigma_i(t)$ should be based on the *history* observed at node i only and the *sensing* information available at node i about which of its neighbors are transmitting at a given time.

Congestion Control Problem. In this setup, unlike the scheduling problem, we require each node or queue to control its arrival or data generation process. Specifically, at each node i , an algorithm decides the rate $\lambda_i(t) \in [0, 1]$ at each time t . The data

are generated at node i as per a deterministic process with rate $\lambda_i(t)$ at time t . That is, for any $0 \leq s < t$

$$A_i(s, t) = \int_s^t \lambda_i(r) dr.$$

Given the arrival or data generation process, the remaining problem is similar to scheduling. That is, an algorithm is required to make decisions $\sigma(t) \in \mathcal{I}(G)$ for all t using only local information and so as to keep queues stable. Now in order to determine the right rate allocation, we assume that all nodes have some utility. Let $U_i : [0, 1] \rightarrow \mathbb{R}$ be a strictly concave and increasing utility function of node i , with $U_i(x)$ representing the value of its utility when it is allocated rate $x \in [0, 1]$. Then, ideally, we wish nodes to allocate rates $\lambda^* = [\lambda_i^*]$ where

$$\lambda^* = \operatorname{argmax}_{\lambda \in \Lambda} \sum_{i=1}^n U_i(\lambda_i) \quad \text{over } \lambda \in \Lambda. \quad (3)$$

In summary, the problem of congestion control requires designing an algorithm that makes decisions $\lambda(t) \in [0, 1]^n$ and $\sigma(t) \in \mathcal{I}(G)$ for all t so that $\lambda(t) \rightarrow \lambda^*$ and the network of queues is stable, i.e., rate stable or positive Harris recurrent. The algorithm should utilize only local information, i.e., both $\lambda_i(t)$ and $\sigma_i(t)$ should be based on the *history* observed at node i only and the *sensing* information available at node i about which of its neighbors are transmitting at a given time.

III. MAIN RESULTS

This section describes our algorithms and theorems stating their performance guarantees for scheduling and congestion control. The algorithms presented here are variants of algorithms proposed in an earlier work [29]. As noted earlier, this paper provides an alternate proof of the rate stability established in [30] and the new result of Harris recurrence.

A. Scheduling Algorithm

The algorithm to decide $\sigma(t)$ through local decisions $\sigma_i(t)$ can be classified as based on random access using carrier sense information or CSMA. The basic operation of each node under such an algorithm can be described as follows. In between two transmissions, a node waits for a random amount of time—also known as *backoff*. Each node can sense the medium perfectly and instantly, i.e., knows if any other interfering node is transmitting at a given time instance. If a node that finishes waiting senses the medium to be busy, it starts waiting for another random amount of time; else, it starts transmitting for a random amount of time. The nodes repeat this operation. The difference between all such protocols lies in the selection of the random waiting time and random transmission time.

In this paper, we assume that node i 's random waiting time and transmission time have exponential distributions with mean $1/R_i$ and 1, respectively. Therefore, the performance of the algorithm is solely determined by the parameters $R_i, 1 \leq i \leq n$. In essence, our scheduling algorithm will *learn* a good value for R_i at each node i using only local information, so that the

³Note that, if $\lambda_i = 0$ for some i , then algorithm will ignore such queues by setting their access probability to 0.

performance of the algorithm is throughput optimal. It is somewhat surprising that such a simple class of algorithms can indeed achieve the optimal throughput.

More precisely, let $R_i(t)$ be the value of parameter R_i at time t . Given that $R_i(t)$ changes over time, the waiting time becomes distributed according to an exponential distribution with time-varying rate. A convenient way to think of this is as follows. Suppose node i starts its new waiting period at time t_1 and is still waiting at time $t > t_1$. Then, given the history until time t , the waiting time ends during $(t, t + \varepsilon)$ with probability $R_i(t)\varepsilon + o(\varepsilon)$.

Given the above description, the scheduling algorithm is completely determined once we describe how $R_i(t)$ are decided for all i and all $t \in \mathbb{R}_+$. For convenience, we describe the algorithm for selecting $r_i(t) \triangleq \ln R_i(t)$. The algorithm, at each node i , updates $r_i(t)$ at time instances $l(j), j \in \mathbb{Z}_+$ with $l(0) = 0$. Also, $r_i(t)$ remains the same between times $l(j)$ and $l(j+1)$ for all $j \in \mathbb{Z}_+$. To begin with, the algorithm sets $r_i(0) = 0$ for all i . With an abuse of notation, from now onwards, we denote by $r_i(j)$ the value of $r_i(t)$ for all $t \in [l(j), l(j+1))$. Finally, define $T(j) = l(j+1) - l(j)$ for $j \geq 0$. Note that $T(0) = l(1) - l(0) = l(1)$.

In what follows, we describe two variants that differ in the choice of $l(j)$ and the update procedure $r_i(\cdot)$. The first variant uses strictly local information while the second variant uses information about the number of nodes in the network and a performance parameter $\varepsilon > 0$. We provide theorems quantifying the performance of these variants as well.

Scheduling Algorithm 1. In this variant, we use a varying update interval $T(j)$. Specifically, we select

$$T(j) = \exp(\sqrt{j}), \quad \text{for } j \geq 1.$$

Also, we choose a step size $\alpha(j)$ of the algorithm as

$$\alpha(j) = \frac{1}{j}, \quad \text{for } j \geq 1.$$

Given this, node i updates $r_i(\cdot)$ as follows. Let $\hat{\lambda}_i(j), \hat{s}_i(j)$ be empirical arrival and service observed at queue i in $[l(j), l(j+1))$. That is

$$\begin{aligned} \hat{\lambda}_i(j) &= \frac{1}{T(j)} A_i(l(j), l(j+1)) \\ \hat{s}_i(j) &= \frac{1}{T(j)} \left[\int_{l(j)}^{l(j+1)} \sigma_i(t) dt \right]. \end{aligned}$$

Then, the update $r_i(j+1)$ of $r_i(j)$ is defined by

$$r_i(j+1) = r_i(j) + \alpha(j)(\hat{\lambda}_i(j) - \hat{s}_i(j)) \quad (4)$$

with initial condition $r_i(0) = 0$. This update rule is essentially an approximate gradient algorithm for the optimization problem (26).

Note that, under this update rule, the algorithm at each node i uses only its local history. Despite this, we establish that this algorithm is rate stable. Formally, we obtain the following result.

Theorem 1: The scheduling algorithm with updating rule (4) as described above is rate stable for any $\lambda \in \Lambda^o$.

Remark: In the above algorithm, the update interval $T(j)$'s are chosen to be quite large. This may result in large convergence time as well as queue lengths. The purpose of choosing large $T(j)$'s is to ensure that the Markov chain corresponding to the algorithm decisions reaches its stationary distribution in a given interval so as to allow for accurate estimates of mean service rates. According to the conservative estimation in Section IV-B, the worst case mixing time is exponential in the network size n [cf., (23)], which naturally calls for large $T(j)$'s. This issue may be related to the inherent complexity of scheduling: Shah, Tse, and Tsitsiklis [57] have established that in the worst case, no throughput-optimal scheduling policy can achieve polynomial queue lengths on average. On the other hand, in reality, the network topologies (i.e., the conflict graphs) have certain structure and may not be the worst case. For example, when the conflict graph is a complete graph, the mixing time is polynomial in n . In such cases, it is possible to modify the parameters of algorithms to achieve smaller convergence time and queue lengths.

Scheduling Algorithm 2. In this variant, we use $T(j) = T$ for some fixed T . The choice of T will depend on two quantities—the number n of nodes in the network (we assume $n > 3$ here) and $\varepsilon > 0$ that characterizes the approximate stability of the system. Specifically

$$T \triangleq T(n, \varepsilon) = \exp\left(\Theta\left(\frac{n^2}{\varepsilon} \log \frac{n}{\varepsilon}\right)\right). \quad (5)$$

Then, the updating rule becomes

$$r_i(j+1) = [r_i(j) + \alpha(\hat{\lambda}_i(j) + \varepsilon - \hat{s}_i(j))]_{\frac{n}{\varepsilon}} \quad (6)$$

where $\alpha = \alpha(n, \varepsilon) = \varepsilon^2 n^{-2} / 72(K+1)^2$ (here, K is the Lipschitz constant⁴ for the cumulative arrival process) and if $\hat{x}_i = [x_i]_{\frac{n}{\varepsilon}}$ then

$$\hat{x}_i = \begin{cases} \frac{n}{\varepsilon}, & \text{if } x_i > \frac{n}{\varepsilon} \\ -\frac{n}{\varepsilon}, & \text{if } x_i < -\frac{n}{\varepsilon} \\ x_i, & \text{otherwise.} \end{cases} \quad (7)$$

We state the following throughput optimal property of the algorithm using this rule.

Theorem 2: For given $\varepsilon > 0$, under the above described scheduling algorithm, the network is positive Harris recurrent if $\lambda + \mathbf{1} \cdot 2\varepsilon \in \Lambda^o$.

B. Congestion Control Algorithm

The algorithm for congestion control has to select the appropriate values of $r_i(\cdot)$ and the arrival rates $\lambda_i(\cdot)$. These decisions have to be taken so that the arrival rates maximize overall network utility while keeping the queues stable.

Like in the scheduling problem, the algorithm for congestion control updates its choice of $r_i(t)$ and $\lambda_i(t)$ at time instances $l(j), j \in \mathbb{Z}_+$ with $l(0) = 0$. To begin with, it sets $r_i(0) = 0$ and $\lambda_i(0) = 1$ for all i . Again, with an abuse of notation, from now onwards, we denote by $r_i(j)$ [resp., $\lambda_i(j)$] the value of $r_i(t)$

⁴A function $f: \mathbb{R}^M \rightarrow \mathbb{R}^N$ for some $M, N \geq 1$, is called Lipschitz with constant K if $\|f(x) - f(y)\| \leq K\|x - y\|$.

[resp., $\lambda_i(t)$] for all $t \in [l(j), l(j+1))$). As before, define $T(j) = l(j+1) - l(j)$ for $j \geq 0$. Note that $T(0) = l(1) - l(0) = l(1)$.

In what follows, we describe two algorithms for congestion control. Like the two scheduling algorithms, the first variant does not utilize any global information while the second variant utilizes information about number of nodes and a performance parameter.

Congestion Control Algorithm 1. Here, $T(j) = \exp(\sqrt{j})$, $\alpha(j) = \frac{1}{j}$ for $j \in \mathbb{N}$. The $r_i(\cdot)$, $\lambda_i(\cdot)$ are updated as follows: for all i

$$\begin{aligned} r_i(j+1) &= [r_i(j) + \alpha(j)(\lambda_i(j) - \hat{s}_i(j))]_+ \\ \lambda_i(j+1) &= \arg \max_{y \in [0,1]} (\beta \cdot U_i(y) - r_i(j+1)y) \end{aligned} \quad (8)$$

with initially $\mathbf{r}(0) = \mathbf{0}$ and $\lambda(0) = \mathbf{1}$. Here $\beta > 0$ is an algorithm parameter and it plays a role in determining the efficiency of the algorithm. As before, each node updates its parameters based only on local information. Recall that each node i accepts data at rate $\lambda_i(j)$ in $[l(j), l(j+1))$ deterministically. We state the following result about the performance of this algorithm.

Theorem 3: Under the above described algorithm, the queues $\mathbf{Q}(\cdot)$ and arrival rates $\lambda(\cdot)$ are such that

$$\lim_{t \rightarrow \infty} \frac{Q_i(t)}{t} = 0 \quad \text{and} \quad \lim_{j \rightarrow \infty} \lambda(j) = \bar{\lambda}$$

with probability 1, where $\bar{\lambda}$ is such that

$$\left(\sum_i U_i(\bar{\lambda}_i) \right) \geq \left(\sum_i U_i(\lambda_i^*) \right) - \frac{\log |\mathcal{I}(G)|}{\beta}. \quad (9)$$

Here λ^* represents a solution to the utility maximization problem (3).

Congestion Control Algorithm 2. Here, the step size $T(j)$ is constant, and equals a large value T , for all j . In addition to the above, we assume that $U_i(\cdot)$ are such that

$$V = \max_i U_i'(0) < \infty \quad (10)$$

and V is known to all nodes. The algorithm performance parameter is $\varepsilon > 0$. The step size α is a small, fixed constant in $(0, 1)$. Let $\beta = 4n/\varepsilon$. Select T such that

$$T = \exp(\Theta(\beta n V)) \Theta \left(\frac{(\beta V + \alpha)n^2}{\beta \varepsilon} \right).$$

The updating rule is as follows. For all i

$$\begin{aligned} r_i(j+1) &= [r_i(j) - \alpha \hat{s}_i(j)]_+ + \alpha \lambda_i(j) \\ \lambda_i(j+1) &= \arg \max_{y \in [0,1]} (\beta \cdot U_i(y) - r_i(j+1)y) \end{aligned} \quad (11)$$

with initially $\mathbf{r}(0) = \mathbf{0}$ and $\lambda(0) = \mathbf{1}$.

Remark: The dynamics (11) of $r_j(\cdot)$ is similar to that of queue lengths in [49]. The primary difference arises from the fact that $\hat{s}_i(j)$ here is based on empirical service rate induced by the CSMA algorithm while in [49] it is obtained through a maximal weight scheduling.

We state the following result about this algorithm.

Theorem 4: Under the above described algorithm 2, the queue lengths $\mathbf{Q}(\cdot)$ are such that

$$Q_i(t) \leq \frac{T(\beta V + 2\alpha)}{\alpha}, \quad \text{for all } t \geq 0, \quad \text{for all } i.$$

Further, define $\tilde{\lambda}(J)$ as

$$\tilde{\lambda}_i(J) = \frac{1}{J} \left(\sum_{j=0}^{J-1} \lambda_i(j) \right).$$

Then, with probability 1

$$\liminf_{J \rightarrow \infty} \sum_i U_i(\tilde{\lambda}_i(J)) \geq \sum_i U_i(\lambda_i^*) - \varepsilon. \quad (12)$$

Remark: Assuming T is fixed (i.e., ignoring its dependence on β), the above theorem indicates a tradeoff between the utility gap $\varepsilon = O(1/\beta)$ and the backlog $O(\beta)$. This is similar to the observation made in [49]. In particular, a larger β gives more emphasis on the utility, at the cost of larger backlog. The extra T factor here in the queue length is due to the required mixing time for the CSMA Markov chain to approach its stationary distribution. Indeed, the exponential dependence of T on β captures the hardness of problem similar to that observed in [57] in the context of scheduling. This factor does not show up in [49] because there the maximal weight scheduling decision is assumed to be instantaneous.

IV. PRELIMINARIES

This section recalls relevant known results about establishing bound on mixing time of Markov chains. We will start by setting up basic notations and recalling known definitions.

A. Markov Chain and Mixing Time

Consider a discrete-time, time-homogeneous Markov chain over a finite state space Ω . Let an $|\Omega| \times |\Omega|$ matrix P be its transition probability matrix. If P is irreducible and aperiodic, then the Markov chain has a unique stationary distribution and it is ergodic in the sense that $\lim_{\tau \rightarrow \infty} P^\tau(j, i) \rightarrow \pi_i$ for any $i, j \in \Omega$, where $P^\tau(j, i)$ means the element indexed by (j, i) in the matrix P^τ (where $\tau \in \mathbb{Z}_+$). Let $\pi = [\pi_i]$ denote the stationary distribution of the Markov chain. The adjoint of the transition matrix P , also called the time reversal of P , is denoted by P^* and defined as: for any $i, j \in \Omega$, $\pi(i)P^*(i, j) = \pi(j)P(j, i)$. By definition, P^* has π as its stationary distribution as well. If $P = P^*$ then P is called *reversible*, and in this paper we will be primarily interested in such reversible Markov chains.

As noted earlier, the distribution of the irreducible and aperiodic Markov chain converges to its stationary distribution π starting from any initial condition. To establish our results, we will need quantifiable bounds on the time it takes for the Markov chain to reach close to stationary distribution—popularly known as *mixing time*. To make this notion precise and recall known bound on mixing time, we start with definition of distance between probability distributions.

Definition 1 (Two Distances): Given two probability distributions $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ on a finite space Ω , we define the following two distances. The total variation distance, denoted as $\|\boldsymbol{\mu} - \boldsymbol{\nu}\|_{TV}$, is

$$\|\boldsymbol{\mu} - \boldsymbol{\nu}\|_{TV} = \frac{1}{2} \sum_{i \in \Omega} |\mu_i - \nu_i|.$$

The χ^2 distance, denoted as $\|\frac{\boldsymbol{\nu}}{\boldsymbol{\mu}} - 1\|_{2,\boldsymbol{\mu}}$, is

$$\left\| \frac{\boldsymbol{\nu}}{\boldsymbol{\mu}} - 1 \right\|_{2,\boldsymbol{\mu}}^2 = \|\boldsymbol{\nu} - \boldsymbol{\mu}\|_{2,\frac{1}{\boldsymbol{\mu}}}^2 = \sum_{i \in \Omega} \mu_i \left(\frac{\nu_i}{\mu_i} - 1 \right)^2.$$

We make note of the following relation between the above defined two distances: for any probability distributions $\boldsymbol{\mu}, \boldsymbol{\nu}$, using the Jensen's inequality, we have

$$\begin{aligned} \left\| \frac{\boldsymbol{\nu}}{\boldsymbol{\mu}} - 1 \right\|_{2,\boldsymbol{\mu}} &= \sqrt{\sum_{i \in \Omega} \mu_i \left(\frac{\nu_i}{\mu_i} - 1 \right)^2} \\ &\geq \sum_{i \in \Omega} \mu_i \left| \frac{\nu_i}{\mu_i} - 1 \right| = \sum_{i \in \Omega} |\nu_i - \mu_i| \\ &= 2\|\boldsymbol{\nu} - \boldsymbol{\mu}\|_{TV}. \end{aligned} \quad (13)$$

In general, for any two vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}_+^{|\Omega|}$, we define norm

$$\|\mathbf{v}\|_{2,\mathbf{u}}^2 = \sum_{i \in \Omega} u_i v_i^2.$$

This norm naturally induces a matrix norm that will be useful in determining rate of convergence or mixing time of a finite state Markov chain.

Definition 2 (Matrix Norm): Consider an $|\Omega| \times |\Omega|$ nonnegative valued matrix $A \in \mathbb{R}_+^{|\Omega| \times |\Omega|}$ and a vector $\mathbf{u} \in \mathbb{R}_+^{|\Omega|}$. Then, the matrix norm of A with respect to \mathbf{u} is defined as follows:

$$\|A\|_{\mathbf{u}} = \sup_{\mathbf{v}: \mathbb{E}_{\mathbf{u}}[\mathbf{v}] = 0} \frac{\|A\mathbf{v}\|_{2,\mathbf{u}}}{\|\mathbf{v}\|_{2,\mathbf{u}}}$$

where $\mathbb{E}_{\mathbf{u}}[\mathbf{v}] = \sum_i u_i v_i$.

It can be easily checked that the above definition of matrix norm satisfies the following properties.

P1) For matrices $A, B \in \mathbb{R}_+^{|\Omega| \times |\Omega|}$ and $\boldsymbol{\pi} \in \mathbb{R}_+^{|\Omega|}$

$$\|A + B\|_{\boldsymbol{\pi}} \leq \|A\|_{\boldsymbol{\pi}} + \|B\|_{\boldsymbol{\pi}}.$$

P2) For matrix $A \in \mathbb{R}_+^{|\Omega| \times |\Omega|}$, $\boldsymbol{\pi} \in \mathbb{R}_+^{|\Omega|}$ and $c \in \mathbb{R}$

$$\|cA\|_{\boldsymbol{\pi}} = |c| \|A\|_{\boldsymbol{\pi}}.$$

P3) Let A and B be transition matrices of reversible Markov chains, i.e., $A = A^*$ and $B = B^*$. Let both of them have $\boldsymbol{\pi}$ as their unique stationary distribution. Then

$$\|AB\|_{\boldsymbol{\pi}} \leq \|A\|_{\boldsymbol{\pi}} \|B\|_{\boldsymbol{\pi}}.$$

P4) Let A be the transition matrix of a reversible Markov chain, i.e., $A = A^*$. Let $\boldsymbol{\pi}$ be its stationary distribution. Then

$$\|A\|_{\boldsymbol{\pi}} = \lambda_{\max}$$

where $\lambda_{\max} = \max\{|\lambda| \mid \lambda \neq 1 \text{ is an eigenvalue of } A\}$.

For a probability matrix P , mostly in this paper, we will be interested in the matrix norm of P with respect to its stationary distribution $\boldsymbol{\pi}$, i.e., $\|P\|_{\boldsymbol{\pi}}$. Therefore, unless stated otherwise, if we use matrix norm for a probability matrix without mentioning the reference measure, then it is with respect to the stationary distribution. That is, in the above example, $\|P\|$ will mean $\|P\|_{\boldsymbol{\pi}}$.

With these definitions and fact that P and P^* have the same stationary distribution, say $\boldsymbol{\pi}$, it follows that for any distribution $\boldsymbol{\mu}$ on Ω

$$\left\| \frac{\boldsymbol{\mu}P}{\boldsymbol{\pi}} - 1 \right\|_{2,\boldsymbol{\pi}} \leq \|P^*\| \left\| \frac{\boldsymbol{\mu}}{\boldsymbol{\pi}} - 1 \right\|_{2,\boldsymbol{\pi}} \quad (14)$$

where we have used (abused) notation $\|P^*\| = \|P^*\|_{\boldsymbol{\pi}}$ and since $\mathbb{E}_{\boldsymbol{\pi}}[\frac{\boldsymbol{\mu}}{\boldsymbol{\pi}} - 1] = 0$, with interpretation $\frac{\boldsymbol{\mu}}{\boldsymbol{\pi}} = [\mu(i)/\pi(i)]$. Also, in the above (and throughout the paper), in the left multiplication of a vector with a matrix, the vector should be thought of as a row vector. Therefore, for a reversible Markov chain ($P = P^*$) starting with initial distribution $\boldsymbol{\mu}(0)$, the distribution $\boldsymbol{\mu}(\tau)$ at time τ is such that

$$\left\| \frac{\boldsymbol{\mu}(\tau)}{\boldsymbol{\pi}} - 1 \right\|_{2,\boldsymbol{\pi}} \leq \|P\|^\tau \left\| \frac{\boldsymbol{\mu}(0)}{\boldsymbol{\pi}} - 1 \right\|_{2,\boldsymbol{\pi}}. \quad (15)$$

Now starting from any state i , i.e., probability distribution with unit mass on state i , the initial distance $\left\| \frac{\boldsymbol{\mu}(0)}{\boldsymbol{\pi}} - 1 \right\|_{2,\boldsymbol{\pi}}$ in the worst case is bounded above by $\sqrt{1/\pi_{\min}}$ where $\pi_{\min} = \min_i \pi_i$. Therefore, for any $\delta > 0$, we have $\left\| \frac{\boldsymbol{\mu}(\tau)}{\boldsymbol{\pi}} - 1 \right\|_{2,\boldsymbol{\pi}} \leq \delta$ for any τ such that

$$\tau \geq \frac{\log 1/\pi_{\min} + \log 1/\delta}{\log 1/\|P\|} = O\left(\frac{\log 1/\pi_{\min} + \log 1/\delta}{1 - \|P\|}\right).$$

This suggests that the ‘‘mixing time,’’ i.e., time to reach (close to) stationary distribution of the Markov chain scales inversely with $1 - \|P\|$. Therefore, we will define the ‘‘mixing time’’ of a Markov chain with transition matrix P as $1/(1 - \|P\|)$. This also suggests that in order to bound the distance between a Markov chain's distribution after some steps and its stationary distribution, it is sufficient to obtain a bound on $\|P\|$.

B. CSMA Markov Chain and Its Mixing Time

The *backbone* of our algorithms, for scheduling and congestion control, is a Markov chain with state space being $\mathcal{I}(G)$, where $\mathcal{I}(G)$ is the set of independent sets of G as defined in (1). In recent years, this was considered in the context of CSMA by Wang and Kar [66]. Its transition matrix is determined by the vector $\mathbf{r}(\cdot)$ and hence is time varying. However, if $\mathbf{r}(\cdot)$ were fixed, then it would be a time-homogeneous reversible Markov chain. In the context of CSMA, the vector of $\mathbf{r}(\cdot)$ corresponds to the *aggressiveness* of backoff. In what follows, we will describe this time-homogeneous version [i.e., assuming fixed $\mathbf{r}(\cdot)$] of Markov chain, which was implicit in the description of the scheduling/congestion control algorithm, its stationary distribution, and a bound on its mixing time.

To this end, let $\mathbf{r}(\cdot) = \mathbf{r}$ be fixed. Recall that, under scheduling/congestion control algorithm, each node does the following. Each node i is either in ‘‘transmission’’ state (denoted as $\sigma_i = 1$) or ‘‘waiting’’ state (denoted by $\sigma_i = 0$). In

a waiting state, node has an exponential clock ticking at rate $R_i = \exp(r_i)$ (mean $1/R_i$): when it ticks, if medium is free, it acquires and starts transmitting (i.e., now $\sigma_i = 1$); else if medium is busy, it continues the waiting state (i.e., retains $\sigma_i = 0$). In a transmission state, node has an exponential clock ticking at rate 1: when it ticks, it frees the medium and enters waiting state (i.e., now $\sigma_i = 0$).

This is a continuous-time Markov chain over a finite state space. It can be easily checked that it has the following *product form* stationary distribution $\pi^{\mathbf{r}} = [\pi_{\sigma}^{\mathbf{r}}]$: for any $\sigma \in \mathcal{I}(G)$

$$\begin{aligned} \pi_{\sigma}^{\mathbf{r}} &\propto \exp(\sigma \cdot \mathbf{r}) \\ &= \frac{\exp(\sigma \cdot \mathbf{r})}{\sum_{\sigma \in \mathcal{I}(G)} \exp(\sigma \cdot \mathbf{r})}. \end{aligned} \quad (16)$$

Here, for vectors \mathbf{a}, \mathbf{b} , we use notation of dot product $\mathbf{a} \cdot \mathbf{b} = \sum_{i=1}^n a_i b_i$. Under this stationary distribution, the average fraction of time node i ends up transmitting, which is its “service rate,” is given by

$$s_i(\mathbf{r}) = \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_i] = \sum_{\sigma \in \mathcal{I}(G)} \sigma_i \cdot \pi_{\sigma}^{\mathbf{r}}. \quad (17)$$

Throughout, we will call $\mathbf{s}(\mathbf{r}) = [s_i(\mathbf{r})]$ as the service rate vector induced by $\pi^{\mathbf{r}}$. To understand the “mixing time” of this continuous-time Markov chain, first consider its following discrete-time version with transition matrix P on $\mathcal{I}(G)$. Under P , the transition from current state $\sigma \in \mathcal{I}(G)$ to the next state $\sigma^* \in \mathcal{I}(G)$ happens as follows.

- Choose a node $i \in V$ with probability $\frac{\max\{\exp(r_i), 1\}}{\sum_{k \in V} \max\{\exp(r_k), 1\}}$.
- If $\sigma_i = 1$ (equivalently, $i \in \sigma$), then

$$\sigma_i^* = \begin{cases} 0, & \text{with probability } \min\{1/\exp(r_i), 1\} \\ 1, & \text{otherwise} \end{cases}$$

and $\sigma_j^* = \sigma_j$ for $j \neq i$.

- If $\sigma_i = 0$ and $\sigma_k = 0$ for all $k \in \mathcal{N}(i)$ (i.e., $i \notin \sigma$ and $k \notin \sigma$ for all $k \in \mathcal{N}(i)$), then

$$\sigma_i^* = \begin{cases} 1, & \text{with probability } \min\{\exp(r_i), 1\} \\ 0, & \text{otherwise} \end{cases}$$

and $\sigma_j^* = \sigma_j$ for all $j \neq i$.

- Otherwise $\sigma^* = \sigma$.

The above discrete version of the continuous-time Markov chain is reversible, i.e., $P = P^*$. It can be checked that P is indeed the discretized version, i.e., $\pi^{\mathbf{r}}$ is stationary distribution of P .

The continuous-time Markov chain relates to the above described discrete-time Markov chain with transition matrix P as follows: think of continuous-time Markov chain making its transitions when a clock of net rate $R = \sum_{k \in V} \max\{\exp(r_k), 1\}$ ticks. When its clock ticks, the next state for transition is chosen as per transition matrix P . Given this, let $\mu(t)$ be the distribution over $\mathcal{I}(G)$ under the continuous CSMA Markov chain at time t . Then, the dynamics of $\mu(\cdot)$ is described as

$$\begin{aligned} \mu(t) &= \sum_{i=0}^{\infty} \Pr(\zeta = i) \mu(0) P^i \\ &= \frac{1}{e^{Rt}} \mu(0) e^{RtP} \\ &= \mu(0) e^{Rt(P-I)} \end{aligned} \quad (18)$$

where ζ is Poisson random variable with parameter Rt , which is equal to the number of clock ticks in time $[0, t]$.

Given (18) and earlier discussion on matrix norms, mixing time analysis for discrete-time Markov chain, we obtain that

$$\left\| \frac{\mu(t)}{\pi^{\mathbf{r}}} - 1 \right\|_{2, \pi^{\mathbf{r}}} < \left\| e^{Rt(P-I)} \right\| \left\| \frac{\mu(0)}{\pi^{\mathbf{r}}} - 1 \right\|_{2, \pi^{\mathbf{r}}}.$$

Therefore, to bound the distance between $\mu(t)$ and $\pi^{\mathbf{r}}$, we need to get a bound on $\|e^{Rt(P-I)}\|$.

Lemma 5: The matrix norm of $e^{Rt(P-I)}$ is bounded as

$$\left\| e^{Rt(P-I)} \right\| \leq \left(1 - \frac{1}{\exp(\Theta(n\|\mathbf{r}\|_{\infty} + n))} \right)^{\lfloor t \rfloor}.$$

Proof: Define partition function or normalization constant $Z(\mathbf{r})$ of $\pi^{\mathbf{r}}$ as

$$Z(\mathbf{r}) = \sum_{\sigma \in \mathcal{I}(G)} \exp(\sigma \cdot \mathbf{r}).$$

It follows that

$$Z(\mathbf{r}) \leq |\mathcal{I}(G)| \exp(n\|\mathbf{r}\|_{\infty}) \leq \exp(n(1 + \|\mathbf{r}\|_{\infty})).$$

Therefore, for any $\sigma \in \mathcal{I}(G)$

$$\begin{aligned} \pi_{\sigma}^{\mathbf{r}} &= \frac{1}{Z(\mathbf{r})} \exp(\sigma \cdot \mathbf{r}) \\ &\geq \exp(-n(1 + 2\|\mathbf{r}\|_{\infty})). \end{aligned} \quad (19)$$

Now for any $\sigma, \rho \in \mathcal{I}(G)$ such that they differ in only one component, i.e., it is possible to transit from σ to ρ and *vice versa* in one step, we have

$$\begin{aligned} \left(e^{R(P-I)} \right)_{\sigma\rho} &\geq \Pr(\zeta = 1) P_{\sigma\rho} \\ &\geq \exp(-\Theta(n\|\mathbf{r}\|_{\infty} + n)). \end{aligned}$$

In above, we used the fact that ζ is Poisson random variable with parameter $\sum_i \max\{\exp(r_i), 1\}$, which is at most $n(1 + \exp(\|\mathbf{r}\|_{\infty}))$.

Given above calculations, we are ready to bound the conductance Φ of $W = e^{R(P-I)}$ defined as

$$\begin{aligned} \Phi &= \min_{S \subset \mathcal{I}(G)} \frac{Q(S, \mathcal{I}(G) \setminus S)}{\pi^{\mathbf{r}}(S) \pi^{\mathbf{r}}(\mathcal{I}(G) \setminus S)} \\ &> \min_{\sigma, \rho \in \mathcal{I}(G)} \pi_{\sigma}^{\mathbf{r}} W_{\sigma\rho} \\ &> \exp(-\Theta(n\|\mathbf{r}\|_{\infty} + n)) \end{aligned} \quad (20)$$

where $Q(A, B) = \sum_{\sigma \in A, \rho \in B} \pi_{\sigma}^{\mathbf{r}} W_{\sigma\rho}$. By Cheeger’s inequality [27], [12], [38], [47], it is well known that⁵

$$\begin{aligned} \lambda_{\max} &\leq 1 - \frac{\Phi^2}{2} \\ &< 1 - \exp(-\Theta(n\|\mathbf{r}\|_{\max} + n)). \end{aligned} \quad (21)$$

⁵Cheeger’s inequality is about the second largest eigenvalue which is not equal to λ_{\max} in general. By adding self loop of probability $1/2$, the resulting “lazy” version of the Markov chain has mixing time that is at most constant factor larger compared to that of the original Markov chain; it has all eigenvalues nonnegative and hence the second largest eigenvalue equals λ_{\max} . In the subsequent use of bound on mixing time in this paper, the characterization of mixing time up to constant factor is sufficient, and the use of (21) is justified.

where $|r|_{\max} := \|\mathbf{r}\|_{\infty}$.

Hence, from the properties P3) and P4) of the matrix norm, we can conclude that

$$\begin{aligned} \|e^{Rt(P-I)}\| &\leq \|e^{R(P-I)}\|^{\lfloor t \rfloor} \\ &\leq \lambda_{\max}^{\lfloor t \rfloor} \\ &< (1 - \exp(-\Theta(n|r|_{\max} + n)))^{\lfloor t \rfloor}. \end{aligned} \quad (22)$$

□

Using Lemma 5 and the fact that $\|\frac{\boldsymbol{\mu}^{(0)}}{\boldsymbol{\pi}^r} - 1\|_{2, \boldsymbol{\pi}^r} < \sqrt{1/\boldsymbol{\pi}_{\min}} < \exp(\Theta(n\|\mathbf{r}\|_{\infty} + n))$, we obtain

$$\left\| \frac{\boldsymbol{\mu}^{(t)}}{\boldsymbol{\pi}^r} - 1 \right\|_{2, \boldsymbol{\pi}^r} < \delta, \quad \text{for } t > \exp(\Theta(n\|\mathbf{r}\|_{\infty} + n)) \log \frac{1}{\delta}. \quad (23)$$

C. Positive Harris Recurrence and Its Implication

For completeness, we define the well-known notion of positive Harris recurrence (e.g., see [10] and [11]). We also state its useful implications to explain its desirability. In this paper, we will be concerned with discrete-time, time-homogeneous Markov process or chain evolving over a complete, separable metric space or Polish space X . Let \mathcal{B}_{X} denote the Borel σ -algebra on X . We assume that the space X is endowed with a norm,⁶ denoted by $|\cdot|$. Let $X(\tau)$ denote the state of Markov chain at time $\tau \in \mathbb{Z}_+$.

Consider any $A \in \mathcal{B}_{\mathsf{X}}$. Define stopping time $T_A = \inf\{\tau \geq 1 : X(\tau) \in A\}$. Then, the set A is called Harris recurrent if

$$\Pr_{\mathbf{x}}(T_A < \infty) = 1, \quad \text{for any } \mathbf{x} \in \mathsf{X}$$

where $\Pr_{\mathbf{x}}(\cdot) \equiv \Pr(\cdot | X(0) = \mathbf{x})$. A Markov chain is called Harris recurrent if there exists a σ -finite measure μ on $(\mathsf{X}, \mathcal{B}_{\mathsf{X}})$ such that whenever $\mu(A) > 0$ for $A \in \mathcal{B}_{\mathsf{X}}$, A is Harris recurrent. It is well known that if X is Harris recurrent then an essentially unique invariant measure exists (e.g., see [18]). If the invariant measure is finite, then it may be normalized to obtain a unique invariant probability measure (or stationary probability distribution); in this case, X is called positive Harris recurrent.

Now we describe a useful implication of positive Harris recurrence. Let π be the unique invariant (or stationary) probability distribution of the positive Harris recurrent Markov chain X . Then, the following ergodic property is satisfied: for any $x \in \mathsf{X}$ and nonnegative measurable function $f : \mathsf{X} \rightarrow \mathbb{R}_+$

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{\tau=0}^{T-1} f(X(\tau)) \rightarrow \mathbb{E}_{\pi}[f], \quad \Pr_{\mathbf{x}}\text{-almost surely.}$$

Here $\mathbb{E}_{\pi}[f] = \int f(z)\pi(z)$. Note that $\mathbb{E}_{\pi}[f]$ may not be finite.

1) *A Criterion for Positive Harris Recurrence:* Here we introduce a well-known criterion for establishing the positive Harris recurrence based on existence of a Lyapunov function and an appropriate petite set.

We will need some definitions to begin with. Given a probability distribution (also called sampling distribution) a on \mathbb{N} ,

⁶One may assume it to be induced by the metric of X , denoted by d . For example, for any $\mathbf{x} \in \mathsf{X}$, $|\mathbf{x}| = d(\mathbf{0}, \mathbf{x})$ with respect to a fixed $\mathbf{0} \in \mathsf{X}$.

the a -sampled transition matrix of the Markov chain, denoted by K_a , is defined as

$$K_a(\mathbf{x}, B) = \sum_{\tau \geq 0} a(\tau) P^{\tau}(\mathbf{x}, B), \quad \text{for any } \mathbf{x} \in \mathsf{X}, B \in \mathcal{B}_{\mathsf{X}}.$$

Now, we define a notion of a *petite* set. A nonempty set $A \in \mathcal{B}_{\mathsf{X}}$ is called μ_a -petite if μ_a is a nontrivial measure on $(\mathsf{X}, \mathcal{B}_{\mathsf{X}})$ and a is a probability distribution on \mathbb{N} such that for any $\mathbf{x} \in A$

$$K_a(\mathbf{x}, \cdot) \geq \mu_a(\cdot).$$

A set is called a *petite* set if it is μ_a -petite for some such non-trivial measure μ_a , i.e., there exists a measurable set A' so that $\mu_a(A') > 0$. A known sufficient condition to establish positive Harris recurrence of a Markov chain is to establish positive Harris recurrence of closed petite sets as stated in the following lemma. We refer an interested reader to the book by Meyn and Tweedie [44] or the recent survey by Foss and Konstantopoulos [16] for details.

Theorem 6: Let B be a closed petite set. Suppose B is recurrent, i.e.,

$$\Pr(T_B < \infty | X(0) = \mathbf{x}) = 1, \quad \text{for any } \mathbf{x} \in \mathsf{X}$$

where $T_B = \inf\{\tau \geq 1 : X(\tau) \in B\}$. Further, let

$$\sup_{\mathbf{x} \in B} \mathbb{E}_{\mathbf{x}}[T_B] < \infty.$$

Then, the Markov chain is positive Harris recurrent. In the above, $\mathbb{E}_{\mathbf{x}}[\cdot]$ is an expectation condition on the initial state being \mathbf{x} .

Theorem 6 suggests that to establish the positive Harris recurrence of the network Markov chain, it is sufficient to find a closed petite set that satisfies the conditions of Theorem 6. To establish recurrence property of a set, the following Lyapunov and Foster's criteria will be useful.

Lemma 7: Let there exist functions $h, g : \mathsf{X} \rightarrow \mathbb{R}_+$ and $L : \mathsf{X} \rightarrow \mathbb{R}_+$ such that for any $\mathbf{x} \in \mathsf{X}$

$$\mathbb{E}[L(X(g(\mathbf{x}))) - L(X(0)) | X(0) = \mathbf{x}] \leq -h(\mathbf{x})$$

and

- a) $\inf_{\mathbf{x} \in \mathsf{X}} h(\mathbf{x}) > -\infty$;
- b) $\liminf_{L(\mathbf{x}) \rightarrow \infty} h(\mathbf{x}) > 0$;
- c) $\sup_{L(\mathbf{x}) \leq \gamma} g(\mathbf{x}) < \infty$ for all $\gamma > 0$;
- d) $\limsup_{L(\mathbf{x}) \rightarrow \infty} g(\mathbf{x})/h(\mathbf{x}) < \infty$.

Then, there exists finite $\kappa > 0$ so that the set $B_{\kappa} = \{\mathbf{x} : L(\mathbf{x}) \leq \kappa\}$, and the following holds:

$$\begin{aligned} \mathbb{E}_{\mathbf{x}}[T_{B_{\kappa}}] &< \infty, \quad \text{for any } \mathbf{x} \in \mathsf{X} \\ \sup_{\mathbf{x} \in B_{\kappa}} \mathbb{E}_{\mathbf{x}}[T_{B_{\kappa}}] &< \infty. \end{aligned}$$

V. THROUGHPUT PROPERTY OF SCHEDULING ALGORITHMS

This section establishes throughput optimality for the two scheduling algorithms proposed in Section III-A. Specifically, we present proof of Theorem 1 to establish rate stability of the Scheduling Algorithm 1 in Section V-A and proof of Theorem 2

to establish positive Harris recurrence of the Scheduling Algorithm 2 in Section V-B. As noted earlier, Algorithm 1 does not utilize any global information while Algorithm 2 utilizes only global information in terms of number of nodes in the network.

A. Proof of Theorem 1: Rate Stability

The proof of Theorem 1 consists of three parts. First, we introduce and study a relevant optimization problem whose parameters are the vector of backoff parameters $\mathbf{r}(\cdot)$. On the one hand, it is related to the classical variational principle studied in the context of Gibbs distributions or Markov random fields (e.g., [17, ch. 15.4]). On the other hand, it will suggest that the optimal solution corresponding to $\mathbf{r}(\cdot)$, say \mathbf{r}^* , will be such that the service rate vector $\mathbf{s}(\mathbf{r}^*)$, induced by the Markov chain's stationary distribution, is the same as the arrival rate vector $\boldsymbol{\lambda}$. Therefore, if Algorithm 1 adjusts the $\mathbf{r}(\cdot)$ appropriately so that $\mathbf{r}(\cdot)$ converges to \mathbf{r}^* , then there is a possibility establishing rate stability. In the second part, we do so by showing that Algorithm 1 is a stochastic gradient algorithm for the optimization problem of interest. Finally, in the third part, we conclude the proof of Theorem 1 by establishing that the system is rate stable for any $\boldsymbol{\lambda} \in \Lambda^\circ$.

A Relevant Optimization Problem and Its Properties. We begin by introducing the optimization problem of interest. Its relation to variational principle will be alluded to later. To this, given an arrival rate vector $\boldsymbol{\lambda} \in \Lambda^\circ$ and $\mathbf{r} \in \mathbb{R}^n$, define function $F(\mathbf{r}, \boldsymbol{\lambda})$, where

$$F(\mathbf{r}, \boldsymbol{\lambda}) = \boldsymbol{\lambda} \cdot \mathbf{r} - \log \left(\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r}) \right). \quad (24)$$

The interpretation of $F(\mathbf{r}, \boldsymbol{\lambda})$ is as follows. Assume that $\boldsymbol{\lambda}$ is strictly feasible, i.e., $\boldsymbol{\lambda} \in \Lambda^\circ$, so that it can be written as a positive combination of feasible transmission vectors. That is

$$\boldsymbol{\lambda} = \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \nu_{\boldsymbol{\sigma}} \boldsymbol{\sigma}$$

for $\nu = [\nu_{\boldsymbol{\sigma}}] \in \mathbb{R}_+^{|\mathcal{I}(G)|}$. Therefore, if $\boldsymbol{\sigma} \in \mathcal{I}(G)$ is scheduled for $\nu_{\boldsymbol{\sigma}}$ fraction of the time, then effective service rate is the same as arrive rate $\boldsymbol{\lambda}$. Clearly, ν can be thought of as a probability distribution on $\mathcal{I}(G)$ as well.

Now consider the KL-divergence or relative entropy between this distribution ν and $\pi^{\mathbf{r}}$, the stationary distribution of CSMA Markov chain with parameters \mathbf{r} , defined as follows:

$$d(\nu, \pi^{\mathbf{r}}) = \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \nu_{\boldsymbol{\sigma}} \log \left(\frac{\nu_{\boldsymbol{\sigma}}}{\pi_{\boldsymbol{\sigma}}^{\mathbf{r}}} \right).$$

It is well known (by Pinsker's inequality) that

$$\|\nu - \pi^{\mathbf{r}}\|_{TV}^2 / 2 \leq d(\nu, \pi^{\mathbf{r}}).$$

However, $d(\cdot, \cdot)$ is not a metric and it is only premetric. Consider the following relation between $F(\mathbf{r}, \boldsymbol{\lambda})$ and $d(\nu, \pi^{\mathbf{r}})$

$$F(\mathbf{r}, \boldsymbol{\lambda}) = \left(\sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \nu_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot \mathbf{r} \right) - \log \left(\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r}) \right)$$

$$\begin{aligned} &= \sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \nu_{\boldsymbol{\rho}} \log \left(\frac{\exp(\boldsymbol{\rho} \cdot \mathbf{r})}{\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r})} \right) \\ &= \sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \nu_{\boldsymbol{\rho}} \log \left(\frac{\pi_{\boldsymbol{\rho}}^{\mathbf{r}}}{\nu_{\boldsymbol{\rho}}} \right) + \sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \nu_{\boldsymbol{\rho}} \log \nu_{\boldsymbol{\rho}} \\ &= -d(\nu, \pi^{\mathbf{r}}) - H_{ER}(\nu) \end{aligned} \quad (25)$$

where $H_{ER}(\nu) := -\sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \nu_{\boldsymbol{\rho}} \log \nu_{\boldsymbol{\rho}}$ denotes the entropy of the distribution ν .

Thus, for a given fixed $\boldsymbol{\lambda}$, we have that

$$d(\nu, \pi^{\mathbf{r}}) + F(\mathbf{r}, \boldsymbol{\lambda}) = \text{constant}.$$

Therefore, minimizing $d(\nu, \pi^{\mathbf{r}})$ with respect to parameter \mathbf{r} is equivalent to maximizing $F(\mathbf{r}, \boldsymbol{\lambda})$. As we will show, this optimization of \mathbf{r} leads to \mathbf{r}^* so that the $\mathbf{s}(\mathbf{r}^*)$ equals $\boldsymbol{\lambda}$ as long as $\boldsymbol{\lambda} \in \Lambda^\circ$. For this reason, the following is the optimization problem of interest:

$$\begin{aligned} &\text{maximize} && F(\mathbf{r}, \boldsymbol{\lambda}) \\ &\text{subject to} && \mathbf{r} \in \mathbb{R}^n. \end{aligned} \quad (26)$$

Now we state the following useful properties of this optimization problem.

Lemma 8: Consider a given $\boldsymbol{\lambda} \in \mathbb{R}_+^n$. Then, the following holds.

- 1) The objective function $F(\mathbf{r}, \boldsymbol{\lambda})$, as a function of \mathbf{r} is strictly concave. Moreover

$$\frac{\partial F(\mathbf{r}, \boldsymbol{\lambda})}{\partial r_i} = \lambda_i - s_i(\mathbf{r}) \quad (27)$$

and

$$\frac{\partial^2 F}{\partial r_i \partial r_j} = \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_i \sigma_j] - \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_i] \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_j]. \quad (28)$$

- 2) For $\boldsymbol{\lambda} \in \Lambda^\circ$, the optimization problem (26) has a unique solution $\mathbf{r}^*(\boldsymbol{\lambda})$ that is attained and $F(\mathbf{r}^*(\boldsymbol{\lambda})) < 0$. Let $\mathbf{r} = \mathbf{r}^*(\boldsymbol{\lambda})$. Then, under $\pi^{\mathbf{r}}$ the "service rate vector" [as defined in (17)] $\mathbf{s}(\mathbf{r}) = [s_i(\mathbf{r})]$ equals $\boldsymbol{\lambda}$. That is

$$\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \sigma_i \pi_{\boldsymbol{\sigma}}^{\mathbf{r}} = \lambda_i, \quad \text{for all } i.$$

- 3) Further, for any $\varepsilon > 0$ such that $\boldsymbol{\lambda} + \varepsilon \mathbf{1} \in \Lambda$

$$\|\mathbf{r}^*(\boldsymbol{\lambda})\|_\infty \leq \frac{\log |\mathcal{I}(G)|}{\min\{\varepsilon, \lambda_{\min}\}}.$$

Proof: For simplicity of notation, we will drop the reference to $\boldsymbol{\lambda}$ in $F(\mathbf{r}, \boldsymbol{\lambda})$ and simply denote it as $F(\mathbf{r})$ as we have $\boldsymbol{\lambda}$ fixed throughout the proof. We will use additional notation of the partition function $Z(\mathbf{r})$ of $\pi^{\mathbf{r}}$ defined as

$$Z(\mathbf{r}) = \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r}).$$

Proof of 1). We wish to establish that $F(\mathbf{r})$ is strictly concave as a function of \mathbf{r} . To this end, its first derivative can be calculated as

$$\begin{aligned}\frac{\partial F}{\partial r_i} &= \lambda_i - \frac{\sum_{\sigma \in \mathcal{I}(G)} \sigma_i \exp(\sigma \cdot \mathbf{r})}{\sum_{\rho \in \mathcal{I}(G)} \exp(\rho \cdot \mathbf{r})} \\ &= \lambda_i - \sum_{\sigma \in \mathcal{I}(G)} \sigma_i \pi_{\sigma}^{\mathbf{r}} \\ &= \lambda_i - \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_i] = \lambda_i - s_i(\mathbf{r}).\end{aligned}\quad (29)$$

Here, we have used the definition of $\pi^{\mathbf{r}}$ in (16).

To obtain strict concavity, we would like to show that the Hessian of F is negative definite. Now, we compute the second derivative as [using (29)]

$$\begin{aligned}-\frac{\partial^2 F}{\partial r_i \partial r_j} &= \frac{\partial}{\partial r_j} \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_i] \\ &= \frac{\partial}{\partial r_j} \left(\sum_{\sigma \in \mathcal{I}(G)} \sigma_i \exp(\sigma \cdot \mathbf{r}) \frac{1}{Z(\mathbf{r})} \right) \\ &= \sum_{\sigma \in \mathcal{I}(G)} \left(\sigma_i \sigma_j \exp(\sigma \cdot \mathbf{r}) \frac{1}{Z(\mathbf{r})} \right) \\ &\quad + \sum_{\sigma \in \mathcal{I}(G)} \left(\sigma_i \exp(\sigma \cdot \mathbf{r}) \frac{\partial}{\partial r_j} \frac{1}{Z(\mathbf{r})} \right) \\ &= \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_i \sigma_j] \\ &\quad - \sum_{\sigma \in \mathcal{I}(G)} \sigma_i \exp(\sigma \cdot \mathbf{r}) \frac{1}{Z(\mathbf{r})^2} \left(\sum_{\rho \in \mathcal{I}(G)} \rho_j \exp(\rho \cdot \mathbf{r}) \right) \\ &= \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_i \sigma_j] \\ &\quad - \left(\sum_{\sigma \in \mathcal{I}(G)} \sigma_i \frac{\exp(\sigma \cdot \mathbf{r})}{Z(\mathbf{r})} \right) \times \left(\sum_{\rho \in \mathcal{I}(G)} \rho_j \frac{\exp(\rho \cdot \mathbf{r})}{Z(\mathbf{r})} \right) \\ &= \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_i \sigma_j] - \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_i] \mathbb{E}_{\pi^{\mathbf{r}}}[\sigma_j].\end{aligned}\quad (30)$$

Thus, the Hessian of F , denoted by $M = [M_{ij}]$ with $M_{ij} = \frac{\partial^2 F}{\partial r_i \partial r_j}$, is the negative covariance matrix of a random vector with distribution $\pi^{\mathbf{r}}$. It is well known that covariance matrices are positive semidefinite, i.e., M is negative semidefinite. For strict concavity of F , we need to show that M is negative definite or the covariance matrix of $\pi^{\mathbf{r}}$ is positive definite. To this end, let \mathbf{X} be a vector (of n binary) random variables with the joint distribution $\pi^{\mathbf{r}}$. Let $\boldsymbol{\mu} = \mathbb{E}[\mathbf{X}] \in \mathbb{R}_+^n$ be the vector of its mean. Then, from the above, we have that $-M = \mathbb{E}[(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^T]$. Now consider any vector $\zeta \in \mathbb{R}^n$. To establish the positive definiteness of $-M$, we need to show that

$$\zeta^T (-M) \zeta > 0 \Leftrightarrow \zeta \neq \mathbf{0}.$$

Suppose to the contrary that there exists a vector $\zeta \neq \mathbf{0}$ such that $\zeta^T (-M) \zeta = 0$. Clearly

$$\begin{aligned}\zeta^T (-M) \zeta &= \zeta^T \mathbb{E}[(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^T] \zeta \\ &= \mathbb{E}[\zeta^T (\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^T \zeta] \geq 0.\end{aligned}$$

Therefore, let us assume that

$$\zeta^T (-M) \zeta = \mathbb{E}[\zeta^T (\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^T \zeta] = 0.$$

That is, the random variable $\zeta^T (\mathbf{X} - \boldsymbol{\mu}) = 0$ with probability 1 with respect to $\pi^{\mathbf{r}}$. Now consider n vectors e_1, \dots, e_n , where in e_i only node i is selected; i.e., $e_i \in \{0, 1\}^n$ with i th component 1 and all other components 0. Now, by definition, $\pi^{\mathbf{r}}(e_i) > 0$ for any \mathbf{r} . Therefore, the above condition implies that for all i , $\zeta^T (e_i - \boldsymbol{\mu}) = 0$. That is, for all i

$$\zeta_i (1 - \mu_i) - \sum_{j \neq i} \zeta_j \mu_j = 0 \Rightarrow \zeta_i = \sum_{j=1}^n \zeta_j \mu_j. \quad (31)$$

That is, for all i , $\zeta_i = c$. Now applying the same argument with the choice of $\sigma = \mathbf{0}$, we obtain that

$$c \mathbf{1}^T \boldsymbol{\mu} = 0.$$

This immediately implies that $c = 0$ since $\boldsymbol{\mu}^T \mathbf{1} > 0$ for any \mathbf{r} . Thus, we have proved that if $\zeta^T (-M) \zeta \leq 0$, then it must be that $\zeta = \mathbf{0}$. That is, M is negative definite and hence F is strictly concave. This completes the proof of 1) of Lemma 8.

Proof of 2) and 3). We wish to establish that for $\boldsymbol{\lambda} \in \Lambda^\circ$, the optimization problem has a unique solution that is attained. We will establish this by showing that the optimal solution must lie inside a closed, bounded, and convex set since $\boldsymbol{\lambda} \in \Lambda^\circ$. As a by-product, this will provide (3). Then, the strict concavity of F will immediately lead to the existence of a unique solution, and the claim that $\boldsymbol{\lambda} = \mathbf{s}(\mathbf{r}^*(\boldsymbol{\lambda}))$ as a result of the local optimality condition. As the first step towards this, we establish that $F(\mathbf{r}^*(\boldsymbol{\lambda})) < 0$.

To this end, since $\boldsymbol{\lambda} \in \Lambda^\circ$, it can be easily checked that there exists a distribution ν on $\mathcal{I}(G)$ such that

$$\boldsymbol{\lambda} = \sum_{\sigma \in \mathcal{I}(G)} \nu_{\sigma} \sigma. \quad (32)$$

Therefore, using (32) in the definition of F , we have

$$\begin{aligned}F(\mathbf{r}) &= \sum_{\rho \in \mathcal{I}(G)} \nu_{\rho} \rho \cdot \mathbf{r} - \log \left(\sum_{\sigma \in \mathcal{I}(G)} \exp(\sigma \cdot \mathbf{r}) \right) \\ &= \left(\sum_{\rho \in \mathcal{I}(G)} \nu_{\rho} \log \exp(\rho \cdot \mathbf{r}) \right) - \log \left(\sum_{\sigma \in \mathcal{I}(G)} \exp(\sigma \cdot \mathbf{r}) \right) \\ &= \sum_{\rho \in \mathcal{I}(G)} \nu_{\rho} \log \left(\frac{\exp(\rho \cdot \mathbf{r})}{\sum_{\sigma \in \mathcal{I}(G)} \exp(\sigma \cdot \mathbf{r})} \right) \\ &< 0.\end{aligned}\quad (33)$$

The last step follows because 1) for any $\rho \in \mathcal{I}(G)$, $\exp(\rho \cdot \mathbf{r}) < Z(\mathbf{r})$ since any graph has at least two independent sets; and 2) for some $\rho \in \mathcal{I}(G)$, $\nu_{\rho} > 0$.

Next, we will show that if $\boldsymbol{\lambda} + \varepsilon \mathbf{1} \in \Lambda$, then

$$\sup_{\mathbf{r} \in \mathbb{R}^n} F(\mathbf{r}) = \sup_{\mathbf{r} \in [-K, K]^n} F(\mathbf{r}),$$

$$\text{where } K = \frac{\log |\mathcal{I}(G)|}{\min\{\varepsilon, \boldsymbol{\lambda}_{\min}\}}. \quad (34)$$

To establish (34), we will show that for any $\mathbf{r} \in \mathbb{R}^n$, if a) $\mathbf{r}_{\max} := \max_{1 \leq i \leq n} r_i > K \geq \log |\mathcal{I}(G)| / \varepsilon$ or b) $\mathbf{r}_{\min} := \min_{1 \leq i \leq n} r_i < -K$, then $F(\mathbf{r}) \leq F(\mathbf{0}) = -\log |\mathcal{I}(G)|$. As a by-product, this will imply 3) of Lemma 8.

First, for case a), consider a given \mathbf{r} so that $\mathbf{r}_{\max} > \log |\mathcal{I}(G)|/\varepsilon$. Since $\boldsymbol{\lambda} + \varepsilon \cdot \mathbf{1} \in \boldsymbol{\Lambda}$ and $\boldsymbol{\sigma} \in \{0, 1\}^n$, there exists a nonnegative valued measure ν on $\mathcal{I}(G)$ such that

$$\boldsymbol{\lambda} = \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \nu_{\boldsymbol{\sigma}} \boldsymbol{\sigma} \quad \text{and} \quad \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \nu_{\boldsymbol{\sigma}} \leq 1 - \varepsilon. \quad (35)$$

This implies the existence of a distribution $\hat{\nu}$ on $\mathcal{I}(G)$ defined as

$$\hat{\nu}_{\boldsymbol{\sigma}} = \begin{cases} \nu_{\boldsymbol{\sigma}} + \left(1 - \sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \nu_{\boldsymbol{\rho}}\right), & \text{if } \boldsymbol{\sigma} = \mathbf{0} \\ \nu_{\boldsymbol{\sigma}}, & \text{otherwise.} \end{cases}$$

Note that $\boldsymbol{\lambda} = \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \hat{\nu}_{\boldsymbol{\sigma}} \boldsymbol{\sigma}$. Therefore

$$\begin{aligned} F(\mathbf{r}) &= \boldsymbol{\lambda} \cdot \mathbf{r} - \log \left(\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r}) \right) \\ &= \sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \hat{\nu}_{\boldsymbol{\rho}} \log \frac{\exp(\boldsymbol{\rho} \cdot \mathbf{r})}{\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r})} \\ &\leq \hat{\nu}_{\mathbf{0}} \log \frac{\exp(\mathbf{0} \cdot \mathbf{r})}{\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r})} \\ &\leq \varepsilon \log \frac{1}{\exp(\mathbf{r}_{\max})} \\ &< -\log |\mathcal{I}(G)| \\ &= F(\mathbf{0}) \\ &\leq \sup_{\mathbf{r}} F(\mathbf{r}). \end{aligned} \quad (36)$$

Now, we prove case b). For this, let \mathbf{r} be such that $\mathbf{r}_{\min} < -\log |\mathcal{I}(G)|/\min\{\varepsilon, \boldsymbol{\lambda}_{\min}\}$. Let i be such that $r_i = \mathbf{r}_{\min}$. Define $\bar{\boldsymbol{\lambda}}$ as $\bar{\lambda}_i = \lambda_i - \min\{\varepsilon, \lambda_i\}$ and $\bar{\lambda}_j = \lambda_j$ for $j \neq i$. Clearly, $\bar{\boldsymbol{\lambda}} + \min\{\varepsilon, \lambda_i\} \cdot \mathbf{1} \in \boldsymbol{\Lambda}^o$. Therefore, similar to (35), there exists a nonnegative valued measure ν^1 on $\mathcal{I}(G)$ so that

$$\bar{\boldsymbol{\lambda}} = \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \nu_{\boldsymbol{\sigma}}^1 \boldsymbol{\sigma}, \quad \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \nu_{\boldsymbol{\sigma}}^1 \leq 1 - \min\{\varepsilon, \lambda_i\}. \quad (37)$$

Now define a distribution ν' on $\mathcal{I}(G)$ such that

$$\nu'_{\boldsymbol{\sigma}} = \begin{cases} \nu_{\boldsymbol{\sigma}}^1 + \min\{\varepsilon, \lambda_i\}, & \text{if } \boldsymbol{\sigma} = e_i \\ \nu_{\boldsymbol{\sigma}}^1 + \left(1 - \sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \nu_{\boldsymbol{\rho}}^1\right) - \min\{\varepsilon, \lambda_i\}, & \text{if } \boldsymbol{\sigma} = \mathbf{0} \\ \nu_{\boldsymbol{\sigma}}^1, & \text{otherwise.} \end{cases}$$

Here, as before, e_i refers to the independent set with only node i transmitting. Note that $\boldsymbol{\lambda} = \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \nu'_{\boldsymbol{\sigma}} \boldsymbol{\sigma}$. Now, combined with the fact that $r_i < -\log |\mathcal{I}(G)|/\min\{\varepsilon, \boldsymbol{\lambda}_{\min}\}$, we have

$$\begin{aligned} F(\mathbf{r}) &= \boldsymbol{\lambda} \cdot \mathbf{r} - \log \left(\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r}) \right) \\ &= \sum_{\boldsymbol{\rho} \in \mathcal{I}(G)} \nu'_{\boldsymbol{\rho}} \log \frac{\exp(\boldsymbol{\rho} \cdot \mathbf{r})}{\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r})} \\ &\leq \nu'_{e_i} \log \frac{\exp(e_i \cdot \mathbf{r})}{\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r})} \end{aligned}$$

$$\begin{aligned} &\leq \min\{\varepsilon, \boldsymbol{\lambda}_{\min}\} \log \frac{\exp(r_i)}{\exp(0)} \\ &< -\log |\mathcal{I}(G)| \\ &= F(\mathbf{0}) \\ &\leq \sup_{\mathbf{r}} F(\mathbf{r}). \end{aligned}$$

This completes the proof of b), and subsequently that of Lemma 8. \square

Convergence of $\mathbf{r}(j)$ to $\mathbf{r}^*(\boldsymbol{\lambda})$. The statement of Lemma 8 suggests that if indeed we have algorithm parameter $\mathbf{r}(j) = \mathbf{r}^*(\boldsymbol{\lambda})$, then we have a desirable situation where the effective service rate equals the arrival rate for all nodes as long as $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^o$. To this end, we establish that indeed $\mathbf{r}(j)$ converges to $\mathbf{r}^*(\boldsymbol{\lambda})$ with probability 1. This is because update (4) of Scheduling Algorithm 1 is essentially step of an approximate gradient algorithm for solving optimization problem (26). This is made precise in the proof of Lemma 9.

Lemma 9: If $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^o$, then under Scheduling Algorithm 1

$$\lim_{j \rightarrow \infty} \mathbf{r}(j) = \mathbf{r}^*(\boldsymbol{\lambda}), \quad \text{componentwise, with probability 1.}$$

Proof: First note that the solution $\mathbf{r}^*(\boldsymbol{\lambda})$ of concave (maximization) optimization problem (26) can be found iteratively using the gradient algorithm with appropriate step size. The objective is $F(\mathbf{r}, \boldsymbol{\lambda})$ —we will drop reference to $\boldsymbol{\lambda}$ since it is fixed in what follows and use $F(\mathbf{r})$ instead for $F(\mathbf{r}, \boldsymbol{\lambda})$. Now the i th component of gradient vector of $F(\mathbf{r})$, $\nabla F(\mathbf{r})$, is

$$\frac{\partial F}{\partial r_i}(\mathbf{r}) = \lambda_i - s_i(\mathbf{r}).$$

For a given i , as per (4), the $r_i(\cdot)$ is updated as

$$\begin{aligned} r_i(j+1) &= r_i(j) + \alpha(j)(\hat{\lambda}_i(j) - \hat{s}_i(j)) \\ &= r_i(j) + \frac{1}{j}(\lambda_i - s_i(\mathbf{r}(j)) + e(j)) \\ &= r_i(j) + \frac{1}{j} \left(\frac{\partial F_1}{\partial r_i}(\mathbf{r}(j)) + e(j) \right) \end{aligned}$$

where $e(j) = (\hat{\lambda}_i(j) - \hat{s}_i(j)) - (\lambda_i - s_i(\mathbf{r}(j)))$ captures the ‘‘approximation’’ error in estimating the actual gradient direction given by $\lambda_i - s_i(\mathbf{r}(j))$. Thus, if $e(j) = 0$, then the update of $\mathbf{r}(j)$ is as per the standard gradient algorithm with step size $\alpha(j) = 1/j$. Then, standard arguments from optimization theory would imply that $\mathbf{r}(j) \rightarrow \mathbf{r}^*(\boldsymbol{\lambda})$. But $e(j)$ is a random vector. Therefore, in order to establish the convergence, we will show that norm of $e(j)$ is sufficiently small. Specifically, we establish the following.

Lemma 10: The following bound holds:

$$\begin{aligned} \mathbb{E}[\|e(j)\|_1] &\leq \mathbb{E}[\|\hat{\boldsymbol{\lambda}}(j) - \boldsymbol{\lambda}\|_1 + \|\hat{\mathbf{s}}(j) - \mathbf{s}(\mathbf{r}(j))\|_1] \\ &= O\left(\frac{n}{j^2}\right) \end{aligned} \quad (38)$$

where constant in O -term in the error may depend on n .

The proof of Lemma 10 is stated later in Section V. Now using the bound of (38), we will establish the convergence of $\mathbf{r}(j) \rightarrow \mathbf{r}^*(\boldsymbol{\lambda})$. To this end, consider evolution of $F(\mathbf{r}(\cdot))$. By Taylor's expansion (with notation $\boldsymbol{\delta}(j) = \nabla F(\mathbf{r}(j)) + e(j)$)

$$\begin{aligned}
F(\mathbf{r}(j+1)) &= F\left(\mathbf{r}(j) + \frac{1}{j}[\nabla F(\mathbf{r}(j)) + e(j)]\right) \\
&= F(\mathbf{r}(j)) + \nabla F(\mathbf{r}(j)) \cdot \frac{1}{j}\boldsymbol{\delta}(j) \\
&\quad + \frac{1}{2j^2}\boldsymbol{\delta}(j)^T M \boldsymbol{\delta}(j) \\
&= F(\mathbf{r}(j)) + \frac{1}{j}\|\nabla F(\mathbf{r}(j))\|_2^2 \\
&\quad + \frac{1}{j}\nabla F(\mathbf{r}(j)) \cdot e(j) + \frac{1}{2j^2}\boldsymbol{\delta}(j)^T M \boldsymbol{\delta}(j) \\
&\geq F(\mathbf{r}(j)) + \frac{1}{j}\|\nabla F(\mathbf{r}(j))\|_2^2 \\
&\quad - \frac{\|\nabla F(\mathbf{r}(j))\|_\infty \|e(j)\|_1}{j} - \frac{\|M\|_\infty \|\boldsymbol{\delta}(j)\|_1^2}{2j^2}.
\end{aligned} \tag{39}$$

Here M is an $n \times n$ matrix as per Taylor's expansion is evaluation of second-order partial derivative of F at some values. Therefore, any element of M , say M_{ab} with $1 \leq a, b \leq n$, is bounded as [using calculations executed in (30)]

$$\begin{aligned}
|M_{ab}| &\leq \sup_{\mathbf{r}} \left| \frac{\partial^2 F}{\partial r_a \partial r_b}(\mathbf{r}) \right| \\
&= \sup_{\mathbf{r}} |\mathbb{E}_{\boldsymbol{\pi}^{\mathbf{r}}}[\sigma_a] \mathbb{E}_{\boldsymbol{\pi}^{\mathbf{r}}}[\sigma_b] - \mathbb{E}_{\boldsymbol{\pi}^{\mathbf{r}}}[\sigma_a \sigma_b]| \\
&\leq 1.
\end{aligned} \tag{40}$$

We also note that each component of vectors $\nabla F(\mathbf{r}(j))$ and $e(j)$ is bounded by a constant since the cumulative arrival process is Lipschitz and service process is bounded above by unit rate. Specifically, for any j

$$\begin{aligned}
\|\nabla F(\mathbf{r}(j))\|_\infty &\leq 2, \|e(j)\|_\infty \leq K+1 \\
&\Rightarrow \|\boldsymbol{\delta}(j)\|_\infty \leq K+3.
\end{aligned} \tag{41}$$

Taking expectation on both sides of (39) and using (40), (41), and Lemma 10, for all $j \geq C$

$$\begin{aligned}
&\mathbb{E}[F(\mathbf{r}(j+1))] \\
&\geq \mathbb{E}[F(\mathbf{r}(j))] + \frac{1}{j}\mathbb{E}[\|\nabla F(\mathbf{r}(j))\|_2^2] \\
&\quad - O\left(\mathbb{E}\left[\frac{\|e(j)\|_1}{j}\right]\right) - O\left(\frac{n^2}{j^2}\right) \\
&\geq \mathbb{E}[F(\mathbf{r}(j))] + \frac{1}{j}\mathbb{E}[\|\nabla F(\mathbf{r}(j))\|_2^2] - O\left(\frac{n^2}{j^2}\right).
\end{aligned} \tag{42}$$

Performing summation of (42) from $j = C$ to ∞ , we obtain

$$\begin{aligned}
&\sum_{j=C}^{\infty} \frac{1}{j} \mathbb{E}[\|\nabla F(\mathbf{r}(j))\|_2^2] \\
&\leq O(n^2) - \mathbb{E}[F(\mathbf{r}(C))] + F(\mathbf{r}^*(\boldsymbol{\lambda})) \\
&< \infty
\end{aligned} \tag{43}$$

since $F(\mathbf{r}^*(\boldsymbol{\lambda})) < 0$ from Lemma 8 and by definition of the algorithm and $F(\cdot)$, $\mathbb{E}[F_1(\mathbf{r}(C))] > -\infty$. Now since $\sum_{j=C}^{\infty} 1/j = \infty$, we conclude from (43) that

$$\begin{aligned}
0 &= \liminf_j \mathbb{E}[\|\nabla F(\mathbf{r}(j))\|_2^2] \\
&\geq \mathbb{E}\left[\liminf_j \|\nabla F(\mathbf{r}(j))\|_2^2\right]
\end{aligned} \tag{44}$$

by Fatou's lemma. Therefore, using property of concave maximization, we have that with probability 1

$$\begin{aligned}
\liminf_j \|\nabla F(\mathbf{r}(j))\|_2 &= 0 \\
\Rightarrow \liminf_j \|\mathbf{r}(j) - \mathbf{r}^*(\boldsymbol{\lambda})\|_2 &= 0.
\end{aligned} \tag{45}$$

Thus, in order to complete the proof of Lemma 9, it is enough to show that $\|\mathbf{r}(j) - \mathbf{r}^*(\boldsymbol{\lambda})\|_2$ converges with probability 1. To this end, consider [with notation $\mathbf{r}^* = \mathbf{r}^*(\boldsymbol{\lambda})$, $\boldsymbol{\delta}(j) = \nabla F(\mathbf{r}(j)) + e(j)$]

$$\begin{aligned}
&\|\mathbf{r}(j+1) - \mathbf{r}^*\|_2^2 \\
&= \|\mathbf{r}(j) - \mathbf{r}^*\|_2^2 + \frac{1}{j}(\nabla F(\mathbf{r}(j)) + e(j)) \cdot (\mathbf{r}(j) - \mathbf{r}^*) \\
&\quad + \frac{2\boldsymbol{\delta}(j) \cdot (\mathbf{r}(j) - \mathbf{r}^*)}{j} \\
&\stackrel{(a)}{\leq} \|\mathbf{r}(j) - \mathbf{r}^*\|_2^2 + O\left(\frac{1}{j^2}\right) \\
&\quad + \frac{2\nabla F(\mathbf{r}(j)) \cdot (\mathbf{r}(j) - \mathbf{r}^*) + 2e(j) \cdot (\mathbf{r}(j) - \mathbf{r}^*)}{j} \\
&\stackrel{(b)}{\leq} \|\mathbf{r}(j) - \mathbf{r}^*\|_2^2 + O\left(\frac{1}{j^2}\right) + \frac{2e(j) \cdot (\mathbf{r}(j) - \mathbf{r}^*)}{j} \\
&\stackrel{(c)}{\leq} \|\mathbf{r}(j) - \mathbf{r}^*\|_2^2 + O\left(\frac{1}{j^2}\right) \\
&\quad + O\left(\frac{(\log j + \|\mathbf{r}^*\|_\infty)\|e(j)\|_1}{j}\right) \\
&\stackrel{(d)}{\leq} \|\mathbf{r}(j) - \mathbf{r}^*\|_2^2 + O\left(\frac{1}{j^2}\right) + O\left(\frac{\log j}{j}\|e(j)\|_1\right).
\end{aligned}$$

In the above, (a) follows from (41), (b) follows from the concavity of F , i.e., $\nabla F(\mathbf{r}(j)) \cdot (\mathbf{r}(j) - \mathbf{r}^*) \leq 0$, (c) follows from property of update rule that $\|\mathbf{r}(j)\|_\infty = O(\log j)$, and (d) from Lemma 8 that $\|\mathbf{r}^*\|_\infty = O(1)$. In application of Lemma 10, we have that

$$\sum_{j=1}^{\infty} \mathbb{E}[\|e(j)\|_1] < \infty.$$

Since the terms in the above are nonnegative, by an application of Fubini's theorem and Markov's inequality, we have that with probability 1

$$\sum_{j=1}^{\infty} \|e(j)\|_1 < \infty.$$

Of course, $\sum_j 1/j^2$ is finite. Using this, we have that

$$\|\mathbf{r}(j+1) - \mathbf{r}^*\|_2^2 \leq \|\mathbf{r}(j) - \mathbf{r}^*\|_2^2 + \gamma_j$$

where $\sum_j \gamma_j < \infty$ with probability 1. Now the following (standard) fact from analysis (proof is omitted) implies that $\|\mathbf{r}(j) - \mathbf{r}^*\|$ convergence with probability 1 and completes the proof of Lemma 9.

Proposition 11: Consider two real-valued, nonnegative sequences $x_k, y_k, k \in \mathbb{N}$ such that for each k

$$x_{k+1} \leq x_k + y_k \quad \text{and} \quad \sum_{k=1}^{\infty} y_k < \infty.$$

Then, $\lim_k x_k$ exists. \square

Wrapping Up: Establishing Rate Stability. As an implication of Lemma 9, we establish the rate stability of the queueing network. The following Lemma implies Theorem 1.

Lemma 12: Given $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^o$, under Scheduling Algorithm 1

$$\lim_{t \rightarrow \infty} \frac{Q_i(t)}{t} = 0, \quad \text{for all } 1 \leq i \leq n.$$

Proof: Given $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^o$, recall that $\mathbf{r}^*(\boldsymbol{\lambda})$ is the unique optimal solution of optimization problem (26) as per Lemma 8. In the remainder of the proof, since $\boldsymbol{\lambda}$ is fixed, we will use notation $F(\mathbf{r}) = F(\mathbf{r}, \boldsymbol{\lambda})$, and $\mathbf{r}^* = \mathbf{r}^*(\boldsymbol{\lambda})$ as before. Now by Lemma 9, we have $\mathbf{r}(j) \rightarrow \mathbf{r}^*(\boldsymbol{\lambda})$ with probability 1 as $j \rightarrow \infty$. Now as noted earlier, $\nabla F(\mathbf{r}) = \boldsymbol{\lambda} - \mathbf{s}(\mathbf{r})$. It can be easily checked that $\mathbf{s}(\mathbf{r})$ is continuous as function of \mathbf{r} . Therefore, with probability 1

$$\begin{aligned} \lim_{j \rightarrow \infty} \nabla F(\mathbf{r}(j)) &= \nabla F(\mathbf{r}^*) \\ &= \mathbf{0} \end{aligned} \quad (46)$$

where the equality to $\mathbf{0}$, the vector of all 0's, is implied by Lemma 8. Thus, effectively

$$\lim_{j \rightarrow \infty} \mathbf{s}(\mathbf{r}(j)) = \boldsymbol{\lambda}. \quad (47)$$

Lemma 10 implies that with probability 1

$$\sum_{j=C}^{\infty} \|\hat{\boldsymbol{\lambda}}(j) - \boldsymbol{\lambda}\|_1 + \sum_{j=C}^{\infty} \|\hat{\mathbf{s}}(j) - \mathbf{s}(\mathbf{r}(j))\|_1 < \infty. \quad (48)$$

That is, with probability 1

$$\lim_{j \rightarrow \infty} \hat{\boldsymbol{\lambda}}(j) = \boldsymbol{\lambda}; \quad \lim_{j \rightarrow \infty} \|\hat{\mathbf{s}}(j) - \mathbf{s}(\mathbf{r}(j))\| = 0. \quad (49)$$

From (47) and (49), with probability 1

$$\lim_{j \rightarrow \infty} \|\hat{\boldsymbol{\lambda}}(j) - \hat{\mathbf{s}}(j)\| = 0. \quad (50)$$

Now consider a node i and any time t . Let $t \in [l(j), l(j+1)] = [l(j), l(j) + T(j)]$ for some j . We will bound $Q_i(t)/t$ next. To begin with, note that

$$A_i(0, t) = \sum_{k=0}^{j-2} T(k) \hat{\lambda}_i(k) + A_i(l(j-1), t).$$

Note that the service provided to the i th node in interval $[L(k), L(k+1)]$ is $T(k) \hat{s}_i(k)$. Now, for the purpose of

upper bounding queue, we will assume that this service can be used only to serve the work that has arrived in interval $[L(k-1), L(k)]$. Given this, we obtain the following upper bound (using $Q_i(0) = 0$):

$$\begin{aligned} Q_i(t) &= A_i(0, t) - \int_0^t \sigma_i(r) \mathbf{1}_{\{Q_i(r) > 0\}} dr \\ &\leq \left[\sum_{k=0}^{j-2} \left(T(k) \hat{\lambda}_i(k) - T(k+1) \hat{s}_i(k+1) \right)_+ \right] \\ &\quad + A_i(l(j-1), t). \end{aligned}$$

Here, we have used definition $[x]_+ = \frac{x+|x|}{2}$, the nonnegative part of x , for any $x \in \mathbb{R}$. Since $t \in [l(j), l(j+1)]$ and the cumulative arrival process is Lipschitz, we have

$$\begin{aligned} A_i(l(j-1), t) &\leq A(l(j-1), l(j+1)) \\ &\leq K(l(j+1) - l(j-1)) \\ &= K(T(j-1) + T(j)). \end{aligned}$$

By definition, $T(k) \leq T(k+1)$. Therefore, putting these together, we obtain

$$\begin{aligned} \frac{Q_i(t)}{t} &\leq \frac{1}{l(j)} \left[\sum_{k=0}^{j-2} (T(k) \hat{\lambda}_i(k) - T(k) \hat{s}_i(k+1))_+ \right] \\ &\quad + \frac{K(T(j-1) + T(j))}{l(j)}. \end{aligned} \quad (51)$$

Consider the first term on the right-hand side of (51). From (49) and (50), it follows that $\hat{\lambda}_i(k) - \hat{s}_i(k+1) \rightarrow 0$ as $k \rightarrow \infty$, and $l(j) \geq \sum_{k=0}^{j-2} T(k)$ as well as $l(j) \rightarrow \infty$. Therefore, it easily follows that as $j \rightarrow \infty$, the first term goes to 0. Now, consider the second term on the right-hand side of (51). Since $T(j) = \exp(\sqrt{j})$, $T(j)/l(j) \rightarrow 0$ as $j \rightarrow \infty$. In summary, from this discussion and (51), we obtain that for any i , with probability 1

$$\lim_{t \rightarrow \infty} \frac{Q_i(t)}{t} = 0.$$

This complete the proof of Lemma 12. \square

Proof of Lemma 10: Note that, as per the update (4) of Scheduling Algorithm 1, $\mathbf{r}(j)$ is such that

$$\|\mathbf{r}(j)\|_{\infty} \leq \sum_{k=1}^j \frac{1}{k} = O(\log j).$$

Therefore, the statement of Lemma 10 follows by establishing existence of C so that for $j \geq C$

$$\mathbb{E}[\|\hat{\boldsymbol{\lambda}}(j) - \boldsymbol{\lambda}\|_1 + \|\hat{\mathbf{s}}(j) - \mathbf{s}(\mathbf{r}(j))\|_1 | \mathbf{r}(j)] = O(1/j^2) \quad (52)$$

for $\|\mathbf{r}(j)\|_{\infty} = O(\log j)$. In the remaining proof, for simplicity of notation, we will drop reference $\mathbf{r}(j)$ and simply use $\mathbb{E}[\cdot]$ in place of $\mathbb{E}[\cdot | \mathbf{r}(j)]$. We will establish that by arguing separately that $\mathbb{E}[\|\hat{\boldsymbol{\lambda}}(j) - \boldsymbol{\lambda}\|_1] = O(1/j^2)$ and $\mathbb{E}[\|\hat{\mathbf{s}}(j) - \mathbf{s}(\mathbf{r}(j))\|_1] = O(1/j^2)$.

First, we consider the deviation in $\hat{\boldsymbol{\lambda}}(j)$. This will immediately follow from the property of arrival process. By definition,

$\hat{\lambda}(j)$ is the empirical arrival rate vector over $[l(j), l(j+1))$. Now for any i

$$\begin{aligned}\hat{\lambda}_i(j) &= \frac{1}{T(j)} A_i(l(j), l(j+1)) \\ &= \frac{1}{T(j)} \left(\sum_{k=1}^{T(j)} A_i(l(j) + k - 1, l(j) + k) \right).\end{aligned}\quad (53)$$

Now, $X_k \triangleq A_i(l(j) + k - 1, l(j) + k)$ are i.i.d. random variables with $\mathbb{E}[X_k] = \lambda_i$, bounded support $[0, K]$ and hence standard deviation at most K . Using this, we have

$$\begin{aligned}\mathbb{E}[|\hat{\lambda}_i(j) - \lambda_i|] &= \mathbb{E} \left[\left| \frac{1}{T(j)} \sum_{k=1}^{T(j)} (X_k - \mathbb{E}[X_k]) \right| \right] \\ &\leq \mathbb{E} \left[\left| \frac{1}{T(j)} \sum_{k=1}^{T(j)} (X_k - \mathbb{E}[X_k]) \right|^2 \right]^{1/2} \\ &\leq \left(\frac{1}{T(j)} \mathbb{E}[X_1^2] \right)^{1/2} \\ &\leq \frac{K}{\sqrt{T(j)}} \\ &= O(1/j^2)\end{aligned}\quad (54)$$

where the last inequality follows from $T(j) = \exp(\sqrt{j})$. This completes the proof of bound on deviation for $\hat{\lambda}(j)$.

Now, we consider deviations in $\hat{\mathbf{s}}(j)$ compared to $\mathbf{s}(\mathbf{r}(j))$. For this, first we establish $\mathbb{E}[\hat{\mathbf{s}}(j)]$ being close to $\mathbf{s}(\mathbf{r}(j))$ and then we establish $\hat{\mathbf{s}}(j)$ being close to $\mathbb{E}[\hat{\mathbf{s}}(j)]$. Therefore, we start by evaluating deviation between $\mathbb{E}[\hat{\mathbf{s}}(j)]$ and $\mathbf{s}(\mathbf{r}(j))$. To this end, consider any i . We will establish that

$$|\mathbb{E}[\hat{s}_i(j)] - s_i(\mathbf{r}(j))| = O\left(\frac{1}{j^4}\right).\quad (55)$$

To establish (55), we will use the mixing time bounds (23) derived in Section IV-B next. To this end, let $\boldsymbol{\mu}(t)$ be the distribution over $\mathcal{I}(G)$ of scheduling decisions at time $t \in [l(j), l(j+1))$. By Lemma 8(2), $s_i(\mathbf{r}(j)) = \mathbb{E}_{\boldsymbol{\pi}^{\mathbf{r}(j)}}[\sigma_i]$. And σ_i is 0-1 valued random variable. Therefore

$$\begin{aligned}|\mathbb{E}_{\boldsymbol{\mu}(t)}[\sigma_i] - s_i(\mathbf{r}(j))| &= |\mathbb{E}_{\boldsymbol{\mu}(t)}[\sigma_i] - \mathbb{E}_{\boldsymbol{\pi}^{\mathbf{r}(j)}}[\sigma_i]| \\ &\leq \left| \boldsymbol{\mu}(t) - \boldsymbol{\pi}^{\mathbf{r}(j)} \right|_{TV} \\ &\leq \left\| \frac{\boldsymbol{\mu}(t)}{\boldsymbol{\pi}^{\mathbf{r}(j)}} - 1 \right\|_{2, \boldsymbol{\pi}^{\mathbf{r}(j)}}\end{aligned}\quad (56)$$

where the last inequality follows from (13). Now, from (23), the right-hand side of (56) is bounded above by $O(1/j^4)$ as long as

$$\begin{aligned}t &\geq l(j) + (\exp(\Theta(n|r|_{\max}) + n)) \log j^4 \\ &= l(j) + j^{O(n)} \log j = l(j) + \Gamma(j)\end{aligned}\quad (57)$$

where $\Gamma(j) = j^{O(n)} \log j$. In the above, while applying (23), we have used the fact $\|\mathbf{r}\|_{\infty} = O(\log j)$. This leads to the following bound:

$$\begin{aligned}|\mathbb{E}[\hat{s}_i(j)] - s_i(\mathbf{r}(j))| &= \frac{1}{T(j)} \left| \int_{l(j)}^{l(j+1)} \mathbb{E}_{\boldsymbol{\mu}(t)}[\sigma_i] - \mathbb{E}_{\boldsymbol{\pi}^{\mathbf{r}(j)}}[\sigma_i] dt \right| \\ &\leq \frac{\Gamma(j)}{T(j)} + O\left(\frac{1}{j^4}\right).\end{aligned}\quad (58)$$

Hence, (55) follows since $\Gamma(j)/T(j) = (j^{O(n)} \log j)/T(j) = O(1/j^4)$ due to the choice of $T(j) = \exp(\sqrt{j})$.

Given (55), as the last step to establish $\mathbb{E}[|\hat{\mathbf{s}}(j) - \mathbf{s}(\mathbf{r}(j))|_1] = O(1/j^2)$, we will show that for any i

$$\mathbb{E}[|\hat{s}_i(j) - \mathbb{E}[\hat{s}_i(j)]|] = O\left(\frac{1}{j^2}\right).\quad (59)$$

Consider (with notation $S = [l(j), l(j+1))$), and $\mathbb{E}[\cdot]^2 = (\mathbb{E}[\cdot])^2$

$$\begin{aligned}T(j)^2 \mathbb{E}[|\hat{s}_i(j) - \mathbb{E}[\hat{s}_i(j)]|^2] &= \mathbb{E} \left[\left| \int_{l(j)}^{l(j+1)} \sigma_i(t) dt - \mathbb{E} \left[\int_{l(j)}^{l(j+1)} \sigma_i(t) dt \right] \right|^2 \right] \\ &\leq \mathbb{E} \left[\left(\int_{l(j)}^{l(j+1)} \sigma_i(t) dt - \mathbb{E} \left[\int_{l(j)}^{l(j+1)} \sigma_i(t) dt \right] \right)^2 \right] \\ &= \mathbb{E} \left[\left(\int_{l(j)}^{l(j+1)} \sigma_i(t) dt \right)^2 \right] - \mathbb{E} \left[\int_{l(j)}^{l(j+1)} \sigma_i(t) dt \right]^2 \\ &= \int_{l(j)}^{l(j+1)} \mathbb{E}[\sigma_i(t)] \\ &\quad \times \left(\int_{l(j)}^{l(j+1)} \mathbb{E}[\sigma_i(s) | \sigma_i(t) = 1] - \mathbb{E}[\sigma_i(s)] ds \right) dt \\ &= 2 \int_{l(j)}^{l(j+1)} \mathbb{E}[\sigma_i(t)] \\ &\quad \times \left(\int_t^{l(j+1)} \mathbb{E}[\sigma_i(s) | \sigma_i(t) = 1] - \mathbb{E}[\sigma_i(s)] ds \right) dt \\ &= 2 \int_{l(j)}^{l(j+1)} \mathbb{E}[\sigma_i(t)] \\ &\quad \times \left(\int_t^{t+\Gamma(j)} \mathbb{E}[\sigma_i(s) | \sigma_i(t) = 1] - \mathbb{E}[\sigma_i(s)] ds \right. \\ &\quad \left. + \int_{t+\Gamma(j)}^{l(j+1)} \mathbb{E}[\sigma_i(s) | \sigma_i(t) = 1] - \mathbb{E}[\sigma_i(s)] ds \right) dt \\ &\stackrel{(a)}{\leq} 2 \int_{l(j)}^{l(j+1)} \mathbb{E}[\sigma_i(t)] \\ &\quad \times \left(\int_t^{t+\Gamma(j)} 1 ds + \int_{t+\Gamma(j)}^{l(j+1)} O\left(\frac{1}{j^4}\right) ds \right) dt \\ &\leq 2T(j) \left(\Gamma(j) + O\left(\frac{T(j)}{j^4}\right) \right) \\ &= O\left(\frac{T(j)^2}{j^4}\right).\end{aligned}\quad (60)$$

In the above, (a) follows from choice of $\Gamma(j)$ as in (56) and (57). If $s \geq t + \Gamma(j)$, then due to the ‘‘mixing effect,’’ $\mathbb{E}[\sigma_i(s) | \sigma_i(t)]$, $\mathbb{E}[\sigma_i(s)]$ are within $O(1/j^4)$ of $s_i(\mathbf{r}(j))$. Now, (60) immediately implies that

$$\mathbb{E}[|\hat{s}_i(j) - \mathbb{E}[\hat{s}_i(j)]|] = O\left(\frac{1}{j^2}\right). \quad (61)$$

To conclude, observe that (54), (55), and (61) imply the result of Lemma 10.

B. Proof of Theorem 2: Positive Harris Recurrence

The goal of this section is to prove Theorem 2, that is, the positive Harris recurrence of the network Markov process under Scheduling Algorithm 2. For a countable Markov chain, positive recurrence means that all states are visited infinitely often, with a finite mean intervisit time. When the state space is not countable (as in our case), one cannot expect every state to be visited infinitely often. However, a small set of states can have that property. If the transition probabilities out of that set are similar, then the set plays the role of a recurrent state. Indeed, the evolution essentially starts afresh once the chain hits that set. This idea is made precise by the definition of a petite⁷ set. Section IV-C has a review of known results about establishing positive Harris recurrence. In particular, Theorem 6 there states that the existence of a positive recurrent closed petite set implies positive Harris recurrence.

The appropriate petite set is the set \mathbf{S} where the sum of the squares of the queue lengths is less than some constant κ . The positive recurrence is proved using the fact that the sum of the squares of the queue lengths is a Lyapunov function which tends to decrease when it is larger than κ (Lemma 13). Intuitively, this is true because Scheduling Algorithm 2 tries to balance $\hat{s}_i(j)$ and $\hat{\lambda}_i(j) + \varepsilon$ for all i , so that on average, the service rate dominates the arrival rate on each queue. The set \mathbf{S} is shown to be petite (Lemma 14) by proving that starting from any state in that set, there is some lower bound θ on the probability that, at some later time T_κ , the queues become empty, no link is active, and the parameters \mathbf{r} of the CSMA backoff delays reach their maximum value (Proposition 17). Thus, the evolution of the Markov chain essentially starts afresh from that set with at least probability θ .

To this end, we start with necessary definitions of the network Markov process under Scheduling Algorithm 2. Let $\tau \in \mathbb{N} \cup \{0\}$ be the index for the discrete time. It can be checked that the tuple $X(j) = (\mathbf{Q}(T_j), \mathbf{r}(T_j), \sigma(T_j))$ forms the state of the time-homogeneous Markov chain operating under the algorithm. Now $X(\tau) \in \mathbf{X}$ where $\mathbf{X} = \mathbb{R}_+^n \times [-\frac{n}{\varepsilon}, \frac{n}{\varepsilon}]^n \times \mathcal{I}(G)$. Clearly, \mathbf{X} is a Polish space endowed with the natural product topology. Let $\mathcal{B}_\mathbf{X}$ be the Borel σ -algebra of \mathbf{X} . Finally, for $\mathbf{x} = (\mathbf{Q}, \mathbf{r}, \sigma) \in \mathbf{X}$, define norm of \mathbf{x} denoted by $|\mathbf{x}|$ as

$$|\mathbf{x}| = |\mathbf{Q}| + |\mathbf{r}| + |\sigma|$$

where $|\mathbf{Q}|$ and $|\mathbf{r}|$ denote the ℓ_1 norm, and $|\sigma|$ is σ 's index in $\{0, \dots, |\mathcal{I}(G)| - 1\}$, assigned arbitrarily. Thus, $|\mathbf{r}|, |\sigma|$ are always bounded. Therefore, in essence, $|\mathbf{x}| \rightarrow \infty$ iff $|\mathbf{Q}| \rightarrow \infty$.

⁷Recall that petite means small in French.

To establish statement of Theorem 2, we need to show that $X(\tau)$ is indeed positive Harris recurrent as long as $\boldsymbol{\lambda} + 2\varepsilon\mathbf{1} \in \boldsymbol{\Lambda}$. By Theorem 6, it is sufficient to find positive recurrent closed petite set. First, we will find closed recurrent set using criterion of Lemma 7 and then establish that the set is indeed petite. To this end, define a Lyapunov function $L : \mathbf{X} \rightarrow \mathbb{R}_+$ as

$$L(\mathbf{x}) = \sum_{i=1}^n Q_i^2 \triangleq \mathbf{Q}^2 \cdot \mathbf{1}, \quad \text{where } \mathbf{x} = (\mathbf{Q}, \mathbf{r}, \sigma) \in \mathbf{X}.$$

We establish the following ‘‘drift’’ property about L .

Lemma 13: Given $\boldsymbol{\lambda}$ so that $\boldsymbol{\lambda} + 2\varepsilon\mathbf{1} \in \boldsymbol{\Lambda}$, define

$$N = N(\varepsilon, n) = \left\lceil \frac{48 \times 16 \times 72n^5}{\varepsilon^6} \right\rceil.$$

Then, for any initial state $X(0) = (\mathbf{Q}(0), \mathbf{r}(0), \sigma(0)) \in \mathbf{X}$

$$\mathbb{E}[L(X(N)) - L(X(0)) | X(0)] \leq -h(X(0)) \quad (62)$$

where $h : \mathbf{X} \rightarrow \mathbb{R}$ is defined as

$$h(\mathbf{x}) = \varepsilon TN(\mathbf{Q}(0) \cdot \mathbf{1}) - n(TN)^2(\varepsilon + K^2 + 2K). \quad (63)$$

Therefore, Lemma 7 implies that for some finite $\kappa > 0$, set $B_\kappa = \{\mathbf{x} : L(\mathbf{x}) \leq \kappa\}$ satisfies

$$\begin{aligned} \mathbb{E}_\mathbf{x}[T_{B_\kappa}] &< \infty, & \text{for any } \mathbf{x} \in \mathbf{X} \\ \sup_{\mathbf{x} \in B_\kappa} \mathbb{E}_\mathbf{x}[T_{B_\kappa}] &< \infty. \end{aligned}$$

Therefore, by Theorem 6, the following is sufficient to complete the proof of Theorem 2.

Lemma 14: Consider any $\kappa > 0$. Then, the set $B_\kappa = \{\mathbf{x} : L(\mathbf{x}) \leq \kappa\}$ is a closed petite set.

In the remainder of this section, we will prove Lemmas 13 and 14.

1) Proof of Lemma 13:

A relevant optimization problem. The basic idea behind the update algorithm (6) is to design a simple gradient procedure for solving the following optimization problem:

$$\begin{aligned} &\text{maximize} && F(\mathbf{r}, \boldsymbol{\lambda} + \varepsilon\mathbf{1}) \triangleq F_\varepsilon(\mathbf{r}) \\ &\text{subject to} && \mathbf{r} \in \mathbb{R}^n. \end{aligned} \quad (64)$$

By Lemma 8, it follows that if $\boldsymbol{\lambda} + 2\varepsilon\mathbf{1} \in \boldsymbol{\Lambda}$, then (64) has a unique solution that is attained; let it be $\mathbf{r}^* = \mathbf{r}^*(\boldsymbol{\lambda} + \varepsilon\mathbf{1})$. Then, from Lemma 8(2), the effective service rate $\mathbf{s}(\mathbf{r}^*)$, under the random access algorithm with fixed \mathbf{r}^* , is such that

$$s_i(\mathbf{r}^*) = \lambda_i + \varepsilon.$$

That is, the arrival rate is less than the service rate by $\varepsilon > 0$ under this idealized setup. In order to establish the positive Harris recurrence, we will need more than this—service rate should dominate arrival rate for small enough time interval to imply appropriate drift condition desired by Lyapunov–Foster’s criteria. This is exactly what we will establish next.

Derivative of F_ε Becomes Small. As per statement of Lemma 13, let initial state be $X(0) = (\mathbf{Q}(0), \mathbf{r}(0), \boldsymbol{\sigma}(0))$. As the first step, we wish to establish the following:

$$\begin{aligned} & \frac{1}{N} \sum_{j=1}^N \mathbb{E} [\|\nabla F_\varepsilon(\mathbf{r}(j))\|_2^2] \\ &= \frac{1}{N} \sum_{j=1}^N \mathbb{E} [\|\boldsymbol{\lambda} + \varepsilon \cdot \mathbf{1} - \mathbf{s}(\mathbf{r}(j))\|_2^2] \\ &\leq \frac{\varepsilon^2}{16}. \end{aligned} \quad (65)$$

In the above and everywhere else in the proof of Lemma 13, the expectation is always assumed to be conditioned on the initial state $X(0)$. For simplicity, we will drop reference to this conditioning. Intuitively, (65) implies that on average and in expectation, the arriving rate $\boldsymbol{\lambda}$ is strictly less than the normalized service rate $\mathbf{s}(\mathbf{r}(j))$ within a time interval of length N . This will allow us to establish drift in Lyapunov function. To this end, we start with definition $G(\mathbf{r}) = F_\varepsilon(\mathbf{r}) - \|\mathbf{r} - \mathbf{r}^*\|_2^2$. We establish the following useful nondecreasing property of $G(\cdot)$ under the ‘‘projection’’ defined in (7).

Lemma 15: For any $\mathbf{r} \in [-\frac{n}{\varepsilon}, \frac{n}{\varepsilon}]^n$ and $\Delta \mathbf{r} \in [-1, 1]^n$, $-\frac{16n^3}{\varepsilon^2} \leq G(\mathbf{r}) < 0$ and $G([\mathbf{r} + \Delta \mathbf{r}]_{\frac{n}{\varepsilon}}) \geq G(\mathbf{r} + \Delta \mathbf{r})$.

Proof: $G(\mathbf{r})$ is upper bounded by 0 since $F_\varepsilon(\mathbf{r}) \leq F_\varepsilon(\mathbf{r}^*) < 0$ by Lemma 8. Further

$$\begin{aligned} G(\mathbf{r}) &= F_\varepsilon(\mathbf{r}) - \|\mathbf{r} - \mathbf{r}^*\|_2^2 \\ &\stackrel{(a)}{\geq} (\boldsymbol{\lambda} + \varepsilon \cdot \mathbf{1}) \cdot \mathbf{r} - \log \left(\sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \exp(\boldsymbol{\sigma} \cdot \mathbf{r}) \right) - \frac{4n^3}{\varepsilon^2} \\ &\geq n \cdot \left(-\frac{n}{\varepsilon} \right) - \log(2^n \exp(n\mathbf{r}_{\max})) - \frac{4n^3}{\varepsilon^2} \\ &\geq -\frac{16n^3}{\varepsilon^2}. \end{aligned} \quad (66)$$

Here (a) follows from Lemma 8(3) for $\mathbf{r}^* = \mathbf{r}^*(\boldsymbol{\lambda} + \varepsilon \mathbf{1})$ (thus $\mathbf{r}^* \in [-\frac{n}{\varepsilon}, \frac{n}{\varepsilon}]^n$), and the last step has used $n > 3$. Now if we set $\mathbf{x} = \mathbf{r} + \Delta \mathbf{r}$, $|\mathbf{x}|_{\max} \leq \frac{n}{\varepsilon} + 1$ and we need to show $G([\mathbf{x}]_{\frac{n}{\varepsilon}}) \geq G(\mathbf{x})$. Note that it is enough to show that for any $i \in V$

$$G([\mathbf{x}]_{\frac{n}{\varepsilon}, i}) \geq G(\mathbf{x}) \quad (67)$$

where the i -projection $\bar{\mathbf{x}} = [\mathbf{x}]_{\frac{n}{\varepsilon}, i}$ is defined as

$$\bar{x}_j = \begin{cases} ([\mathbf{x}]_{\frac{n}{\varepsilon}})_j, & \text{if } j = i \\ x_j, & \text{otherwise.} \end{cases}$$

Then, we can iteratively apply (67) to complete the proof. When $x_i \in [-\frac{n}{\varepsilon}, \frac{n}{\varepsilon}]$, desired claim trivially follows as $[\mathbf{x}]_{\frac{n}{\varepsilon}, i} = \mathbf{x}$. Now suppose $x_i \notin [-\frac{n}{\varepsilon}, \frac{n}{\varepsilon}]$. By definition, it must be that $x_i \in (\frac{n}{\varepsilon}, \frac{n}{\varepsilon} + 1]$ or $x_i \in [-\frac{n}{\varepsilon} - 1, -\frac{n}{\varepsilon}]$. We prove (67) when $x_i \in [\frac{n}{\varepsilon}, \frac{n}{\varepsilon} + 1]$; the other arguments for the other case are very similar. Consider

$$\begin{aligned} & G([\mathbf{x}]_{\frac{n}{\varepsilon}, i}) - G(\mathbf{x}) \\ &= F_\varepsilon([\mathbf{x}]_{\frac{n}{\varepsilon}, i}) - F_\varepsilon(\mathbf{x}) - \left(\frac{n}{\varepsilon} - r_i^* \right)^2 + (x_i - r_i^*)^2 \\ &\stackrel{(a)}{\geq} - \left(x_i - \frac{n}{\varepsilon} \right) + \left(x_i - \frac{n}{\varepsilon} \right) \left(x_i + \frac{n}{\varepsilon} - 2r_i^* \right) \end{aligned}$$

$$\begin{aligned} &\stackrel{(b)}{\geq} - \left(x_i - \frac{n}{\varepsilon} \right) + 2 \left(x_i - \frac{n}{\varepsilon} \right) \\ &\geq 0. \end{aligned}$$

In the above, (a) and (b) are due to $|\frac{\partial F_\varepsilon}{\partial r_i}| \leq 1$ and $|\mathbf{r}^*|_{\max} \leq \frac{n \cdot \log(2)}{\varepsilon} < \frac{n}{\varepsilon} - 1$ (since $n > 3$ by assumption), respectively. This completes the proof of Lemma 15. \square

Now consider the relation between $G(\mathbf{r}(j+1))$ and $G(\mathbf{r}(j))$

$$\begin{aligned} G(\mathbf{r}(j+1)) &= G([\mathbf{r}(j) + \alpha(\nabla F_\varepsilon(\mathbf{r}(j)) + e(j))]_{\frac{n}{\varepsilon}}) \\ &\stackrel{(a)}{\geq} G(\mathbf{r}(j) + \alpha(\nabla F_\varepsilon(\mathbf{r}(j)) + e(j))) \\ &= F_\varepsilon(\mathbf{r}(j) + \alpha(\nabla F_\varepsilon(\mathbf{r}(j)) + e(j))) \\ &\quad - \|\mathbf{r}(j) + \alpha(\nabla F_\varepsilon(\mathbf{r}(j)) + e(j)) - \mathbf{r}^*\|_2^2 \\ &= F_\varepsilon(\mathbf{r}(j)) + \nabla F_\varepsilon(\mathbf{r}(j))\alpha(\nabla F_\varepsilon(\mathbf{r}(j)) + e(j)) \\ &\quad + \frac{1}{2}\alpha(\nabla F_\varepsilon(\mathbf{r}(j)) + e(j))M\alpha(\nabla F_\varepsilon(\mathbf{r}(j)) + e(j)) \\ &\quad - \|\mathbf{r}(j) - \mathbf{r}^*\|_2^2 - 2\alpha\nabla F_\varepsilon(\mathbf{r}(j)) \cdot (\mathbf{r}(j) - \mathbf{r}^*) \\ &\quad - 2\alpha e(j)(\mathbf{r}(j) - \mathbf{r}^*) - \alpha^2\|\nabla F_\varepsilon(\mathbf{r}(j)) + e(j)\|_2^2 \\ &\stackrel{(b)}{\geq} F_\varepsilon(\mathbf{r}(j)) + \alpha\|\nabla F_\varepsilon(\mathbf{r}(j))\|_2^2 + \alpha\nabla F_\varepsilon(\mathbf{r}(j))e(j) \\ &\quad - \frac{\alpha^2(K+1)^2n^2}{2} - \|\mathbf{r}(j) - \mathbf{r}^*\|_2^2 \\ &\quad - 2\alpha e(j)(\mathbf{r}(j) - \mathbf{r}^*) - \alpha^2(K+1)^2n^2 \\ &\stackrel{(c)}{\geq} F_\varepsilon(\mathbf{r}(j)) + \alpha\|\nabla F_\varepsilon(\mathbf{r}(j))\|_2^2 - \alpha\|e(j)\|_1 \\ &\quad - \frac{3\alpha^2(K+1)^2n^2}{2} - \|\mathbf{r}(j) - \mathbf{r}^*\|_2^2 - \frac{4\alpha n}{\varepsilon}\|e(j)\|_1 \\ &= G(\mathbf{r}(j)) + \alpha\|\nabla F_\varepsilon(\mathbf{r}(j))\|_2^2 - \alpha \left(1 + \frac{4n}{\varepsilon} \right) \|e(j)\|_1 \\ &\quad - \frac{3\alpha^2(K+1)^2n^2}{2} \\ &\geq G(\mathbf{r}(j)) + \alpha\|\nabla F_\varepsilon(\mathbf{r}(j))\|_2^2 - \frac{5\alpha n}{\varepsilon}\|e(j)\|_1 \\ &\quad - \frac{3\alpha^2(K+1)^2n^2}{2} \end{aligned} \quad (68)$$

where the random vector $e(j) = \hat{\lambda}_i(j) - \hat{s}_i(j) - (\lambda_i - s_i(\mathbf{r}(j)))$; M is the $n \times n$ with $M_{ab} = \partial^2 F_\varepsilon(\tilde{\mathbf{r}}) / \partial r_a \partial r_b$ for some $\tilde{\mathbf{r}}$ in neighborhood of \mathbf{r} with $M_{ab} \in [-1, 1]$ by (40). In the above, (a) follows from the fact that $\alpha \leq (K+1)^{-2}$, $\nabla F_\varepsilon(\mathbf{r}(j)) \in [-1, 1]^n$, $e(j) \in [-1, K]^n$ and Lemma 15.⁸ For (b), we use that $\|M\|_\infty \leq 1$ and the concavity of F_ε and $\nabla F_\varepsilon(\mathbf{r}(j)) + e(j) \in [-2, (K+1)^n]$. Finally, (c) follows from $\nabla F_\varepsilon(\mathbf{r}(j)) \in [-1, 1]^n$ and $\mathbf{r}(j) - \mathbf{r}^* \in [-\frac{2n}{\varepsilon}, \frac{2n}{\varepsilon}]^n$.

Our choice of the large updating period T is merely for bounding $e(j)$ and we obtain the following lemma which is analogous to Lemma 10.

Lemma 16: If the updating period $T \geq \exp(\Theta(\frac{n^2}{\varepsilon} \log \frac{n}{\varepsilon}))$, then for all $j \in \mathbb{N}$

$$\mathbb{E}[\|\hat{\lambda}(j) - \boldsymbol{\lambda}\|_1 + \|\hat{\mathbf{s}}(j) - \mathbf{s}(\mathbf{r}(j))\|_1] \leq \frac{\varepsilon^3}{240n}.$$

⁸This is the main reason why we consider G instead of F_ε as we cannot establish monotonicity of F_ε under the projection.

Therefore, for all $j \in \mathbb{N}$

$$\mathbb{E}[\|e(j)\|_1] \leq \frac{\varepsilon^3}{240n}.$$

Proof: We provide sketch proof here since the proof of Lemma 16 is essentially the same as that of Lemma 10—replace $T(j) = T$, $\alpha(j) = \alpha$ for all j and use $\|\mathbf{r}\|_{\max} \leq \frac{n}{\varepsilon}$ to obtain bound of

$$\exp\left(\Theta\left(\frac{n^2}{\varepsilon}\right)\right)$$

on mixing time of the Markov chain on $\mathcal{I}(G)$ using (5). As a consequence, it follows that by choice of T with large enough constant in its exponent, as stated in Lemma 16, the expectation of $\|e(j)\|_1$ can be made smaller than any given constant. Specifically, it can be made smaller than $\frac{\varepsilon^3}{240n}$. \square

Summing (68) from $j = 1$ to N

$$\begin{aligned} 0 &\geq G(\mathbf{r}(N+1)) \\ &\geq G(\mathbf{r}(1)) + \alpha \left(\sum_{j=1}^N \|\nabla F_\varepsilon(\mathbf{r}(j))\|_2^2 \right) \\ &\quad - \frac{5\alpha n}{\varepsilon} \left(\sum_{j=1}^N \|e(j)\|_1 \right) - \frac{3\alpha^2(K+1)^2 n^2}{2} N. \end{aligned} \quad (69)$$

Taking expectation on both sides and diving by αN

$$\begin{aligned} &\frac{1}{N} \sum_{j=1}^N \mathbb{E}[\|\nabla F_\varepsilon(\mathbf{r}(j))\|_2^2] \\ &\leq -\frac{1}{\alpha N} G(\mathbf{r}(1)) + \frac{5n}{\varepsilon N} \sum_{j=1}^N \mathbb{E}[\|e(j)\|_1] + \frac{3\alpha(K+1)^2 n^2}{2} \\ &\stackrel{(a)}{\leq} \frac{1}{\alpha N} \frac{16n^3}{\varepsilon^2} + \frac{\varepsilon^2}{48} + \frac{\varepsilon^2}{48} \\ &\leq \frac{\varepsilon^2}{16} \end{aligned} \quad (70)$$

since $N = \lceil \frac{48 \times 16n^3}{\alpha \varepsilon^4} \rceil = \lceil \frac{48 \times 16 \times 72n^5}{\varepsilon^6} \rceil$, $\alpha = \frac{\varepsilon^2}{72(K+1)^2 n^2}$ and Lemmas 15 and 16.

Service Rate Dominates Arrival Rate. Next, we wish to establish that the average of empirical service rate dominates the average arrival rate over time interval of length N . That is, for all i

$$\frac{1}{N} \left(\sum_{j=1}^N (\mathbb{E}[\hat{s}_i(j)]) \right) \geq \lambda_i + \varepsilon/2. \quad (71)$$

To this end, first note that (using Cauchy–Schwartz inequality)

$$\begin{aligned} &\frac{1}{N} \sum_{j=1}^N \mathbb{E} \left[\left\| \frac{\partial F_\varepsilon}{\partial r_i}(\mathbf{r}(j)) \right\| \right] \\ &\leq \sqrt{\frac{1}{N} \sum_{j=1}^N \mathbb{E} \left[\left\| \frac{\partial F_\varepsilon}{\partial r_i}(\mathbf{r}(j)) \right\|^2 \right]} \end{aligned}$$

$$\begin{aligned} &\leq \sqrt{\frac{1}{N} \sum_{j=1}^N \mathbb{E} \left[\left\| \frac{\partial F_\varepsilon}{\partial r_i}(\mathbf{r}(j)) \right\|^2 \right]} \\ &\leq \frac{\varepsilon}{4} \end{aligned} \quad (72)$$

where the last inequality is from (70). Therefore

$$\begin{aligned} &\frac{1}{N} \left(\sum_{j=1}^N (\mathbb{E}[\hat{s}_i(j)] - \lambda_i) \right) \\ &= \frac{1}{N} \left(\sum_{j=1}^N ((\mathbb{E}[s_i(\mathbf{r}(j))] - \lambda_i) + (\mathbb{E}[\hat{s}_i(j)] - \mathbb{E}[s_i(\mathbf{r}(j))])) \right) \\ &\geq \frac{1}{N} \left(\sum_{j=1}^N ((\mathbb{E}[s_i(\mathbf{r}(j))] - \lambda_i) - |\mathbb{E}[\hat{s}_i(j)] - \mathbb{E}[s_i(\mathbf{r}(j))]|) \right) \\ &\stackrel{(a)}{\geq} \frac{1}{N} \left(\sum_{j=1}^N \left(\varepsilon - \mathbb{E} \left[\left\| \frac{\partial F_\varepsilon}{\partial r_i}(\mathbf{r}(j)) \right\| \right] - \frac{\varepsilon^3}{240n} \right) \right) \\ &\geq \frac{3}{4}\varepsilon - \frac{1}{N} \sum_{j=1}^N \mathbb{E} \left[\left\| \frac{\partial F_\varepsilon}{\partial r_i}(\mathbf{r}(j)) \right\| \right] \\ &\stackrel{(b)}{\geq} \frac{\varepsilon}{2}. \end{aligned} \quad (73)$$

In the above, (a) follows from Lemma 8(2), i.e.,

$$s_i(\mathbf{r}(j)) = \lambda_i + \varepsilon - \frac{\partial F_\varepsilon}{\partial r_i}(\mathbf{r}(j))$$

and from Lemma 16; and (b) follows from (72).

Wrapping Up: Negative Drift. Now, consider $Q_i(N)$. For this, suppose $Q_i(0) > TN$. Then, $Q_i(\cdot)$ is strictly positive over interval $[0, TN]$ as service rate is at most 1. Therefore, in that case, the queue $Q_i(\cdot)$ is *fully served* in time $[0, TN]$. Hence, using (73), we conclude that

$$\begin{aligned} \mathbb{E}[Q_i(TN)] &= Q_i(0) + T \left(\sum_{j=1}^N \mathbb{E}[\lambda_i - \hat{s}_i(j)] \right) \\ &\leq Q_i(0) - \frac{\varepsilon}{2} T \cdot N \end{aligned} \quad (74)$$

if $Q_i(0) > TN$. In the above, as usual, we have assumed that the expectation is conditional with respect to $X(0)$. In what follows, we will use this conditioning explicitly. Given (74), we have

$$\begin{aligned} &\mathbb{E}[Q_i^2(TN) - Q_i^2(0) | X(0)] \\ &= \mathbb{E}[(Q_i(TN) - Q_i(0))(Q_i(TN) + Q_i(0)) | X(0)] \\ &= \mathbb{E}[(Q_i(TN) - Q_i(0))^2 + 2Q_i(0)(Q_i(TN) - Q_i(0)) | X(0)] \\ &\stackrel{(a)}{\leq} (KTN)^2 + 2Q_i(0)\mathbb{E}[Q_i(TN) - Q_i(0) | X(0)] \\ &\stackrel{(b)}{\leq} \begin{cases} (KTN)^2 - \frac{\varepsilon}{2} TN \times 2Q_i(0), & \text{if } Q_i(0) > TN \\ (K^2 + 2K)(TN)^2, & \text{if } Q_i(0) \leq TN \end{cases} \quad (75) \\ &\leq -\varepsilon TN Q_i(0) + \varepsilon(TN)^2 + (K^2 + 2K)(TN)^2 \end{aligned} \quad (76)$$

for all $\mathbf{Q}(0)$. In the above, (a) is from boundedness of arrival process and (b) is from (74). Hence

$$\begin{aligned} & \mathbb{E}[L(X(N)) - L(X(0)) | X(0)] \\ &= \mathbb{E} \left[\sum_{i=1}^n Q_i^2(TN) - \sum_{i=1}^n Q_i^2(0) \middle| X(0) \right] \\ &\leq -\varepsilon TN \left(\sum_{i=1}^n Q_i(0) \right) + n(\varepsilon + K^2 + 2K)T^2N^2. \end{aligned}$$

This completes the proof of Lemma 13.

C. Proof of Lemma 14

We wish to establish that set $B_\kappa = \{\mathbf{x} : L(\mathbf{x}) \leq \kappa\}$ is a closed petite set. By definition, it is closed. To establish that it is a petite set, we need to find a nontrivial measure μ on (X, \mathcal{B}_X) and a sampling distribution a on \mathbb{N} so that for any $\mathbf{x} \in B_\kappa$,

$$K_a(\mathbf{x}, \cdot) \geq \mu(\cdot).$$

To construct such a measure μ , we will use Proposition 17.

Proposition 17: Let the network Markov chain $X(\cdot)$ start with state $\mathbf{x} \in B_\kappa$ at time 0, $X(0) = \mathbf{x}$. Then, there exists $T_\kappa \geq 1$ and $\gamma_\kappa > 0$ such that

$$\sum_{\tau=1}^{T_\kappa} \Pr_{\mathbf{x}}(X(\tau) = \mathbf{y}) \geq \gamma_\kappa, \quad \forall \mathbf{x} \in B_\kappa.$$

Here $\mathbf{y} = (\mathbf{0}, \lfloor \frac{n}{\varepsilon} \rfloor, \mathbf{0}) \in X$ denote the state where all components of \mathbf{Q} and σ (i.e., the schedule is the empty independent set) and $r_i = \frac{n}{\varepsilon}$ for all $i \in V$.

Proof: Consider any $\mathbf{x} \in B_\kappa$. By definition, total amount of work in each queue is no more than $\sqrt{\kappa} + 1$. Consider some large enough (soon to be determined) T_κ . By the property of the assumed arrival process, there is a positive probability $\theta_\kappa^0 > 0$ of no arrivals happening to the system in time T_κ . Assuming no arrivals happen, we will show that in large enough time t_κ^1 , with probability $\theta_\kappa^1 > 0$, each queue receives at least $\sqrt{\kappa} + 1$ amount of service; and after that in additional time t^2 with positive probability $\theta^2 > 0$, the empty set schedule is reached. Now, after the empty set schedule is reached, in additional time t^3 with positive probability $\theta^3 > 0$, the empty set schedule remains; i.e., the scheduling does not change in this time. Since the empty set schedule remains and no new data arrives, r_i is increasing by ε from (6) and finally reaches $\frac{n}{\varepsilon}$ for a large enough t^3 which depends on n . This will imply that by defining $T_\kappa \triangleq t_\kappa^1 + t^2 + t^3$ the state $\mathbf{y} \in X$ is reached with probability at least

$$\gamma_\kappa \triangleq \theta_\kappa^0 \theta_\kappa^1 \theta^2 \theta^3 > 0.$$

This will immediately imply the desired result of Proposition 17. To this end, we need to show existence of t_κ^1, t^2, t^3 and $\theta_\kappa^1, \theta^2, \theta^3$ with properties stated above to complete the proof of Proposition 17.

First, show the existence of $t_\kappa^1, \theta_\kappa^1$. For this, note that the Markov chain corresponding to the scheduling algorithm has always bounded transition probabilities (since \mathbf{r} is bounded in terms of n) and is irreducible over the space of all independent

sets $\mathcal{I}(G)$. Therefore, it follows that starting from any initial scheduling configuration, there exists finite time \hat{t} such that a schedule is reached so that any given queue i is scheduled for at least unit amount of time with probability at least $\hat{\theta} > 0$. Here, both $\hat{t}, \hat{\theta}$ depend on only n (and ε), not κ . Therefore, it follows that in time $t_\kappa^1 \triangleq (\sqrt{\kappa} + 1)n\hat{t}$ all queues become empty with probability at least $\theta_\kappa^1 \triangleq (\hat{\theta})^{n(\sqrt{\kappa}+1)}$. Next, in the existence of t^2, θ^2 also follows from the bounded property of our Markov chain. Finally, for t^3, θ^3 , consider the interpretation of the Markov chain as in Section IV-B using the clock ticks. Note that no clock ticks in time t^3 with probability $\theta^3 > 0$ since its rate is bounded in terms of n . Hence, the empty set schedule remains in time t^3 with probability $\theta^3 > 0$, where t^3 and θ^3 depend only on n . This completes the proof of Proposition 17. \square

In what follows, Proposition 17 will be used to complete the proof of Lemma 14. To this end, consider Geometric(1/2) as the sampling distribution a , i.e.,

$$a(\ell) = 2^{-\ell}, \quad \ell \geq 1.$$

Let $\delta_{\mathbf{y}}$ be the delta distribution on element $\mathbf{y} \in X$. Then, define μ as

$$\mu = 2^{-T_\kappa} \gamma_\kappa \delta_{\mathbf{y}}, \quad \text{that is } \mu(\cdot) = 2^{-T_\kappa} \gamma_\kappa \delta_{\mathbf{y}}(\cdot).$$

Clearly, μ is nontrivial measure on (X, \mathcal{B}_X) . With these definitions of a and μ , Proposition 17 immediately implies that for any $\mathbf{x} \in B_\kappa$

$$K_a(\mathbf{x}, \cdot) \geq \mu(\cdot).$$

This establishes that set B_κ is a closed petite set and this completes the proof of Lemma 14.

VI. THROUGHPUT AND FAIRNESS OF CONGESTION CONTROL ALGORITHMS

A. Proof of Theorem 3: Rate Stable Congestion Control

The proof of Theorem 3 is similar to that of Theorem 1. In a nutshell, the basic idea is to show that the update (8) solves a relevant optimization problem through a subgradient algorithm. That is, $\lambda(j), \mathbf{r}(j)$ converge to the solution of the appropriate optimization problem with probability 1. The property of the optimization problem will imply the goodness of utility of the convergent arrival rates. Using this convergence property, it will in turn imply rate stability of queue size.

A Relevant Optimization Problem and Its Properties. Let \mathcal{M} be space of all probability distributions on $\mathcal{I}(G)$. Given a distribution $\mu \in \mathcal{M}$, recall that $H_{\text{ER}}(\mu)$ denotes its entropy

$$H_{\text{ER}}(\mu) = - \sum_{\sigma \in \mathcal{I}(G)} \mu_\sigma \log \mu_\sigma.$$

Consider the following optimization problem:

$$\begin{aligned} & \text{maximize} && H_{\text{ER}}(\mu) + \beta \left(\sum_i U_i(\lambda_i) \right) \\ & \text{over} && \mu \in \mathcal{M}, \lambda \in [0, 1]^n \\ & \text{subject to} && \mathbb{E}_\mu[\sigma_i] \geq \lambda_i, \quad \text{for all } i. \end{aligned} \quad (77)$$

Associate a dual variable $r_i \geq 0$ to constraint $\mathbb{E}_{\boldsymbol{\mu}}[\sigma_i] \geq \lambda_i$. Here the use of r_i for dual variable is an intentional abuse of notation and the reason behind this will soon become clear to the reader. Given this, the result Lagrangian is given by

$$\begin{aligned} \mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\lambda}; \mathbf{r}) &= \text{H}_{\text{ER}}(\boldsymbol{\mu}) + \beta \left(\sum_i U_i(\lambda_i) \right) + \sum_i r_i (\mathbb{E}_{\boldsymbol{\mu}}[\sigma_i] - \lambda_i) \\ &= \text{H}_{\text{ER}}(\boldsymbol{\mu}) + \sum_i r_i \mathbb{E}_{\boldsymbol{\mu}}[\sigma_i] + \sum_i [\beta U_i(\lambda_i) - r_i \lambda_i]. \end{aligned} \quad (78)$$

Therefore, the dual function is given by

$$\begin{aligned} \mathcal{D}(\mathbf{r}) = \sup_{\boldsymbol{\mu} \in \mathcal{M}, \boldsymbol{\lambda} \in [0, 1]^n} \mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\lambda}; \mathbf{r}) \end{aligned} \quad (79)$$

Finally, the dual optimization of (77) is given by

$$\text{minimize } \mathcal{D}(\mathbf{r}) \text{ over } \mathbf{r} \in \mathbb{R}_+^n. \quad (80)$$

Now we are ready to state useful properties of the optimization problems (77) and (80). These properties were presented in earlier work [29].

Lemma 18: The optimization problem (77) is concave maximization while the optimization problem (80) is convex minimization. There is no duality gap and hence both have the same optimal value. They satisfy the following properties.

- 1) Given dual feasible $\mathbf{r} \in \mathbb{R}_+^n$, the associate primal feasible assignment $\boldsymbol{\mu}(\mathbf{r}), \boldsymbol{\lambda}(\mathbf{r})$ is given as follows:

$$\mu_{\boldsymbol{\sigma}} \propto \exp(\boldsymbol{\sigma} \cdot \mathbf{r}), \quad \text{for all } \boldsymbol{\sigma} \in \mathcal{I}(G). \quad (81)$$

That is, $\boldsymbol{\mu}(\mathbf{r}) = \boldsymbol{\pi}^{\mathbf{r}}$ and

$$\lambda_i(\mathbf{r}) = \arg \max_{y \in [0, 1]} (\beta U_i(y) - r_i y), \quad \forall i. \quad (82)$$

- 2) The subgradient for $\mathcal{D}(\mathbf{r})$, represented as $g(\mathbf{r}) = [g_i(\mathbf{r})]$, is given by

$$g_i(\mathbf{r}) = \mathbb{E}_{\boldsymbol{\mu}(\mathbf{r})}[\sigma_i] - \lambda_i(\mathbf{r}).$$

- 3) Both problems have unique optimal solutions.

Proof: To begin with, observe that the objective of (77) is strictly concave as entropy is a strictly concave function over \mathcal{M} and so are U_i for all i under our setup. Therefore, given the constraints of (77), the unique optimal exists and is achieved. To observe the lack of duality gap, note that there exists a $\boldsymbol{\mu} \in \mathcal{M}$ and a $\boldsymbol{\lambda} \in [0, 1]^n$ that is strictly feasible. Therefore, Slater's condition will imply lack of duality gap. We defer the proof of uniqueness of the dual optimal solution until a little later.

Proof of 1). Given the dual feasible $\mathbf{r} \in \mathbb{R}_+^n$, let $\boldsymbol{\mu}(\mathbf{r}), \boldsymbol{\lambda}(\mathbf{r})$ be the corresponding primal feasible solutions that maximize the Lagrangian \mathcal{L} . Given structure of \mathcal{L} as in (78), it follows that $\boldsymbol{\lambda}(\mathbf{r})$ must be such that

$$\lambda_i(\mathbf{r}) = \arg \max_{y \in [0, 1]} (\beta U_i(y) - r_i y), \quad \text{for all } i.$$

For $\boldsymbol{\mu}(\mathbf{r})$, observe that

$$\frac{\partial \mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\lambda}; \mathbf{r})}{\partial \mu_{\boldsymbol{\sigma}}} = -\log \mu_{\boldsymbol{\sigma}} - 1 + \boldsymbol{\sigma} \cdot \mathbf{r}.$$

Since $\boldsymbol{\mu}(\mathbf{r})$ is maximizing \mathcal{L} , from the above, it follows that $\mu_{\boldsymbol{\sigma}}(\mathbf{r}) \in (0, 1)$ for all $\boldsymbol{\sigma} \in \mathcal{I}(G)$. Therefore, for any $\boldsymbol{\sigma}, \boldsymbol{\rho} \in \mathcal{I}(G)$ and $\boldsymbol{\sigma} \neq \boldsymbol{\rho}$, it must be that

$$\frac{\partial \mathcal{L}(\boldsymbol{\mu}(\mathbf{r}), \boldsymbol{\lambda}(\mathbf{r}); \mathbf{r})}{\partial \mu_{\boldsymbol{\sigma}}} = \frac{\partial \mathcal{L}(\boldsymbol{\mu}(\mathbf{r}), \boldsymbol{\lambda}(\mathbf{r}); \mathbf{r})}{\partial \mu_{\boldsymbol{\rho}}}.$$

That is

$$\mu_{\boldsymbol{\sigma}}(\mathbf{r}) \propto \exp(\boldsymbol{\sigma} \cdot \mathbf{r}), \quad \text{for all } \boldsymbol{\sigma} \in \mathcal{I}(G).$$

Thus, $\boldsymbol{\mu}(\mathbf{r}) = \boldsymbol{\pi}^{\mathbf{r}}$.

Proof of 2). Given 1), it follows that

$$\mathcal{D}(\mathbf{r}) = \mathcal{L}(\boldsymbol{\mu}(\mathbf{r}), \boldsymbol{\lambda}(\mathbf{r}); \mathbf{r}).$$

Now the dual variables \mathbf{r} capture ‘‘slack’’ in the corresponding constraints of (77). Specifically, for a given \mathbf{r} , if the corresponding primal solutions are $\boldsymbol{\mu}(\mathbf{r}), \boldsymbol{\lambda}(\mathbf{r})$, then the slack in the i th constraint is $s_i(\mathbf{r}) - \lambda_i(\mathbf{r})$: if it is positive, r_i should be decreased and if it is negative, r_i should be increased. This intuition is formalized in the optimization theory (e.g., see [7]) by establishing that a subgradient of the dual optimization at \mathbf{r} is given by vector $g(\mathbf{r}) \in \mathbb{R}^n$ with

$$g_i(\mathbf{r}) = s_i(\mathbf{r}) - \lambda_i(\mathbf{r}).$$

Proof of 3). The uniqueness of solution of (77) was already explained. To understand uniqueness of \mathbf{r}^* , consider independent set e_i , which has only node i in it; and the null set $\mathbf{0}$. Then, since $\boldsymbol{\mu}(\mathbf{r}^*) = \boldsymbol{\pi}^{\mathbf{r}^*}$, it follows that

$$\mu_{e_i}(\mathbf{r}^*) = \mu_{e_0}(\mathbf{r}^*) \exp(r_i^*).$$

Now suppose to contrary that there is another optimal solution of (80), $\hat{\mathbf{r}} \neq \mathbf{r}^*$. Then, it will immediately contradict the above as $\boldsymbol{\mu}^*$ is unique as discussed above. This completes the proof of 3). \square

Convergence of $\mathbf{r}(j), \boldsymbol{\lambda}(j)$. In light of Lemma 18(2), it follows that the algorithm (8) is motivated by the standard projected dual subgradient algorithm. The algorithm uses estimated $\hat{\mathbf{s}}(\mathbf{r}(j))$ in place of $\mathbf{s}(\mathbf{r}(j))$; but exact update for $\boldsymbol{\lambda}(\mathbf{r}(j))$. That is, for all i

$$r_i(j+1) = [r_i(j) + \alpha(j)(\lambda_i(j) - \hat{s}_i(j))]_{+}.$$

To this end, define ‘‘error’’ vector

$$e(j) \triangleq -\hat{\mathbf{s}}(j) + \mathbf{s}(\mathbf{r}(j))$$

and let

$$d(j) \triangleq \|\mathbf{r}(j) - \mathbf{r}^*\|_2^2.$$

Now consider the relation between $d(j+1)$ and $d(j)$. Since the projection $[\cdot]_+$ is nonexpansive

$$\begin{aligned} d(j+1) &\leq \left\| \mathbf{r}(j) - \mathbf{r}^* + \frac{1}{j} [\boldsymbol{\lambda}(j) - \mathbf{s}(\mathbf{r}(j)) + e(j)] \right\|_2^2 \\ &\leq d(j) + 2[\mathbf{r}(j) - \mathbf{r}^*]^T \cdot \frac{1}{j} [\boldsymbol{\lambda}(j) - \mathbf{s}(\mathbf{r}(j)) + e(j)] + O\left(\frac{n}{j^2}\right) \end{aligned}$$

where we have used the fact that each component of $\boldsymbol{\lambda}(j) - \mathbf{s}(\mathbf{r}(j)) + e(j)$ is $O(1)$. Define, the error in optimal cost at the j th step as

$$\Delta(j) = \mathcal{D}(\mathbf{r}(j)) - \mathcal{D}(\mathbf{r}^*).$$

By definition, $\Delta(j) \geq 0$. Since the dual objective \mathcal{D} is convex, and $\mathbf{s}(\mathbf{r}(j)) - \boldsymbol{\lambda}(j)$ is its subgradient at $\mathbf{r}(j)$, we have

$$[\mathbf{r}(j) - \mathbf{r}^*]^T \cdot [\boldsymbol{\lambda}(j) - \mathbf{s}(\mathbf{r}(j))] \leq -\Delta(j). \quad (83)$$

Also, as used earlier, $r_i(j) = O(\log j)$ for all i . Therefore, from the above, we obtain that

$$\begin{aligned} d(j+1) &\leq d(j) - \frac{2\Delta(j)}{j} \\ &\quad + O\left(\frac{\log(j) + \|\mathbf{r}^*\|_\infty \|e(j)\|_1}{j}\right) + O\left(\frac{n}{j^2}\right). \end{aligned} \quad (84)$$

Note that the analysis of Lemma 10 applies to bound $\|e(j)\|_1$ as is. That is

$$\mathbb{E}[\|e(j)\|_1] = O\left(\frac{n}{j^2}\right). \quad (85)$$

Using this and taking expectation on both sides of inequality (84), we obtain

$$\begin{aligned} \mathbb{E}[d(j+1)] &\leq \mathbb{E}[d(j)] - \frac{2}{j} \mathbb{E}[\Delta(j)] \\ &\quad + O\left(\frac{n \cdot (\log(j) + \|\mathbf{r}^*\|_\infty)}{j^3}\right) + O\left(\frac{n}{j^2}\right). \end{aligned}$$

Summing the above inequality from 1 to ∞ , it follows that

$$\begin{aligned} 0 &\leq \mathbb{E}[d(\infty)] \\ &\leq \mathbb{E}[d(1)] - \left(\sum_{j=1}^{\infty} \frac{2}{j} \mathbb{E}[\Delta(j)] \right) + O(n). \end{aligned}$$

By rearranging the terms and using $\mathbb{E}[d(1)] < \infty$, it follows that $\sum_{j=1}^{\infty} \frac{1}{j} \mathbb{E}[\Delta(j)] < \infty$. Since $\sum_{j=1}^{\infty} \frac{1}{j} = \infty$, we can conclude that with probability 1

$$\begin{aligned} \liminf_j \mathbb{E}[\Delta(j)] = 0 &\Rightarrow \liminf_j \Delta(j) = 0, \\ &\Rightarrow \liminf_j \|\mathbf{r}(j) - \mathbf{r}^*\| = 0 \end{aligned} \quad (86)$$

where we have used the fact that dual optimization (80) has a unique solution and it is convex minimization problem. Now, the rest of the proof of $\mathbf{r}(j) \rightarrow \mathbf{r}^*(j)$ with probability 1 follows

exactly the same set of arguments as those used in the proof of Theorem 1. The convergence of $\boldsymbol{\lambda}(j) \rightarrow \boldsymbol{\lambda}(\mathbf{r}^*) = \bar{\boldsymbol{\lambda}}$ follows due to continuity of solution of concave maximization (82) with respect to \mathbf{r} .

Utility of $\bar{\boldsymbol{\lambda}}$, Rate Stability. To begin with, we observe that convergence $\mathbf{r}(j) \rightarrow \mathbf{r}^*(j)$ and $\boldsymbol{\lambda}(j) \rightarrow \boldsymbol{\lambda}(\mathbf{r}^*) = \bar{\boldsymbol{\lambda}}$ with probability 1 implies the rate stability using exactly the same arguments as those used in Lemma 12.

To establish goodness of the $\bar{\boldsymbol{\lambda}}$, note that along with $\boldsymbol{\mu}^*$ it optimizes (77). Now $\boldsymbol{\lambda}^*$, the optimal allocation [as per (3)] along with appropriate distribution, say $\boldsymbol{\nu}^*$ on $\mathcal{I}(G)$ is a feasible solution. Therefore, it follows that

$$\begin{aligned} \beta \sum_i U_i(\lambda_i^*) &\leq H_{\text{ER}}(\boldsymbol{\nu}^*) + \beta \sum_i U_i(\lambda_i^*) \\ &\leq H_{\text{ER}}(\boldsymbol{\mu}^*) + \beta \sum_i U_i(\bar{\lambda}_i) \\ &\leq \log |\mathcal{I}(G)| + \beta \sum_i U_i(\bar{\lambda}_i). \end{aligned} \quad (87)$$

In the above, we have used the fact that the entropy is nonnegative and the maximum value of a discrete valued random variable's entropy is at most the logarithm of the cardinality of the support set. Equation (87) immediately implies the desired result. This completes the proof of Theorem 3.

B. Proof of Theorem 4

The proof of Theorem 4 in a nutshell requires us to establish that the average rate allocation $\bar{\boldsymbol{\lambda}}$ has near optimal total utility. This follows using similar arguments that we used in proving Theorem 3. That is, establish that the $\bar{\boldsymbol{\lambda}}$ ends up approximately solving optimization problem (77). This property follows primarily because the Congestion Control Algorithm 2 with update (11) is primarily designed as a constant step size dual ‘‘subgradient’’ algorithm. We will formalize this in the rest of this section. We begin with a useful property that establishes uniform bound on components of $\mathbf{r}(\cdot)$ and subsequently implies uniform bound on the components of the queue-size vector $\mathbf{Q}(\cdot)$ for all time duration. This will be followed by proof of the goodness of average rate $\bar{\boldsymbol{\lambda}}$ to conclude the proof of Theorem 4.

Uniform Bound on $\|\mathbf{r}(j)\|_\infty$. We state and prove the following bound on $\|\mathbf{r}(j)\|_\infty$ starting with $\mathbf{r}(0) = \mathbf{0}$.

Lemma 19: Under the update rule (11), for all $1 \leq i \leq n$

$$r_i(j) \in [0, \beta V + \alpha], \quad \text{for all } j$$

where recall that V is defined in (10) and α is the constant step size used in the update (11).

Proof: To prove this Lemma, consider any $i, 1 \leq i \leq n$. Now for any $j, r_i(j) \geq 0$ by the definition [cf., (11)]. To prove $r_i(j) \leq \beta V + \alpha$, we will use the principle of mathematical induction. To this end, for the base case, $j = 0$ and $r_i(0) = 0$ by definition. Suppose, as the inductive hypothesis, that the property $r_i(j) \leq \beta V + \alpha$ is true for all $j \leq J$. Now we wish to establish this property for $j = J + 1$. To this end, we consider

two cases: a) $r_i(J) \leq \beta V$, or b) $r_i(J) \in (\beta V, \beta V + \alpha]$. First, consider case a). By (11), it follows that

$$\begin{aligned} r_i(J+1) &= [r_i(J) - \alpha \hat{s}_i(J)]_+ + \alpha \lambda_i(J) \\ &\leq r_i(J) + \alpha \lambda_i(J) \\ &\leq r_i(J) + \alpha \\ &\leq \beta V + \alpha. \end{aligned}$$

In the above, we have used the fact that $\lambda_i(J) \in [0, 1]$ by definition. Now consider case b). For this note that if $r_i(J) \in [\beta V, \beta V + \alpha]$, then $\lambda_i(J) = 0$. This is because by (11), $\lambda_i(J)$ solves

$$\lambda_i(J) \in \arg \max_{y \in [0,1]} \{\beta U'_i(y) - r_i(J)y\} \quad (88)$$

and for any $y \in [0, 1]$

$$\begin{aligned} \frac{d}{dy}(\beta U'_i(y) - r_i(J)y) &= \beta U''_i(y) - r_i(J) \\ &\leq \beta V - r_i(J) \\ &< 0. \end{aligned} \quad (89)$$

That is, the optimal solution of (88) is 0. This completes the proof of Lemma 19. \square

Uniform Bound on $\|\mathbf{Q}(j)\|_\infty$. We state and prove the following bound on $\|\mathbf{Q}(j)\|_\infty$ starting with $\mathbf{Q}(0) = \mathbf{0}$.

Lemma 20: Under the Congestion Control Algorithm 2, starting with an empty queue, i.e., $\mathbf{Q}(0) = \mathbf{0}$, the following holds for all $t \geq 0$:

$$Q_i(t) \leq \frac{T}{\alpha}(\beta V + 2\alpha).$$

Proof: In what follows, we will show that for time instances $t = jT$, for $j \geq 0$, the queue size is bounded as

$$Q_i(jT) \leq \frac{T}{\alpha} r_i(j), \quad \text{for all } i. \quad (90)$$

Equation (90) along with the bound on $r_i(\cdot)$ implied by Lemma 19 will imply

$$Q_i(jT) \leq \frac{T}{\alpha}(\beta V + \alpha), \quad \text{for all } i. \quad (91)$$

Finally, by noticing that $\lambda_i(j) \in [0, 1]$ for all i, j , it follows that for any $t \in [jT, (j+1)T)$, $Q_i(t) \leq Q_i(jT) + T$. Therefore, we will obtain the desired result of Lemma 20.

Now we prove the remaining bounded as stated in (90). To this end, note that

$$Q_i((j+1)T) \leq [Q_i(jT) - \hat{s}_i(j)T]_+ + \lambda_i(j)T. \quad (92)$$

This follows by imagining that all the arrival traffic in $[jT, (j+1)T)$, $\lambda_i(j)T$ amount of data, is added to the queue at the end of the interval; service $\hat{s}_i(j)T$ is used only to serve data that were present at time jT .

Based on (92), we will establish (90) by means of the principle of mathematical induction. For the based case of $j = 0$,

we have $Q_i(0) = 0$ and $r_i(0) = 0$. For induction hypothesis, assume it to hold true for all $j \leq J$. For $j = J+1$, we wish to establish that the relation holds. To this end, using (92), it follows that

$$\begin{aligned} Q_i((J+1)T) &\leq [Q_i(JT) - \hat{s}_i(J)T]_+ + \lambda_i(J)T \\ &\leq \left[\frac{T}{\alpha} r_i(J) - \hat{s}_i(J)T \right]_+ + \lambda_i(J)T \\ &= \frac{T}{\alpha} ([r_i(J) - \alpha \hat{s}_i(J)]_+ + \alpha \lambda_i(J)) \\ &= \frac{T}{\alpha} r_i(J+1). \end{aligned} \quad (93)$$

Here the last equality follows by definition (11). This completes the proof of (90) and Lemma 20. \square

A Useful Variational Characterization. We state the Gibbs variational characterization (e.g., see [17]) of the distribution $\pi^{\mathbf{r}}$ that will be useful later in the proof.

Lemma 21: Given $\mathbf{r} \in \mathbb{R}^n$, $\pi^{\mathbf{r}}$ is the unique solution of

$$\begin{aligned} &\text{maximize} \quad \mathbb{E}_\mu[\boldsymbol{\sigma} \cdot \mathbf{r}] + H_{\text{ER}}(\boldsymbol{\mu}) \\ &\text{over} \quad \boldsymbol{\mu} \in \mathcal{M} \end{aligned} \quad (94)$$

where recall that \mathcal{M} is the space of probability distributions over $\mathcal{I}(G)$. Further

$$\mathbb{E}_{\pi^{\mathbf{r}}}[\boldsymbol{\sigma} \cdot \mathbf{r}] \geq \max_{\boldsymbol{\lambda} \in \Lambda} \boldsymbol{\lambda} \cdot \mathbf{r} - \log |\mathcal{I}(G)|. \quad (95)$$

Proof: Equation (94) was established implicitly in Lemma 18. To see an explicit proof, consider the following. For any $\boldsymbol{\mu} \in \mathcal{M}$

$$\begin{aligned} &\mathbb{E}_\mu[\boldsymbol{\sigma} \cdot \mathbf{r}] + H_{\text{ER}}(\boldsymbol{\mu}) \\ &= \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} (\boldsymbol{\sigma} \cdot \mathbf{r}) \mu_\sigma - \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \mu_\sigma \log \mu_\sigma \\ &\stackrel{(a)}{=} \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} (\log \pi_\sigma^{\mathbf{r}} + \log Z(\mathbf{r})) \mu_\sigma - \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \mu_\sigma \log \mu_\sigma \\ &= \log Z(\mathbf{r}) + \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \mu_\sigma \log \frac{\pi_\sigma^{\mathbf{r}}}{\mu_\sigma} \\ &\stackrel{(b)}{\leq} \log Z(\mathbf{r}). \end{aligned} \quad (96)$$

In the above, (a) follows from the fact that

$$\pi_\sigma^{\mathbf{r}} = \frac{1}{Z(\mathbf{r})} \exp(\boldsymbol{\sigma} \cdot \mathbf{r})$$

and (b) follows from an application of Jensen's inequality. The above suggests that the optimal cost of (94) is $\log Z(\mathbf{r})$ and is achieved iff the $\boldsymbol{\mu} = \pi^{\mathbf{r}}$. This establishes the first claim of Lemma 21.

To see (95), define $\boldsymbol{\mu}^*$ as

$$\mu_\sigma^* = \begin{cases} 1, & \text{if } \boldsymbol{\sigma} = \boldsymbol{\sigma}^* \\ 0, & \text{o.w.} \end{cases}$$

Here $\boldsymbol{\sigma}^* = \arg \max_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \boldsymbol{\sigma} \cdot \mathbf{r}$. Then, using the above, it follows that

$$\begin{aligned} \mathbf{s}(\mathbf{r}) \cdot \mathbf{r} &= \mathbb{E}_{\boldsymbol{\pi}^{\mathbf{r}}}[\boldsymbol{\sigma} \cdot \mathbf{r}] \\ &\geq \mathbb{E}_{\boldsymbol{\mu}^*}[\boldsymbol{\sigma} \cdot \mathbf{r}] + H_{\text{ER}}(\boldsymbol{\mu}^*) - H_{\text{ER}}(\boldsymbol{\pi}^{\mathbf{r}}) \\ &\stackrel{(a)}{\geq} \boldsymbol{\sigma}^* \cdot \mathbf{r} + 0 - \log |\mathcal{I}(G)| \\ &\stackrel{(b)}{=} \max_{\boldsymbol{\lambda} \in \boldsymbol{\Lambda}} \boldsymbol{\lambda} \cdot \mathbf{r} - \log |\mathcal{I}(G)|. \end{aligned} \quad (97)$$

In the above, (a) follows from the definition of $\boldsymbol{\mu}^*$ and the fact that for any distribution on \mathcal{M} , the entropy is at the most $\log |\mathcal{I}(G)|$; and (b) follows because any $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}$ is a convex combination of elements in $\mathcal{I}(G)$. \square

Some Properties. Here we state some useful properties that will be useful in completing the proof of Theorem 4. To begin with, let $\boldsymbol{\lambda}^*$ be the optimal solution to congestion control problem (3). At any stage j , $\boldsymbol{\lambda}(j)$ is obtained as

$$\lambda_i(j) \in \arg \max_{y \in [0,1]} \{\beta U_i(y) - r_i(j)y\}, \quad \text{for all } i.$$

Therefore, it follows that for any j

$$\beta U_i(\lambda_i(j)) - r_i(j)\lambda_i(j) \geq \beta U_i(\lambda_i^*) - r_i(j)\lambda_i^*. \quad (98)$$

Since $\boldsymbol{\lambda}^* \in \boldsymbol{\Lambda}$, we have

$$\sum_i \lambda_i^* r_i(j) \leq \max_{\boldsymbol{\lambda} \in \boldsymbol{\Lambda}} \boldsymbol{\lambda} \cdot \mathbf{r}(j). \quad (99)$$

Define notation $m^*(\mathbf{r}) = \max_{\boldsymbol{\lambda} \in \boldsymbol{\Lambda}} \boldsymbol{\lambda} \cdot \mathbf{r}$. From (98) and (99), we have

$$\mathbf{r}(j) \cdot \boldsymbol{\lambda}(j) \leq \beta \left(\sum_i U_i(\lambda_i(j)) \right) - \beta \left(\sum_i U_i(\lambda_i^*) \right) + m^*(\mathbf{r}(j)). \quad (100)$$

We will observe another useful property. By Lemma 19, we have $\|\mathbf{r}(j)\|_\infty$ bounded by $\beta V + \alpha$. Therefore, using the mixing time bounds and arguments utilized in Lemma 10, we obtain that by the choice of appropriately large T as

$$T = \exp(\Theta(\beta n V)) \Theta \left(\frac{(\beta V + \alpha)n^2}{\beta \varepsilon} \right) \quad (101)$$

we have that for all j

$$|\mathbb{E}[\hat{s}_i(j)|\mathcal{F}_j] - s_i(\mathbf{r}(j))| \leq \frac{\beta \varepsilon}{2(\beta V + \alpha)n}, \quad \forall i. \quad (102)$$

In the above, the conditioning \mathcal{F}_j represents the filtration (or information) until time $l(j)$, while recall that the random variable $\hat{s}_i(j)$ is the empirical service rate in $[l(j), l(j+1))$.

Wrapping Up: Completing the Proof of Theorem 4. A key element of the following proof is a drift analysis of $\|\mathbf{r}(j)\|^2$ similar to that in [49].⁹ Now, let us start with the algorithm's

⁹Two important differences with the proof of [49] are worth noting. 1) As mentioned earlier, the random access CSMA algorithm takes time to approach the stationary distribution (which approximates the maximal-weight schedule). This time contributes to the queue lengths and needs to be quantified. 2) Theorem 4 establishes pathwise performance guarantees instead of in expectation, even though the state space of $\mathbf{r}(j)$ is uncountable.

update rule (11). Specifically, for a given i , squaring both sides of (11) for $r_i(\cdot)$ gives us

$$\begin{aligned} r_i^2(j+1) &= ([r_i(j) - \alpha \hat{s}_i(j)]_+ + \alpha \lambda_i(j))^2 \\ &= [r_i(j) - \alpha \hat{s}_i(j)]_+^2 + 2\alpha \lambda_i(j)[r_i(j) - \alpha \hat{s}_i(j)]_+ + \alpha^2 \lambda_i^2(j) \\ &\stackrel{(a)}{\leq} [r_i(j) - \alpha \hat{s}_i(j)]^2 + 2\alpha \lambda_i(j)r_i(j) + \alpha^2 \\ &\stackrel{(b)}{\leq} r_i(j)^2 + 2\alpha r_i(j)[\lambda_i(j) - \hat{s}_i(j)] + 2\alpha^2. \end{aligned} \quad (103)$$

In the above, (a) follows from the fact that $[x]_+^2 \leq x^2$ and $\lambda_i(j) \in [0, 1]$ for all i, j ; and (b) follows from the fact that $\hat{s}_i(j) \in [0, 1]$ for all i, j . From (103), we have that

$$\begin{aligned} &\left(\sum_i r_i^2(j+1) - r_i^2(j) \right) \\ &\leq 2\alpha \left(\sum_i r_i(j)(\lambda_i(j) - \hat{s}_i(j)) \right) + 2n\alpha^2 \\ &= 2\alpha \left[\sum_i r_i(j)(\lambda_i(j) - s_i(\mathbf{r}(j))) \right. \\ &\quad \left. + \sum_i r_i(j)(s_i(\mathbf{r}(j)) - \hat{s}_i(j)) + n\alpha \right] \\ &= 2\alpha[\mathbf{r}(j) \cdot \boldsymbol{\lambda}(j) - \mathbf{r}(j) \cdot \mathbf{s}(\mathbf{r}(j)) \\ &\quad + \mathbf{r}(j) \cdot (\mathbf{s}(\mathbf{r}(j)) - \hat{\mathbf{s}}(j)) + n\alpha]. \end{aligned} \quad (104)$$

By (97) and since $|\mathcal{I}(G)| \leq 2^n$, we have

$$-2\alpha \mathbf{r}(j) \cdot \mathbf{s}(\mathbf{r}(j)) \leq -2\alpha m^*(\mathbf{r}(j)) + 2\alpha n. \quad (105)$$

Therefore, using (100), we have

$$\begin{aligned} &2\alpha \mathbf{r}(j) \cdot \boldsymbol{\lambda}(j) - 2\alpha \mathbf{r}(j) \cdot \mathbf{s}(\mathbf{r}(j)) \\ &\leq 2\alpha \beta \left(\sum_i U_i(\lambda_i(j)) - U_i(\lambda_i^*) \right) + 2\alpha n. \end{aligned} \quad (106)$$

Now using (106) in (104) and the fact that $\alpha^2 \leq \alpha$ because $\alpha \in (0, 1)$, we have

$$\begin{aligned} &\left(\sum_i r_i^2(j+1) - r_i^2(j) \right) \\ &\leq 2\alpha \beta (U(\boldsymbol{\lambda}(j)) - U(\boldsymbol{\lambda}^*)) + 2\alpha \mathbf{r}(j) \\ &\quad \cdot (\mathbf{s}(\mathbf{r}(j)) - \hat{\mathbf{s}}(j)) + 4\alpha n \end{aligned} \quad (107)$$

where we have used notation $U(\boldsymbol{\lambda}) = \sum_i U_i(\lambda_i)$. Now taking its summation from $j = 0$ until $J - 1$ on both sides of (107), the fact that $\mathbf{r}(0) = \mathbf{0}$ and dividing both side by J , we have

$$\begin{aligned} &\frac{1}{J} \sum_i r_i^2(J) \\ &\leq 2\alpha \beta \left(\sum_{j=0}^{J-1} \frac{U(\boldsymbol{\lambda}(j))}{J} - U(\boldsymbol{\lambda}^*) \right) \\ &\quad + \frac{1}{J} \left(\sum_{j=0}^{J-1} 2\alpha \mathbf{r}(j) \cdot (\mathbf{s}(\mathbf{r}(j)) - \hat{\mathbf{s}}(j)) \right) + 4\alpha n. \end{aligned} \quad (108)$$

Now, define $\Delta(j) = 2\alpha\mathbf{r}(j) \cdot (\mathbf{s}(\mathbf{r}(j)) - \hat{\mathbf{s}}(j))$ and $X(j) = \Delta(j) - \mathbb{E}[\Delta(j)|\mathcal{F}_j]$. By definition, $S(j) = \sum_{k=0}^{j-1} X(k)$ is a martingale with respect to filtration $\{\mathcal{F}_j\}_{j \geq 1}$. With this notation, we have that for any J

$$\begin{aligned} & \frac{1}{J} \left(\sum_{j=0}^{J-1} 2\alpha\mathbf{r}(j) \cdot (\mathbf{s}(\mathbf{r}(j)) - \hat{\mathbf{s}}(j)) \right) \\ &= \frac{1}{J} \left(\sum_{j=0}^{J-1} X(j) + \mathbb{E}[\Delta(j)|\mathcal{F}_j] \right) \\ &\stackrel{(a)}{\leq} \frac{1}{J} S(J) + \alpha\beta\varepsilon. \end{aligned} \quad (109)$$

In the above, (a) follows from (102) and bound on $\mathbf{r}(\cdot)$ using Lemma 19. Finally, note that $S(\cdot)$ is a martingale with bounded increment due to uniform bound on $\mathbf{r}(\cdot)$, the fact that $\mathbf{s}(\cdot), \hat{\mathbf{s}}(\cdot)$ are vectors in $[0, 1]^n$ and $\alpha \in (0, 1)$. Therefore, by strong law of large numbers for martingales with bounded increments, it follows that

$$\lim_{J \rightarrow \infty} \frac{1}{J} S(J) = 0, \quad \text{with probability 1.} \quad (110)$$

That is, with probability 1

$$\limsup_{J \rightarrow \infty} \frac{1}{J} \left(\sum_{j=0}^{J-1} 2\alpha\mathbf{r}(j) \cdot (\mathbf{s}(\mathbf{r}(j)) - \hat{\mathbf{s}}(j)) \right) \leq \alpha\beta\varepsilon. \quad (111)$$

Using (111) in (108) along with Lemma 19, and then taking $J \rightarrow \infty$, we have that with probability 1

$$\liminf_{J \rightarrow \infty} \sum_{j=0}^{J-1} \frac{U(\lambda(j))}{J} \geq U(\lambda^*) - \frac{\varepsilon}{2} - \frac{2n}{\beta}. \quad (112)$$

Finally, observe that by concavity of function $\sum_i U_i(\cdot)$ along with Jensen's inequality, we have that for $\tilde{\lambda}_i(J) = (\sum_{j=0}^{J-1} \lambda_i(j))/J$

$$U(\tilde{\lambda}(J)) \geq \sum_{j=0}^{J-1} \frac{U(\lambda(j))}{J}.$$

Therefore, the following desired conclusion of Theorem 4 follows from (112) along with choice of $\beta = 4n/\varepsilon$: with probability 1

$$\liminf_{J \rightarrow \infty} U(\tilde{\lambda}(J)) \geq U(\lambda^*) - \varepsilon. \quad (113)$$

VII. CONCLUSION

In this paper, we have presented a simple, distributed randomized algorithm for scheduling and congestion control in a network. Our algorithm is essentially a random access protocol with time-varying access probabilities. Our algorithm for scheduling, in the presence of exogenous arrivals, achieves throughput optimality while our algorithm for scheduling with congestion controlled arrivals achieves near-optimal resource allocation when nodes have concave utilities. We believe that

the algorithmic method presented in this paper should be of general interest.

REFERENCES

- [1] N. Abramson, "The aloha system," in *Computer-Communication Networks*, N. Abramson and F. Kuo, Eds. Englewood Cliffs, NJ: Prentice-Hall, 1973.
- [2] D. J. Aldous, "Ultimate instability of exponential back-off protocol for acknowledgement-based transmission control of random access communication channels," *IEEE Trans. Inf. Theory*, vol. IT-33, no. 2, pp. 219–223, Mar. 1987.
- [3] T. Anderson, S. Owicki, J. Saxe, and C. Thacker, "High-speed switch scheduling for local-area networks," *ACM Trans. Comput. Syst.*, vol. 11, no. 4, pp. 319–352, 1993.
- [4] C. Bordenave, D. McDonald, and A. Proutiere, "Performance of random medium access—An asymptotic approach," in *Proc. ACM Sigmetrics Int. Conf. Meas. Model. Comput. Sci.*, 2008, pp. 1–12.
- [5] V. Borkar, "Stochastic approximation with controlled Markov noise," *Syst. Control Lett.*, vol. 55, pp. 139–145, 2006.
- [6] V. Borkar, *Stochastic Approximation, a Dynamical System View Point*. Cambridge, U.K.: Cambridge Univ. Press, 2008.
- [7] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [8] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," *Proc. IEEE*, vol. 95, no. 1, pp. 255–312, Jan. 2007.
- [9] J. Dai and B. Prabhakar, "The throughput of data switches with and without speedup," in *Proc. 19th Annu. Joint Conf. IEEE Comput. Commun. Soc.*, 2000, vol. 2, pp. 556–564.
- [10] J. G. Dai, "On positive Harris recurrence of multiclass queueing networks," *Ann. Appl. Probab.*, vol. 5, pp. 49–77, 1995.
- [11] J. G. Dai, "Stability of fluid and stochastic processing networks," Maphysto Lecture Notes, Jan. 1999 [Online]. Available: www.maphysto.dk, ISSN: 1398-5957.
- [12] P. Diaconis and D. Stroock, "Geometric bounds for eigenvalues of Markov chains," *Ann. Appl. Probab.*, vol. 1, no. 1, pp. 36–61, 1991.
- [13] A. Dimakis and J. Walrand, "Sufficient conditions for stability of longest-queue-first scheduling: Second-order properties using fluid limits," *Adv. Appl. Probab.*, vol. 38, no. 2, pp. 505–521, 2006.
- [14] A. Eryilmaz and R. Srikant, "Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control," in *Proc. 24th Annu. Joint Conf. IEEE Comput. Commun. Soc.*, Mar. 2005, vol. 3, pp. 1794–1803.
- [15] A. Eryilmaz and R. Srikant, "Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control," *IEEE/ACM Trans. Netw.*, vol. 15, no. 6, pp. 1333–1344, Dec. 2007.
- [16] S. Foss and T. Konstantopoulos, "An overview of some stability methods," *J. Oper. Res. Soc. Jpn.*, vol. 47, no. 4, pp. 275–303, 2004.
- [17] H. O. Georgii, *Gibbs Measures and Phase Transitions*. Berlin, Germany: de Gruyter, 1988.
- [18] R. K. Gettoor, J. Azma and M. Yor, Eds., "Transience and recurrence of Markov processes," *Séminaire de Probabilités XIV*, pp. 397–409, 1979.
- [19] P. Giaccone, B. Prabhakar, and D. Shah, "Randomized scheduling algorithms for high-aggregate bandwidth switches," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 4, pp. 546–559, May 2003.
- [20] L. Goldberg, M. Jerrum, S. Kannan, and M. Paterson, "A bound on the capacity of backoff and acknowledgement-based protocols," Dept. Comput. Sci., Univ. Warwick, Coventry, U.K., Res. Rep. 365, CV4 7AL, Jan. 2000.
- [21] L. A. Goldberg, "Design and analysis of contention-resolution protocols," Epsrc Res. Grant gr/I60982, Oct. 2002 [Online]. Available: <http://www.csc.liv.ac.uk/leslie/contention.html>
- [22] P. Gupta and A. L. Stolyar, "Optimal throughput allocation in general random-access networks," in *Proc. IEEE 40th Annu. Conf. Inf. Sci. Syst.*, Princeton, NJ, 2006, pp. 1254–1259.
- [23] B. Hajek and T. van Loon, "Decentralized dynamic control of a multi-access broadcast channel," *IEEE Trans. Autom. Control*, vol. 27, no. 3, pp. 559–569, Jun. 1982.
- [24] J. M. Harrison, *Brownian Motion and Stochastic Flow Systems*. New York: Wiley, 1985.
- [25] J. M. Harrison, "Brownian models of open processing networks: Canonical representation of workload," *Ann. Appl. Probab.*, vol. 10, pp. 75–103, 2000.
- [26] J. Hastad, T. Leighton, and B. Rogoff, "Analysis of backoff protocols for multiple access channels," *SIAM J. Comput.*, vol. 25, no. 4, pp. 740–774, 1996.

- [27] M. Jerrum and A. Sinclair, "Approximating the permanent," *SIAM J. Comput.*, vol. 18, pp. 1149–1178, 1989.
- [28] L. Jiang, J. Ni, R. Srikant, and J. Walrand, "Performance bounds of distributed csma scheduling," in *Proc. Inf. Theory Appl. Workshop*, San Diego, CA, 2010, DOI: 10.1109/ITA.2010.5454123.
- [29] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," in *Proc. 46th Annu. Allerton Conf. Commun. Control Comput.*, 2008, pp. 1511–1519.
- [30] L. Jiang and J. Walrand, "Convergence and stability of a distributed CSMA algorithm for maximal network throughput," *Electr. Eng. Comput. Sci. Dept., Univ. California, Berkeley, CA, Tech. Rep. UCB/EECS-2009-43*, Mar. 2009.
- [31] C. Joo, X. Lin, and N. Shroff, "Understanding the capacity region of the greedy maximal scheduling algorithm in multi-hop wireless networks," in *Proc. 27th Conf. Comput. Commun.*, 2008, pp. 1103–1111.
- [32] F. Kelly and I. MacPhee, "The number of packets transmitted by collision detect random access schemes," *Ann. Probab.*, vol. 15, no. 4, pp. 1557–1568, 1987.
- [33] F. P. Kelly, "Stochastic models of computer communication systems," *J. R. Stat. Soc. B*, vol. 47, no. 3, pp. 379–395, 1985.
- [34] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, no. 3, pp. 237–252, Mar. 1998.
- [35] M. Leconte, J. Ni, and R. Srikant, "Improved bounds on the throughput efficiency of greedy maximal scheduling in wireless networks," in *Proc. 10th ACM Int. Symp. Mobile Ad Hoc Netw. Comput.*, 2009, pp. 165–174.
- [36] X. Lin and N. Shroff, "Joint rate control and scheduling in multihop wireless networks," in *Proc. IEEE Conf. Decision Control*, 2004, vol. 2, pp. 1484–1489.
- [37] J. Liu, Y. Yi, A. Proutiere, M. Chiang, and V. Poor, "Convergence and tradeoff of utility-optimal CSMA," *IEEE Commun. Lett.*, Feb. 2009, submitted for publication.
- [38] L. Lovasz and P. Winkler, "Mixing times," in *Microsurveys in Discrete Probability*, ser. Discrete Mathematics and Theoretical Computer Science, D. Aldous and J. Propp, Eds. New Brunswick, NJ: DIMACS, 1998, vol. 41, pp. 85–133.
- [39] S. H. Low and D. E. Lapsley, "Optimization flow control, I: Basic algorithm and convergence," *IEEE/ACM Trans. Netw.*, vol. 7, no. 6, pp. 861–874, Dec. 1999.
- [40] I. MacPhee, "On optimal strategies in stochastic decision processes," D.Phil. dissertation, Dept. Math., Univ. Cambridge, Cambridge, U.K., 1989.
- [41] P. Marbach, A. Eryilmaz, and A. Ozdaglar, "Achievable rate region of CSMA schedulers in wireless networks with primary interference constraints," in *Proc. IEEE Conf. Decision Control*, 2007, pp. 1156–1161.
- [42] N. McKeown, "The iSLIP: A scheduling algorithm for input-queued switches," *IEEE Trans. Netw.*, vol. 7, no. 2, pp. 188–201, Apr. 1999.
- [43] R. Metcalfe and D. Boggs, "Distributed packet switching for local computer networks," *Commun. ACM*, vol. 19, pp. 395–404, 1976.
- [44] S. P. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability*. London, U.K.: Springer-Verlag, 1993.
- [45] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 556–567, Oct. 2000.
- [46] E. Modiano, D. Shah, and G. Zussman, "Maximizing throughput in wireless network via gossiping," *ACM SIGMETRICS/Performance Eval. Rev.*, vol. 34, no. 1, pp. 27–38, Jun. 2006.
- [47] R. Montenegro and P. Tetali, "Mathematical aspects of mixing times in Markov chains," *Found. Trends Theor. Comput. Sci.*, vol. 1, no. 3, pp. 237–354, 2006.
- [48] J. Mosely and P. Humblet, "A class of efficient contention resolution algorithms for multiple access channels," *IEEE Trans. Commun.*, vol. 33, no. 2, pp. 145–151, 1985.
- [49] M. Neely, E. Modiano, and C.-P. Li, "Fairness and optimal stochastic control for heterogeneous networks," in *Proc. 24th Conf. IEEE Comput. Commun. Soc.*, 2005, vol. 3, pp. 1723–1734.
- [50] M. J. Neely, E. Modiano, and C. P. Li, "Fairness and optimal stochastic control for heterogeneous networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 2, pp. 396–409, Apr. 2008.
- [51] S. Rajagopalan and D. Shah, "Distributed algorithm and reversible network," in *Proc. 42nd Annu. Conf. Inf. Sci. Syst.*, 2008, pp. 498–502.
- [52] S. Rajagopalan, D. Shah, and J. Shin, "A network adiabatic theorem: An efficient randomized protocol for contention resolution," in *Proc. 11th Int. Joint Conf. Meas. Model. Comput. Sci.*, 2009, pp. 133–144.
- [53] J. Roberts and L. Massoulié, "Bandwidth sharing and admission control for elastic traffic," *Telecommun. Syst.*, vol. 15, pp. 185–201, 2000.
- [54] D. Shah, "Network scheduling and message passing," in *Performance Modeling and Engineering*, Z. Liu and C. Xia, Eds. New York: Springer-Verlag, 2008.
- [55] D. Shah, "Gossip algorithms," in *Foundations and Trends in Networking*. Boston, MA: Now Publishers, Jun. 2009 [Online]. Available: <http://web.mit.edu/devavrat/www/Gossipbook.pdf>
- [56] D. Shah and J. Shin, "Randomized scheduling algorithm for queueing networks," 2009 [Online]. Available: <http://arxiv.org/abs/0908.3670>
- [57] D. Shah, D. Tse, and J. N. Tsitsiklis, "Hardness of low delay network scheduling," *IEEE Trans. Inf. Theory*, Aug. 2009, submitted for publication.
- [58] D. Shah and D. J. Wischik, "Optimal scheduling algorithm for input queued switch," in *Proc. 25th IEEE Int. Conf. Comput. Commun.*, 2006, DOI: 10.1109/INFOCOM.2006.238.
- [59] D. Shah and D. J. Wischik, "Switched networks with maximum weight policies: fluid approximation and multiplicative state space collapse," 2007–2009 [Online]. Available: <http://arxiv.org/abs/1004.1995>
- [60] S. Shakkottai and R. Srikant, "Network optimization and control," in *Foundations and Trends in Networking*. Boston, MA: Now Publishers, 2007.
- [61] A. Stolyar, "Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm," *Queueing Syst.*, vol. 50, no. 4, pp. 401–457, 2005.
- [62] A. L. Stolyar, "Dynamic distributed scheduling in random access networks," *J. Appl. Probab.*, vol. 45, no. 2, pp. 297–313, 2008.
- [63] L. Tassiulas, "Linear complexity algorithms for maximum throughput in radio networks and input queued switches," in *Proc. IEEE 17th Annu. Joint Conf. IEEE Comput. Commun. Soc.*, 1998, vol. 2, pp. 533–539.
- [64] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.
- [65] B. Tsybakov and N. B. Likhanov, "Upper bound on the capacity of a random multiple-access system," *Problemy Peredachi Informatsii*, vol. 23, no. 3, pp. 64–78, 1987.
- [66] X. Wang and K. Kar, "Throughput modelling and fairness issues in CSMA/CA based ad-hoc networks," in *Proc. 24th Annu. Joint Conf. IEEE Comput. Commun. Soc.*, 2005, vol. 1, pp. 23–34.

Libin Jiang (S'07) received the B.Eng. degree in electronic engineering and information science from the University of Science and Technology of China, Hefei, China, in 2003, the M.Phil. degree in information engineering from the Chinese University of Hong Kong, Shatin, Hong Kong, in 2005, and the Ph.D. degree in electrical engineering and computer sciences from the University of California at Berkeley, Berkeley, in 2009.

His research interests include wireless networks, communications, and game theory.

Dr. Jiang received the David Sakrison Memorial Prize for outstanding doctoral research, and the best presentation award at the 2009 ACM Mobihoc S3 Workshop.

Devavrat Shah (M'05) received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology (IIT), Bombay, India, in 1999, with the honor of the President of India Gold Medal, and the Ph.D. degree from the Computer Science Department, Stanford University, Stanford, CA, in 2004.

He was a Postdoctoral Researcher at the Statistics Department, Stanford University, in 2004–2005. Currently, he is a Jamieson Career Development Associate Professor with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (MIT), Cambridge. He is a member of the Laboratory of Information and Decision Systems (LIDS) and affiliated with the Operations Research Center (ORC). He has authored monograph *Gossip Algorithms* (Boston, MA: Now Publishers, 2008). His research focus is on theory of large complex networks which includes network algorithms, stochastic networks, network information theory, and large scale statistical inference.

Dr. Shah was coawarded the best paper awards at the 2004 Conference of the IEEE Communications Society and the 2006 ACM SIGMETRICS/Performance, and best student paper awards at the 2008 Neural Information Processing Systems and the 2009 ACM SIGMETRICS/Performance. He received the 2005 George B. Dantzig best dissertation award from the INFORMS. He received the first 2008 ACM SIGMETRICS Rising Star Award for his work on network scheduling algorithms.

Jinwoo Shin received the B.Sc. degree in mathematics and computer science (double majors) from Seoul National University, Seoul, Korea in 2001, and the Ph.D. degree in mathematics from the Massachusetts Institute of Technology (MIT), Cambridge, in 2010.

His primary research interest is in understanding large complex networks including communication, societal, and statistical networks.

Dr. Shin won a silver medal at the International Mathematical Olympiad in 1995. He was a corecipient of the 2009 ACM SIGMETRICS/Performance Kenneth C. Sevcik outstanding student paper award. He received the 2010 George Sprowl's award for outstanding doctoral thesis in computer science at MIT.

Jean Walrand (S'71–M 80–SM'90–F'93) received the Ph.D. degree in electrical engineering and computer science from the University of California at Berkeley, Berkeley, in 1979.

He has been a Professor at the University of California at Berkeley since 1982. He is the author of *An Introduction to Queueing Networks* (Englewood Cliffs, NJ: Prentice-Hall, 1988) and *Communication Networks: A First Course* (New York: McGraw-Hill, 1998, 2nd ed.) and coauthor of *High Performance Communication Networks* (San Mateo, CA: Morgan Kaufman, 2000, 2nd ed.) and *Communication Networks: A Concise Introduction* (San Rafael, CA: Morgan-Claypool, 2010).

Prof. Walrand is a Fellow of the Belgian American Education Foundation and a recipient of the Lanchester Prize and the Stephen O. Rice Prize.