

Scheduling for Small Delay in Multi-rate Multi-channel Wireless Networks

Shrees Shankar Bodas
Department of EECS

Massachusetts Institute of Technology
bodas@mit.edu

Sanjay Shakkottai
Department of ECE

The University of Texas at Austin
shakkott@mail.utexas.edu

Lei Ying
Department of ECE

Iowa State University
leiying@iastate.edu

R. Srikant

Department of ECE

Univ. of Illinois at Urbana-Champaign
rsrikant@illinois.edu

Abstract—This paper considers the problem of designing scheduling algorithms for multi-channel (e.g., OFDM-based) wireless downlink systems. We show that the Server-Side Greedy (SSG) rule introduced in earlier papers for ON-OFF channels performs well even for more general channel models. The key contribution in this paper is the development of new mathematical techniques for analyzing Markov chains that arise when studying general channel models. These techniques include a way of calculating the distribution of the maximum of a multi-dimensional Markov chain (note that the maximum does not have the Markov property on its own), and also a Markov chain stochastic dominance result using coupling arguments.

Index Terms—Scheduling algorithms, large deviations, small buffer, Markov chain stochastic dominance

I. INTRODUCTION

Scheduling for OFDM (Orthogonal Frequency Division Multiplexing) wireless systems (e.g., WiMax [6] and LTE [1]) is an active area of research in both academia and industry. These systems use an OFDM-based wireless downlink, where the bandwidth available at the base-station is partitioned into hundreds or thousands of orthogonal frequency bands. In every timeslot, a given frequency band can be allocated to one and only one user, but a given user can be served by multiple frequency bands simultaneously, and the allocation can change over time, depending upon the channel quality and the queue backlogs, among other parameters.

The challenge is to design scheduling algorithms for allocating the resources (frequency bands) to the users based on the wireless channel quality and traffic requirements, and with performance guarantees for all users. We want the scheduling rule to be throughput-optimal, and also result in a small per-user delay. Delay is a particularly important performance metric for real-time traffic such as voice or video, and is closely related to maintaining small queue-lengths at the base-station where the incoming packets to a given user are temporarily stored.

Thus, our main objective is to investigate the small-queue characteristics of the OFDM-based wireless downlink system, under the assumption that it has a large number of users and a proportionally large bandwidth. The well-known MaxWeight-type algorithms [12] stabilize the system under a very general class of arrival and channel processes if there is any other scheduling algorithm than can do so. However, we showed in [3] that the MaxWeight algorithm results in a very poor

delay performance for the system under consideration. We then proposed an algorithm called SSG (Server-Side Greedy) that, in addition to being throughput-optimal, results in a very good per-user delay performance.

The proofs of the good small-queue performance of the SSG algorithm crucially depended upon a sample-path dominance property of the algorithm: if there are two (multi-queue) queuing systems \underline{Q} and \underline{R} with queues $\{Q_i\}_{i=1}^n$ and $\{R_i\}_{i=1}^n$, with sample-path coupled arrivals and channels, and $Q_i(t-1) \leq R_i(t-1)$ for all i and for some t , then $Q_i(t) \leq R_i(t)$ for all i . That is, if in a given timeslot the queue-length vector of a queuing system dominates that of the other queuing system element-by-element, and if both the systems use the SSG scheduling rule, then the queue-length dominance continues to hold for all the future timeslots. This sample-path property fails to hold for the case when the channel service rates are more general than 0 or 1 packets per timeslot, rendering the earlier analysis techniques useless for the analysis of the more general systems. As a result, it was unclear if the SSG algorithm provided good delay performance when more realistic channel models are considered. In this paper, we present a new framework for analyzing the rate-function performance of the SSG-like algorithms that do not necessarily have the sample-path dominance property, and consequently establish that the SSG algorithm has good large-deviations performance properties even for the general channel models.

A. Related Work

This paper continues the line of work initiated in [2], [3]. The main difference between this paper and our prior work is that we now consider channels with multiple states, instead of simple ON-OFF models. Related work on optimal scheduling algorithms for multi-channel networks can be found in [7], but the analysis there is only for the case of two users. Apart from these references, to best of our knowledge, the vast prior literature on the performance of channel state-aware scheduling in wireless networks seems to only address heavy traffic behavior or large-buffer asymptotics [14], [11], [13], [10], [9], [8].

B. Main Contributions

The good delay (small-queue) performance results for the proposed algorithms in [2] and [3] were derived for a system where each frequency band could serve only 0 or 1 packets per timeslot. These results crucially depended upon the sample-path dominance property of the proposed algorithms as mentioned in Section I. Even for a system where the frequency bands can (each) serve 0 or 2 packets per timeslot, the earlier analysis techniques fail and it is not clear if there exists any (let alone “simple”) scheduling rule that guarantees the sample-path dominance property. In this paper, we develop new analysis techniques that make possible the analysis of systems and scheduling rules that do not necessarily have the sample-path dominance property, but satisfy much milder conditions. We believe that these Markov chain-based techniques are not only specific to the problem and are of independent interest.

II. SYSTEM MODEL

We study a queuing system with multiple queues and servers, and operating in discrete-time, as shown in Figure 1. Each server S_k models a frequency band or a collection of OFDM subcarriers. We build on the model, performance metrics and notation described in [3], and this section is provided here for completeness. Table I summarizes




Fig. 1. System Model

the notation used throughout this paper. If no confusion

Q_i	=	The entity, queue number i
S_i	=	The entity, server number i
\mathcal{Q}	=	$\{Q_1, Q_2, \dots, Q_n\}$
\mathcal{S}	=	$\{S_1, S_2, \dots, S_n\}$
$A_i(t)$	=	The number of packet arrivals to Q_i at the beginning of timeslot t
$X_{i,j}(t)$	=	The number of packets in Q_i that can potentially be served by S_j , in timeslot t
$Q_i(t)$	=	The length of Q_i at the end of timeslot t
$Q_i^{(k)}(t)$	=	The length of Q_i after $k \geq 1$ rounds of service in timeslot t
$Q_i^{(0)}(t)$	=	$Q_i(t-1) + A_i(t)$, i.e. the length of Q_i after immediately after arrivals, in timeslot t
a^+	=	$\max(a, 0)$
\mathbb{R}	=	The set of real numbers
\mathbb{R}_+	=	The set of nonnegative real numbers
\mathbb{Z}_+	=	The set of nonnegative integers
w.p.	=	with probability
$H(x y)$	=	$x \log \frac{x}{y} + (1-x) \log \frac{1-x}{1-y}$
$M_1(\Sigma)$	=	The probability simplex in \mathbb{R}^k for appropriate k

TABLE I
NOTATION

is possible, we denote $X_{i,j}(t)$ by $X_{ij}(t)$. We assume that $A_i(t), X_{ij}(t), Q_i(t), Q_i^{(k)}(t)$ take values in the set of non-negative integers. For this system, we make the following assumptions on the arrival and channel processes.

Assumption 1 (Multi-level Channels and Arrivals).

The number of packet arrivals to queue Q_i in timeslot t is the random variable $A_i(t)$, where $A_i(t) = r$ with probability p_r for $0 \leq r \leq M$ for some integer $M \geq 1$.

In timeslot t , the server S_j can potentially serve $X_{ij}(t)$ packets from Q_i , where $X_{ij}(t)$ are modeled as random variables with $X_{ij}(t) = \ell$ with probability q_ℓ for $0 \leq \ell \leq K$ for some integer K . We assume that $q_\ell \in (0, 1)$ for all i , and $\sum_i q_i = 1$.

We assume that all the random variables $A_i(t)$ and $X_{jk}(s)$ are mutually independent for all possible values of the involved parameters, that $p_i > 0$ for all $i \in \{0, 1, \dots, M\}$, $\sum_{m=0}^M p_m = 1$, and $\sum_{m=1}^M m p_m < K$. \diamond

For notational convenience, we define $q := q_k$. Our objective is to design a rule for allocating the servers to the queues, based on the current and past arrival and channel process realizations, the past allocation decisions, and any amount of external randomness (if needed). In every timeslot t , this scheduling rule defines the random variables $Y_{i,j}(t)$ where

$$Y_{i,j}(t) = \begin{cases} 1 & \text{if } S_j \text{ is allocated to serve } Q_i \text{ in timeslot } t, \\ 0 & \text{otherwise.} \end{cases}$$

If no confusion is possible, we denote $Y_{i,j}(t)$ by $Y_{ij}(t)$. In real systems, a given OFDM subcarrier can be allocated to serve only one user in a given timeslot. We model this case by imposing the following condition on the server allocations: $\sum_{i=1}^n Y_{ij}(t) \leq 1$ for all t and all $j \in \{1, 2, \dots, n\}$. The individual queues in the system evolve according to the following equation:

$$Q_i(t) = \left(Q_i(t-1) + A_i(t) - \sum_{j=1}^n X_{ij}(t) Y_{ij}(t) \right)^+.$$

As is clear from the above equation, the queues can store any number of packets in the buffers, and the system does not drop any packets.

We want to design a scheduling rule that results in small per-user queues at the base-station, in a probabilistic sense. Mathematically, we want a scheduling rule that results in the maximum possible value of the function

$$I(b) := \liminf_{n \rightarrow \infty} \frac{-1}{n} \log \mathbb{P} \left(\max_{1 \leq i \leq n} Q_i(0) > b \right).$$

Here $b \geq 0$ is a given integer, and can be roughly interpreted as the available per-user buffer-size. The probability measure $\mathbb{P}(\cdot)$ is the stationary measure of the queue-length process (i.e., the queue length at time ‘0’ when the system started at time ‘ $-\infty$ ’). The function $I(\cdot)$ is called the rate function [5], and is a useful surrogate for the actual probability of the “error” event we are interested in (in this case, one of the queues exceeding the given limit of b). Maximizing the rate-function is a surrogate for minimizing the probability of this “error” event, and is mathematically more tractable.

In principle, the scheduling algorithm can depend upon the given buffer-limit, b . We are however interested in an

algorithm that does not explicitly use the value of b , and still results in a good small-queues performance. In Section V, we analyze such a scheduling rule called the Server-Side Greedy rule (SSG, that we introduced in [3]), show that it performs well (results in a positive value of the rate-function and therefore, small per-user queues at the base-station) for the multi-user multi-channel system under consideration.

Remark 1. *The main objective of this paper is to develop new techniques for analyzing the small-queue characteristics of different algorithms where the system has channels that can potentially serve multiple packets per timeslot. This regime is interesting because the earlier sample-path-based techniques are no longer useful. We therefore do not consider the throughput-optimality-related issues for the proposed algorithm, as they are addressed in [3].*

III. PRELIMINARIES

In this section, we present certain basic results regarding the stability of the system under Assumption 1, and also algorithm-independent upper bounds on the rate-function under this assumption.

Lemma 1. *Under Assumption 1, if $\sum_{m=1}^M mp_m > K$, then the system is unstable under any scheduling algorithm. If $\sum_{m=1}^M mp_m < K$, then there exists a constant $n_0 = n_0(p, M, K, q)$ such that for all $n \geq n_0$, the system is stable under some algorithm.*

Proof: This proof is similar to that of Theorem 2 in [2] and has been omitted due to lack of space. Please see [15] for a detailed proof. ■

Theorem 1. *Under Assumption 1, for any scheduling rule,*

$$\mathbb{P} \left(\max_{1 \leq i \leq n} Q_i(0) > b \right) \geq p_M^{\lfloor \frac{b}{M} \rfloor + 1} (q_0)^{n(\lfloor \frac{b}{M} \rfloor + 1)}.$$

Consequently, under any scheduling rule,

$$\limsup_{n \rightarrow \infty} \frac{-1}{n} \log \mathbb{P} \left(\max_{1 \leq i \leq n} Q_i(0) > b \right) \leq \left(\left\lfloor \frac{b}{M} \right\rfloor + 1 \right) \log \frac{1}{q_0}.$$

Proof: Please see Appendix A. ■

Lemma 2. *Under Assumption 1, if $\sum_{i=1}^M p_i \left[\frac{i}{K} \right] > 1$, then under any scheduling algorithm,*

$$\liminf_{n \rightarrow \infty} \frac{-1}{n} \log \mathbb{P} \left(\max_{1 \leq i \leq n} Q_i(0) > 0 \right) = 0.$$

Proof: Please see Appendix B. ■

Remark 2. *It is possible that under Assumption 1, for some constant b large enough, and under some algorithm, we have*

$$\liminf_{n \rightarrow \infty} \frac{-1}{n} \log \mathbb{P} \left(\max_{1 \leq i \leq n} Q_i(0) > b \right) > 0.$$

However, since we are primarily concerned with the small-buffer overflow event, for the purpose of this paper, we are not interested in the case where the system parameters are such

that the rate-function is zero at a finite value of b (even at $b = 0$). We therefore assume throughout the rest of this paper that $\sum_{m=1}^M p_m \left[\frac{m}{K} \right] < 1$ whenever we mention Assumption 1.

IV. STOCHASTIC DOMINANCE OF MARKOV CHAINS

The next three technical theorems are instrumental in characterizing the rate-function performance of the SSG algorithm and are interesting in their own right.

Theorem 2. *Consider a discrete-time, multi-dimensional Markov chain $\mathbf{X}(t) = [X_1(t), X_2(t), \dots, X_n(t)]$. Let $\mathbf{X}(t)$ take values on the countable state-space \mathbb{Z}_+^n , and let the corresponding state-transition probabilities be*

$$p(\mathbf{x}, \mathbf{y}) := \mathbb{P}(\mathbf{X}(t+1) = \mathbf{y} \mid \mathbf{X}(t) = \mathbf{x}), \quad \mathbf{x}, \mathbf{y} \in \mathbb{Z}_+^n.$$

Suppose that the Markov chain $\mathbf{X}(t)$ has a stationary distribution $\sigma(\mathbf{x})$. Further, let

$$\mathcal{W}_k := \left\{ [x_1, x_2, \dots, x_n] : \max_{1 \leq i \leq n} x_i = k \right\}.$$

Then a stationary distribution of $X^(t) := \max_{1 \leq i \leq n} X_i(t)$ is given by a stationary distribution of the (one-dimensional) Markov chain $Y(t)$, taking values in \mathbb{Z}_+ , whose transition probabilities are given by*

$$\mathbb{P}(Y(t+1) = y \mid Y(t) = x) = \mathbb{P}(\mathbf{X}(t+1) \in \mathcal{W}_y \mid \mathbf{X}(t) \in \mathcal{W}_x),$$

where the conditional probability term on the RHS is computed under the stationary distribution $\sigma(\cdot)$.

Proof: Please see Appendix C. ■

The above theorem gives us a way to calculate the stationary distribution of the maximum value of a multi-dimensional Markov chain (which on its own does not have the Markov property) by thinking of it as a one-dimensional Markov chain. Note that we do not need the stationary distribution to be unique for either $\mathbf{X}(t)$ or $Y(t)$, and also that the result holds (with trivial modifications in the notation) for the Markov chains on \mathbb{Z}^n . Note also that if the Markov chain $\mathbf{X}(t)$ is irreducible and has a stationary distribution, then the stationary distribution is unique (Theorem 3.1, Chapter 3 in [4]). Before stating the next result, we remind the reader of the stochastic dominance of random variables: we say that a random variable Y stochastically dominates a random variable X , and write $X \leq_{st} Y$, if $\mathbb{P}(X > z) \leq \mathbb{P}(Y > z)$ for all $z \in \mathfrak{R}$.

Theorem 3. *Consider two discrete-time Markov chains $Y(t)$ and $W(t)$ evolving on the same state-space \mathbb{Z}_+ . Let the random variable $Z_t(i)$ denote the increment in $Y(t)$ when $Y(t) = i$. (Note that $Z_t(i)$ can be negative.) Similarly, let $\tilde{Z}_t(j)$ denote the increment in $W(t)$ when $W(t) = j$. Let $Z_t(i) \leq_{st} \tilde{Z}_t(j)$ for all $i, j, t \in \mathbb{Z}_+$, and $Y(0) \leq_{st} W(0)$. Then $Y(t) \leq_{st} W(t)$ for all $t \geq 0$. In particular, if $Y(t)$ and $W(t)$ are ergodic (aperiodic, irreducible, positive recurrent) and have stationary distributions $\mu_Y(\cdot)$ and $\mu_W(\cdot)$ respectively, and the random variables Y, W are distributed according to $Y \sim \mu_Y(\cdot)$ and $W \sim \mu_W(\cdot)$, then $Y \leq_{st} W$.*

Proof: Please see Appendix D. ■

We now analyze the steady-state distribution of a special class of one-dimensional Markov chains from a rate-function point of view. The Markov chains in this class are similar to birth-death Markov chains, except that there can be multiple (but a finite, bounded number of) “births” in a given timeslot, or at most one “death.” The probability of birth(s) is “small,” and the probability of death is at least a constant. Hence, it is reasonable to expect that the stationary distribution is strongly concentrated around 0. We quantify this intuition in a large-deviations sense in the following theorem.

Theorem 4. *Consider a family of Markov chains $W^{(n)}(t)$ on the set \mathbb{Z}_+ , having the following transition probability structure: there exists an integer n_0 such that for all $n \geq n_0$, for all $x \in \mathbb{Z}_+$, for some fixed integer F and positive real numbers c, η we have*

$$\begin{aligned} & \mathbb{P}(W^{(n)}(t+1) = x+k \mid W^{(n)}(t) = x) \\ &= \begin{cases} \eta e^{-cnk} & \text{for } 1 \leq k \leq F, \\ \frac{1}{2} & \text{for } k = -1, \\ 0 & \text{for } k > F \text{ or } k < -1, \\ \frac{1}{2} - \sum_{k=1}^F e^{-cnk} & \text{for } k = 0. \end{cases} \end{aligned}$$

Then there exists an integer n_1 such that for all $n \geq n_1$, the Markov chain $W^{(n)}(t)$ is positive recurrent, and for any integer $s \geq 0$, we have

$$\liminf_{n \rightarrow \infty} \frac{-1}{n} \log \mathbb{P} \left(W^{(n)}(0) > s \right) \geq c(s+1),$$

where $\mathbb{P}(\cdot)$ denotes the stationary distribution of $W^{(n)}(t)$.

Proof: Omitted due to lack of space. Please see [15] for a detailed proof. ■

Figure 2 shows an example of a Markov chain referred to in Theorem 4, with $F = 2, \eta = 1$, and where the “self-loops” for the transition probabilities are not shown for simplicity. Theorem 4 says that the stationary distribution of this Markov chain is given by $\pi_m \approx \pi_0 e^{-cnm}$, in a large deviations sense as $n \rightarrow \infty$. In other words, the Markov chain is very similar to the birth-death Markov chain in that the steady-state distribution is approximately geometric.

V. ANALYSIS OF THE SSG SCHEDULING RULE

The SSG (Server-Side Greedy) scheduling algorithm was introduced in [3]. This scheduling rule is interesting because of the following reasons: (1) It is throughput-optimal for the system under very general arrival and channel processes (Theorem 5 in [3]). (2) It yields a strictly positive rate-function for the system under Assumption 1 with $K = 1$ (Theorem 7 in [3]). (3) Its computational complexity ($\mathcal{O}(n^2)$ computations per timeslot) is comparable to that of the MaxWeight rule ($\Omega(n^2)$ computations per timeslot), but the MaxWeight rule yields a zero rate-function for all integers $b \geq 0$ for the system under Assumption 1 with $K = 1$ (Theorems 3 and 8 in [3]).

For a formal definition of the rule, please check [3], Definition 3. Here is a description of the SSG rule in words: in every timeslot, the SSG algorithm allocates the servers

to the queues in multiple rounds of allocation. In round j , it allocates the server S_j to the queue that maximizes the product of the queue-length and the corresponding channel rate to the server S_k (i.e., the queue with the maximum value of $Q_i^{(j-1)}(t)X_{ij}(t)$), updates the length of the served queue, and proceeds to the next round. For the allocation decisions in the subsequent rounds, the updated lengths of the served queues are used. Note that $Q_i^{(j-1)}(t)$ denotes the length of the queue Q_i after $j-1$ rounds of server allocation.

We now analyze the rate-function performance of the SSG rule under Assumption 1 by showing that:

- 1) In any given timeslot, and starting with any configuration of queue-lengths, the maximum queue-length increases (from its starting value, before the arrivals) with a very small probability (Lemma 4).
- 2) In a constant k_0 number of timeslots, the maximum queue-length decreases with at least a constant probability (Lemma 5).

We then use Theorems 2, 3 and 4 to conclude the positivity of the rate-function and establish our main large deviations result. We first present a technical lemma that demonstrates a crucial property of the SSG rule. Note that this is a deterministic property of the SSG rule, and it does not require the number of queues or servers to be large in order to be true. This property establishes a stronger version of the following statement: if we have m queues, each connected to m servers with a channel of rate $= K$, then the maximum queue-length decreases by at least K (or becomes 0) at the end of service in that timeslot.

Lemma 3. *Under Assumption 1, let the set of queues be $\mathcal{Q} = \{Q_1, Q_2, \dots, Q_m\}$, the set of servers be $\mathcal{S} = \{S_1, S_2, \dots, S_w\}$. Fix a timeslot t , and let the queue-lengths after arrivals in that timeslot be L_1, L_2, \dots, L_m . Consider the bipartite graph $G(\mathcal{Q} \cup \mathcal{S}, \mathcal{E})$ where \mathcal{E} denotes the set of edges, where an edge is present between a queue Q_i and a server S_j if the corresponding channel supports a rate $= K$ in the timeslot t . Suppose further that every queue Q_i is connected to at least xm servers for some integer $x \geq 1$. Then with the system implementing the SSG rule, for every $i \in \{1, 2, \dots, m\}$, we have $Q_i(t) \leq \max_{1 \leq j \leq m} (L_j - xK)^+$.*

Proof: Please see Appendix E. ■

Let $\zeta(t) := \max_{1 \leq i \leq m} Q_i(t)$ denote the maximum queue-length in the system at the end of timeslot t .

Lemma 4. *Let Assumption 1 hold with $r = \lceil M/K \rceil$ for some integer $r \geq 1$. Fix*

$$\epsilon \in \left(0, \min \left(\frac{p_0}{M}, \frac{1 - \sum_{m=1}^M p_m \lceil \frac{m}{K} \rceil}{rM} \right) \right),$$

and let $\mathcal{B}_\epsilon := \{[x_1, x_2, \dots, x_{M+1}] \subseteq \mathbb{R}^{M+1} : |x_i| < \epsilon \ \forall i\}$

and $\tau := \inf_{z \in \mathcal{M}_1(\Sigma) \setminus \{[p_0, p_1, \dots, p_M] + \mathcal{B}_\epsilon\}} \sum_{m=0}^M z_m \log \frac{z_m}{p_m}$.




Fig. 2. A candidate Markov chain for the rate-function calculation in Theorem 4

Then for a system under the SSG rule, for any fixed $\rho > 0$, for n large enough, and for any timeslot t , we have

$$\begin{aligned} \mathbb{P}(\zeta(t) > \zeta(t-1)) &\leq e^{-n\tau(1-\rho)} + Mn(1-q)^{n\delta} \\ &\quad + Mn\delta \exp\left(-\frac{2n\delta}{q} H\left(\frac{q}{2} \mid q\right)\right). \end{aligned}$$

Proof: Please see Appendix F. ■

Lemma 5. Let Assumption 1 hold. Fix any timeslot t , and let $\zeta(t) = k$ for some integer k . Then for a system using the SSG rule, there exists a constant integer k_0 such that for all n large enough, we have

$$\mathbb{P}(\zeta(t+k_0) < k \mid \zeta(t) = k, k > 0) \geq \frac{1}{2}.$$

Proof: This proof is on the same lines as that of Lemma 7 in [3], and has been omitted to avoid repetition. ■

Now we are in a position to quantify the rate-function performance of the SSG algorithm.

Theorem 5. Let Assumption 1 hold with $r = \left\lfloor \frac{M}{K} \right\rfloor$ for some integer $r \geq 1$. Let ϵ, τ be as specified in the statement of Lemma 4. Then for a system using the SSG rule, for any integer $b \geq 0$, we have

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{-1}{n} \log \mathbb{P}\left(\max_{1 \leq i \leq n} Q_i(0) > b\right) \\ \geq \frac{b+1}{M} \min\left(\tau, \delta \log \frac{1}{1-q}, \frac{2\delta H\left(\frac{q}{2} \mid q\right)}{q}\right) > 0. \end{aligned}$$

Proof: Please see Appendix G. ■

An immediate strengthening of the above result is obtained by maximizing the RHS over the appropriate ranges for τ and δ . We have thus established that under the SSG algorithm, the small-buffer overflow event has a strictly positive rate-function for all integers $b \geq 0$.

VI. SIMULATION RESULTS

We compare the performance of the proposed SSG algorithm with the classic MaxWeight algorithm in [12]. We choose the MaxWeight algorithm as a benchmark because it is the best-known throughput-optimal algorithm. We consider a system with $n = 50$ queues and servers. The channel parameters are set to $K = 2$ and $q = 0.5$. We parameterize the arrival process as follows: set $M = 3$ as the maximum number of arrivals. Vary the parameter p_M , and set $p_i = (1-p_M)/M$

for $0 \leq i \leq M-1$. Let $\eta := \sum_{m=1}^M p_m \left\lceil \frac{m}{K} \right\rceil$. We consider systems with progressively higher values of η , run the simulations for 1×10^5 timeslots, and plot the empirical packet-delay probabilities. The results are summarized in Figure 3.

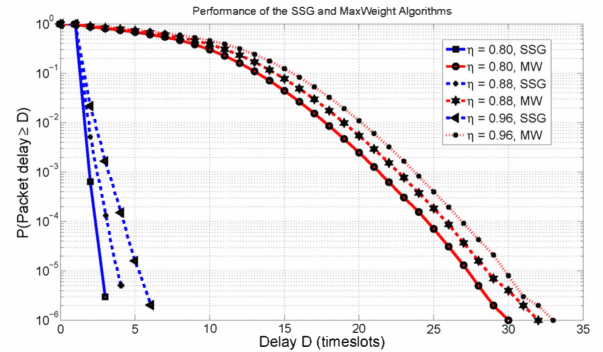


Fig. 3. SSG v/s MaxWeight: Packet delay profiles

We see that there is a significant difference between the performance of the two algorithms, with SSG being the clear winner in this case. Similar results hold for the buffer overflow probabilities under the two algorithms. Due to space constraints, and also because this is mainly a theoretical paper, we have not reported more extensive simulation results, but the SSG algorithm continues to outperform the MaxWeight algorithm under a variety of channel and arrival processes.

VII. CONCLUSIONS

We considered the problem of designing scheduling algorithms for multi-user multi-channel wireless downlink networks. The main result is that the iterative resource allocation rule (the SSG rule) provides a very good delay performance to the users, in addition to network stability. These systems where the channels serve multiple packets per timeslot are inherently more complex than the ones considered in the earlier works ([2] and [3]). Due to space constraints, we considered a simple arrival and channel process that captures the essence of the multi-rate channel allocation problem, and developed Markov chain-based analysis techniques that are applicable to a much wider class of systems and scheduling rules. Indeed, a number of natural extensions of this work are possible, including the cases where the number of queues and the servers is unequal, where the frequency bands are time or frequency-correlated, or when the scheduling rule is more complex (for example, the matching-based iLQF-class scheduling rules introduced in [2]), among other possibilities.

VIII. ACKNOWLEDGMENTS

This work is supported in part by NSF grants CNS-0963807, CNS-0831756, CNS-0721286, CNS-0347400, CNS-0721380, DTRA grants HDTRA1-09-1-0055, HDTRA1-08-1-0016, and the Darpa ITMANET program.

APPENDIX A

PROOF OF THEOREM 1

Consider the following event which implies $\{Q_1(0) > b\}$ under any scheduling rule: for $m := \left(\left\lfloor \frac{b}{M} \right\rfloor + 1\right)$ consecutive timeslots before (and including) timeslot 0, there are M arrivals per timeslot to Q_1 , and all the channels connecting Q_1 to the servers are OFF (i.e., $X_{1j}(t) = 0$) in each of these timeslots. The probability of this event is $p_M^m(q_0)^{nm}$, and the result follows.

APPENDIX B

PROOF OF LEMMA 2

If possible, let $\liminf_{n \rightarrow \infty} \frac{-1}{n} \log \mathbb{P} \left(\max_{1 \leq i \leq n} Q_i(0) > 0 \right) = \delta > 0$

under some scheduling algorithm. Then for n large enough, $\mathbb{P} \left(\max_{1 \leq i \leq n} Q_i(0) > 0 \right) \leq e^{-n\delta/2}$, implying that the long-term average number of timeslots for which we have $\max_{1 \leq i \leq n} Q_i(t) = 0$ is at least $1 - e^{-n\delta/2}$, or greater than 0.8 for n large enough.

Suppose at the end of some timeslot T , we have $\max_{1 \leq i \leq n} Q_i(T) = 0$. Fix $\epsilon > 0$ such that

$$[p'_0, p'_1, \dots, p'_M] = [p_0 + M\epsilon, p_1 - \epsilon, p_2 - \epsilon, \dots, p_M - \epsilon]$$

is a probability vector with strictly positive components, and $\sum_{i=1}^M p'_i \lceil i/K \rceil > 1$.

Let n_m denote the number of queues that receive m packets in timeslot $T+1$. By Sanov's theorem ([5], Theorem 2.1.10), we know that there exists $\zeta > 0$ such that for n large enough,

$$\mathbb{P} \left(\frac{n_m}{n} \geq p'_m \quad \forall m \in \{1, 2, \dots, M\} \right) \geq 1 - e^{-n\zeta}.$$

For a queue with length = m , any algorithm must allocate at least $\lceil m/K \rceil$ servers to serve all packets from it. But if the algorithm tries to allocate this way, then (even assuming that each of the servers has a channel with rate = K to serve any one of the queues, i.e., $X_{ij}(T+1) = K$ for all i, j), the number of servers necessary to drain all the packets is

$$\sum_{m=1}^M n_m \left\lceil \frac{m}{K} \right\rceil \geq n \sum_{m=1}^M p'_m \left\lceil \frac{m}{K} \right\rceil > n.$$

Hence, if at the end of timeslot T we have $\max_{1 \leq i \leq n} Q_i(T) = 0$, then at the end of timeslot $T+1$, with probability at least $1 - e^{-n\zeta}$, we have $\max_{1 \leq i \leq n} Q_i(T+1) \geq 1$ for some $\zeta > 0$. Elementary calculations imply that the long-term fraction of timeslots for which $\{\max_{1 \leq i \leq n} Q_i(t) = 0\}$ is no more than $\frac{1}{2 + e^{-n\zeta}} < \frac{2}{3}$, for n large enough. This contradicts the earlier claim that the long-term fraction of timeslots for which $\max_{1 \leq i \leq n} Q_i(t) = 0$ is at least 0.8, completing the proof.

APPENDIX C

PROOF OF THEOREM 2

A stationary distribution of $X^*(t)$ is given by

$$\mathbb{P}(X^*(t) = x) = \mathbb{P}(\mathbf{X}(t) \in \mathcal{W}_x) = \sum_{z \in \mathcal{W}_x} \sigma(z).$$

We show that this is a stationary distribution of the Markov chain $Y(t)$. By the detailed flow-balance for the Markov chain $\mathbf{X}(t)$, for every non-negative integer x , we have

$$\sum_{z \in \mathcal{W}_x} \sum_{\substack{y \geq 0 \\ y \neq x}} \sum_{v \in \mathcal{W}_y} \sigma(z) p(z, v) = \sum_{z \in \mathcal{W}_x} \sum_{\substack{y \geq 0 \\ y \neq x}} \sum_{v \in \mathcal{W}_y} \sigma(v) p(v, z). \quad (1)$$

In order that $\mathbb{P}(Y(t) = x) = \mathbb{P}(X^*(t) = x) = \sum_{z \in \mathcal{W}_x} \sigma(z)$, it is necessary and sufficient that the proposed stationary distribution satisfies the detailed flow-balance equations for the Markov chain $Y(t)$, is non-negative and adds to 1. In other words, for every non-negative integer x , we want

$$\begin{aligned} & \sum_{\substack{y \geq 0 \\ y \neq x}} \left(\sum_{z \in \mathcal{W}_x} \sigma(z) \right) \mathbb{P}(Y(t+1) = y \mid Y(t) = x) \\ &= \sum_{\substack{y \geq 0 \\ y \neq x}} \left(\sum_{z \in \mathcal{W}_y} \sigma(z) \right) \mathbb{P}(Y(t+1) = x \mid Y(t) = y). \end{aligned}$$

We now verify that the proposed stationary distribution $\mathbb{P}(Y(t) = x) = \mathbb{P}(X^*(t) = x) = \sum_{z \in \mathcal{W}_x} \sigma(z)$ satisfies this condition, as follows:

$$\begin{aligned} & \sum_{\substack{y \geq 0 \\ y \neq x}} \left(\sum_{z \in \mathcal{W}_x} \sigma(z) \right) \mathbb{P}(Y(t+1) = y \mid Y(t) = x) \\ &= \sum_{\substack{y \geq 0 \\ y \neq x}} \mathbb{P}(\mathbf{X}(t+1) \in \mathcal{W}_y, \mathbf{X}(t) \in \mathcal{W}_x) \\ &= \sum_{\substack{y \geq 0 \\ y \neq x}} \sum_{z \in \mathcal{W}_x} \sum_{v \in \mathcal{W}_y} \sigma(z) p(z, v) \\ &\stackrel{(a)}{=} \sum_{\substack{y \geq 0 \\ y \neq x}} \sum_{z \in \mathcal{W}_x} \sum_{v \in \mathcal{W}_y} \sigma(v) p(v, z) \\ &= \sum_{\substack{y \geq 0 \\ y \neq x}} \mathbb{P}(\mathbf{X}(t+1) \in \mathcal{W}_x, \mathbf{X}(t) \in \mathcal{W}_y) \\ &= \sum_{\substack{y \geq 0 \\ y \neq x}} \left(\sum_{z \in \mathcal{W}_y} \sigma(z) \right) \mathbb{P}(Y(t+1) = x \mid Y(t) = y), \end{aligned}$$

where the step (a) follows from Equation (1). Note also that $\sum_{z \in \mathcal{W}_x} \sigma_z \geq 0$ and $\sum_{x \geq 0} \sum_{z \in \mathcal{W}_x} \sigma_z = 1$, since the sets \mathcal{W}_x form a partition of the state-space of the original Markov chain $\mathbf{X}(t)$. Hence, the proof is complete.

APPENDIX D
PROOF OF THEOREM 3

The proof proceeds by induction. Let the claim hold for some $t = k - 1 \geq 0$, i.e., $Y(k - 1) \leq_{st} W(k - 1)$. We have

$$\begin{aligned} Y(k) &= Y(k - 1) + Z_{k-1}(Y(k - 1)), \\ W(k) &= W(k - 1) + \tilde{Z}_{k-1}(W(k - 1)). \end{aligned}$$

By a standard result in stochastic ordering, there exist random variables $\hat{Y}(k - 1)$ and $\hat{W}(k - 1)$ such that

$$\hat{Y}(k - 1) \stackrel{d}{=} Y(k - 1), \quad \hat{W}(k - 1) \stackrel{d}{=} W(k - 1),$$

and $\hat{Y}(k - 1) \leq \hat{W}(k - 1)$, a.s.

$$\begin{aligned} \text{Define } \hat{Y}(k) &:= \hat{Y}(k - 1) + Z_{k-1}(\hat{Y}(k - 1)), \\ \hat{W}(k) &:= \hat{W}(k - 1) + \tilde{Z}_{k-1}(\hat{W}(k - 1)). \end{aligned}$$

Now,

$$\mathbb{P}(W(k) = j) = \sum_i \mathbb{P}(i + \tilde{Z}_{k-1}(i) = j) \mathbb{P}(W(k - 1) = i),$$

and

$$\mathbb{P}(\hat{W}(k) = j) = \sum_i \mathbb{P}(i + \tilde{Z}_{k-1}(i) = j) \mathbb{P}(\hat{W}(k - 1) = i).$$

But $\mathbb{P}(\hat{W}(k - 1) = i) = \mathbb{P}(W(k - 1) = i)$, so $\hat{W}(k) \stackrel{d}{=} W(k)$, and similarly $\hat{Y}(k) \stackrel{d}{=} Y(k)$. By our assumption of $Z_{k-1}(\cdot)$ and $\tilde{Z}_{k-1}(\cdot)$, for all i, j, ℓ ,

$$\begin{aligned} \mathbb{P}(\hat{W}(k) - i > \ell \mid \hat{W}(k - 1) = i, \hat{Y}(k - 1) = j) \\ \geq \mathbb{P}(\hat{Y}(k) - j > \ell \mid \hat{W}(k - 1) = i, \hat{Y}(k - 1) = j), \end{aligned}$$

implying

$$\begin{aligned} \mathbb{P}(\hat{W}(k) > \ell + i \mid \hat{W}(k - 1) = i, \hat{Y}(k - 1) = j) \\ \geq \mathbb{P}(\hat{Y}(k) > \ell + i \mid \hat{W}(k - 1) = i, \hat{Y}(k - 1) = j) \end{aligned}$$

for all (i, j) such that $i \geq j$. (Note that $\{\hat{Y}(k) > \ell + i\} \Rightarrow \{\hat{Y}(k) > \ell + j\}$ for $i \geq j$.) Multiplying both sides by $\mathbb{P}(\hat{W}(k - 1) = i, \hat{Y}(k - 1) = j)$ and summing over all i, j such that $i \geq j$, and noting that $\mathbb{P}(\hat{W}(k - 1) \geq \hat{Y}(k - 1)) = 1$, we get

$$\mathbb{P}(\hat{W}(k) > \ell) \geq \mathbb{P}(\hat{Y}(k) > \ell) \quad \forall \ell \in \mathbb{Z}_+, \text{ i.e., } \hat{Y}(k) \leq_{st} \hat{W}(k).$$

Since $\hat{Y}(k) \stackrel{d}{=} Y(k)$ and $\hat{W}(k) \stackrel{d}{=} W(k)$,

we get $Y(k) \leq_{st} W(k)$, completing the proof of stochastic dominance by induction. Since the limiting distributions of $Y(t)$ and $W(t)$ are the stationary distributions $\mu_Y(\cdot)$ and $\mu_W(\cdot)$ respectively (the convergence to the stationary distribution follows from [4], Chapter 4, Theorem 2.1), for random variables Y, W which are distributed according to $Y \sim \mu_Y(\cdot)$ and $W \sim \mu_W(\cdot)$, we have $Y \leq_{st} W$. Hence the proof of the theorem is complete.

APPENDIX E
PROOF OF LEMMA 3

Fix any $i \in \{1, 2, \dots, m\}$. Let the set of servers connected to Q_i be $\mathcal{S}^* = \{S_1^i, S_2^i, \dots, S_{x_m}^i\}$, in the increasing order of server-indices.

Case 1: The queue Q_i is allocated x or more servers.

In this case, because each server drains K packets, we have $Q_i(t) \leq (L_i - xK)^+$, implying

$$Q_i(t) \leq (L_i - xK)^+ \leq \max_{1 \leq j \leq m} (L_j - xK)^+,$$

and the claim holds.

Case 2: The queue Q_i is allocated at most $x - 1$ servers.

In this case, by the pigeonhole principle, at least one of the queues $Q_j \neq Q_i$ is allocated $x + 1$ (or more) servers from the set \mathcal{S}^* . Consider the round y of SSG when the $(x + 1)^{th}$ server from the set \mathcal{S}^* is allocated to Q_j . Since the SSG rule allocates a server to a longest queue among the queues it can serve, we have $Q_i^{(y-1)}(t) \leq Q_j^{(y-1)}(t) \leq (L_j - xK)^+$, and the result follows since $Q_i(t) \leq Q_i^{(y-1)}(t)$.

APPENDIX F
PROOF OF LEMMA 4

Let $\zeta(t - 1) = c$. By the choice of ϵ ,

$$[p'_0, p'_1, \dots, p'_M] = [p_0 - M\epsilon, p_1 + \epsilon, p_2 + \epsilon, \dots, p_M + \epsilon]$$

is a probability vector with strictly positive components, and $\sum_{i=1}^M p'_i \lceil i/K \rceil < 1$. Further, since the set $\mathcal{M}_1(\Sigma) \setminus \{[p_0, p_1, \dots, p_M] + \mathcal{B}_\epsilon\}$ is compact and the function $g(\mathbf{z}) := \sum_{m=0}^M z_m \log \frac{z_m}{p_m}$ is lower semicontinuous ([5], Chapter 2, Exercise 2.1.22), the infimum of $g(\mathbf{z})$ (in the definition of τ) is achieved and is strictly positive, since $g([p_0, p_1, \dots, p_M]) = 0$ and $g(\mathbf{z}) > 0$ for all other values of \mathbf{z} , and

$$[p_0, p_1, \dots, p_M] \notin \mathcal{M}_1(\Sigma) \setminus \{[p_0, p_1, \dots, p_M] + \mathcal{B}_\epsilon\}$$

After arrivals in timeslot t , let \mathcal{F}_m be the set of queues whose length is at least $c + m$, and define $f_m := |\mathcal{F}_m|/n$. Consider the event

$$E := \left\{ f_m \leq \sum_{i=m}^M p'_i \quad \forall m \in \{1, 2, \dots, M\} \right\}.$$

From Sanov's theorem, we know that for n large enough, $\mathbb{P}(E^c) \leq e^{-n\tau(1-\rho)}$ for any fixed $\rho > 0$. We condition the rest of this proof on the event E .

Claim 1. *Conditioned on the event E , fix any integer $m \geq 1$, and any constant*

$$\delta \in \left(0, q \left(1 - \sum_{i=1}^M p'_i \left\lceil \frac{i}{K} \right\rceil \right) / (M(2 - q)) \right).$$

Then after $n \left(\sum_{j=0}^{\infty} f_{m+jK} + (M - m + 1)\delta \left(\frac{2}{q} - 1 \right) \right)$ rounds of server allocation under the SSG rule, with probability at least

$$1 - Mn(1 - q)^{n\delta} - Mn\delta e^{-\frac{2n\delta}{q} H(\frac{q}{2}|q)},$$

there remain no queues of (updated) length $c + m$ or more.

Note that only a finite number of terms in the above infinite summation are nonzero.

Proof: Let $\mathcal{F}_m^{(i)}$ denote the updated set \mathcal{F}_m after i rounds of server allocation, that is, the set of queues at length $c + m$ or more, after i rounds of server allocation. First we consider the case $m = M$.

Case 1: $|\mathcal{F}_M| = |\mathcal{F}_M^{(0)}| > n\delta$.

For $1 \leq i \leq m_0 := |\mathcal{F}_M^{(0)}| - n\delta$, consider the event

$$W_i = \{X_{r_i}(t) < K \quad \forall Q_r \in \mathcal{F}_M^{(i-1)}\}.$$

By the independence of channel allocation decision and the channel realizations for the higher-indexed servers,

$$\mathbb{P}(W_i | W_1^c, W_2^c, \dots, W_{i-1}^c) = (1-q)^{(|\mathcal{F}_M^{(0)}| - i + 1)} \leq (1-q)^{n\delta}.$$

If the server S_i is connected to any of the queues in the set $\mathcal{F}_M^{(i-1)}$, then (because the set $\mathcal{F}_M^{(i-1)}$ is the set of the ‘‘current’’ longest queues) the server S_i is allocated to a queue in $\mathcal{F}_M^{(i-1)}$, and the served queue is removed from the set $\mathcal{F}_M^{(i-1)}$ to obtain the set $\mathcal{F}_M^{(i)}$. Therefore, using the union bound, with probability at least $1 - (|\mathcal{F}_M^{(0)}| - n\delta)(1-q)^{n\delta} \geq 1 - n(1-q)^{n\delta}$, at the end of $|\mathcal{F}_M^{(0)}| - n\delta$ rounds of service, we have $|\mathcal{F}_M^{(m_0)}| \leq n\delta$.

Consider the set of servers

$$\mathcal{S}^* = \{S_{m_0+1}, S_{m_0+2}, \dots, S_{m_0+2n\delta/q}\}.$$

By the Chernoff bound, the probability that a given queue $Q_i \in \mathcal{F}_M^{(m_0)}$ is connected

- with a channel of rate $= K$
- to at least $n\delta$ of the servers in the set \mathcal{S}^*

is at least $1 - e^{-\frac{2n\delta}{q}H(\frac{q}{2}|q)}$. Therefore, by Lemma 3, at the end of $2n\delta/q$ further rounds of service, with probability at least $1 - n\delta e^{-\frac{2n\delta}{q}H(\frac{q}{2}|q)}$, the maximum queue-length in the system is no more than $c + m - 1$, completing the proof for this case by the union bound.

Case 2: $|\mathcal{F}_M| = |\mathcal{F}_M^{(0)}| \leq n\delta$.

Following the analysis for case 1, it is clear that the claim holds in this case, completing the proof of the claim for the case $m = M$.

The proof of the claim now follows by repeatedly applying the above procedure to the cases $m = M - 1, M - 2, \dots, 1, 0$, using the union bound, and noting that if a queue Q_i of length ℓ is served by a server S_j , then its length decreases by exactly K (or becomes zero), so that it continues to belong to the sets $\mathcal{F}_{\ell-K}^{(j)}, \mathcal{F}_{\ell-K-1}^{(j)}$, and so on. ■

Applying the result of the claim to the case $m = 1$, we get

$$\begin{aligned} & \sum_{j=0}^{\infty} f_{1+jK} + M\delta \left(\frac{2}{q} - 1 \right) \\ &= f_1 + f_{K+1} + \dots + f_{(r-1)K+1} + M\delta \left(\frac{2}{q} - 1 \right) \\ &\stackrel{(a)}{\leq} (p'_1 + p'_2 + \dots + p'_K) + 2(p'_{K+1} + \dots + p'_{2K}) + \dots \\ &\quad + r(p'_{(r-1)K+1} + \dots + p'_M) + M\delta \left(\frac{2}{q} - 1 \right) \end{aligned}$$

$$= \sum_{m=1}^M p'_m \left\lceil \frac{m}{K} \right\rceil + M\delta \left(\frac{2}{q} - 1 \right) < 1$$

where the inequality (a) holds from Sanov’s theorem (event E as defined above) with probability at least $1 - e^{-n\tau(1-\rho)}$, and the last inequality holds by the choice of δ . Hence, by the union bound, we have (for n large enough)

$$\begin{aligned} \mathbb{P}(\zeta(t) > \zeta(t-1)) &\leq e^{-n\tau(1-\rho)} + Mn(1-q)^{n\delta} \\ &\quad + Mn\delta \exp\left(-\frac{2n\delta}{q}H\left(\frac{q}{2} \middle| q\right)\right), \end{aligned}$$

completing the proof.

APPENDIX G PROOF OF THEOREM 5

This proof proceeds in three steps, where we use Theorems 2, 3 and 4 to arrive at the desired conclusion.

Step 1: We have

$$\mathbb{P}\left(\max_{1 \leq i \leq n} Q_i(0) > b\right) = \sum_{k=b+1}^{\infty} \mathbb{P}\left(\max_{1 \leq i \leq n} Q_i(0) = k\right).$$

Our aim is to calculate an upper-bound on the RHS of the above inequality. To this end, we know from Theorem 2 that $\mathbb{P}(\max_{1 \leq i \leq n} Q_i(0) = k)$ is the same as $\mathbb{P}(Y(t) = k)$, where the (one-dimensional) Markov chain has the following transition probability structure:

$$\begin{aligned} \mathbb{P}(Y(t+1) = y | Y(t) = x) \\ = \mathbb{P}\left(\max_{1 \leq i \leq n} Q_i(t+1) = y \middle| \max_{1 \leq i \leq n} Q_i(t) = x\right). \end{aligned}$$

Step 2: As a result of Lemma 5, we know that there exists a constant integer k_0 such that over k_0 consecutive timeslots and for n large enough, the probability that the maximum queue-length in the system makes a transition from a higher value k to a lower value j (provided $k > 0$) is at least 0.5.

Fix any constant $\rho > 0$. From Lemma 4, for n large enough, in a given timeslot, the probability that the maximum queue-length increases is at most

$$e^{-n\tau(1-\rho)} + Mn(1-q)^{n\delta} + Mn\delta \exp\left(-\frac{2n\delta}{q}H\left(\frac{q}{2} \middle| q\right)\right).$$

Let

$$c := \min\left(\tau(1-\rho), \delta \log \frac{1}{1-q}, \frac{2\delta H(\frac{q}{2}|q)}{q}\right).$$

Fix any $\epsilon' \in (0, c)$ and let $c' = c - \epsilon'$. For any given $\epsilon > 0$ there exists n_0 large enough such that for all $n \geq n_0$, we have

$$e^{-n\tau(1-\rho)} + Mn(1-q)^{n\delta} + Mn\delta e^{-\frac{2n\delta}{q}H(\frac{q}{2}|q)} \leq e^{-nc'}.$$

Define $L_i(t) := Q_i(k_0 t)$ and consider a Markov chain

$$L(t) = [L_1(t), L_2(t), \dots, L_n(t)].$$

The Markov chain $L(t)$ is a time-sampled version of the Markov chain $[Q_1(t), Q_2(t), \dots, Q_n(t)]$, and the two Markov

chains have the same stationary distribution. Let $L^*(t) = \max_{1 \leq i \leq n} L_i(t)$. Define

$$p(k, j) := \mathbb{P}(L^*(t+1) = j \mid L^*(t) = k).$$

Then for any integer t and for any integer $k > 0$, we have $\sum_{j=0}^{k-1} p(k, j) \geq 0.5$.

Further, the probability that over a period of k_0 consecutive timeslots the maximum queue-length increases in at least $a \leq k_0$ timeslots is at most $\binom{k_0}{a} e^{-anc'}$ by the union bound. More precisely, for $1 \leq i \leq k_0$, consider the events $E_i := \{\zeta(i) > \zeta(i-1)\}$. We know from Lemma 4 that regardless of the queue-lengths at the beginning of timeslot i , and any other events in any other timeslots, $\mathbb{P}(E_i) \leq e^{-nc'}$ for n large enough. Hence,

$$\mathbb{P}\left(\bigcap_{i=1}^a E_i\right) = \prod_{i=1}^a \mathbb{P}(E_i \mid E_1, E_2, \dots, E_{i-1}) \leq e^{-anc'},$$

and the claimed bound follows from the union bound since there are $\binom{k_0}{a}$ ways to “choose” the candidate timeslots where the maximum queue-length increases.

In a given timeslot, the length of any queue can increase by at most M , so that for any timeslot t if $\zeta(t) = k$, then $\zeta(t+1) = s$ where $s \leq k+M$. It follows that for $1 \leq a \leq k_0$, we have

$$\sum_{j=(a-1)M+1}^{aM} p(k, k+j) \leq \binom{k_0}{a} e^{-anc'} \leq 2^{k_0} e^{-anc'}.$$

We now apply Theorem 3 with the random variables $Z_t(k)$ and $\tilde{Z}_t(k)$ with the following distributions:

$$\mathbb{P}(Z_t(k) = j) = p(k, k+j), \quad j \in \mathbb{Z},$$

for integers $1 \leq a \leq k_0$:

$$\mathbb{P}(\tilde{Z}_t(k) = j) = 2^{k_0} e^{-anc'}, \quad (a-1)M+1 \leq j \leq aM,$$

$$\text{and} \quad \mathbb{P}(\tilde{Z}_t(k) = -1) = 0.5.$$

By the foregoing analysis, we have $Z_t(i) \leq_{st} \tilde{Z}_t(j)$ for all $i, j \in \mathbb{Z}_+$ and all $t \in \mathbb{Z}$, implying $L^*(t) \leq_{st} W(t)$. Further, the Markov chain $W(t)$ has a stationary distribution for n large enough (follows from Theorem 4), and the Markov chain $L(t)$ has a stationary distribution by the stability of the Markov chain $Q(t)$ (follows from Theorem 5 in [3]). Both the Markov chains $W(t)$ and $L(t)$ are clearly aperiodic and irreducible, so is the (hypothetical) Markov chain whose transition probabilities are defined by the transition probabilities of $L^*(t)$. Hence, by Theorem 2.1, Chapter 4 in [4], the distributions of $L^*(t)$ and $W(t)$ converge to their respective stationary distributions, implying that the stationary distribution of $L^*(t)$ is stochastically dominated by the stationary distribution of $W(t)$.

Step 3: We apply Theorem 4 with $\eta = 2^{k_0}$ and $F = k_0 M$ to conclude that

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \frac{-1}{n} \log \mathbb{P}(W(0) > b) \\ &= \liminf_{n \rightarrow \infty} \frac{-1}{n} \log \left(\sum_{k=b+1}^{\infty} \mathbb{P}(W(0) = k) \right) \geq \frac{c'(b+1)}{M}, \end{aligned}$$

implying

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \frac{-1}{n} \log \mathbb{P} \left(\max_{1 \leq i \leq n} Q_i(0) > b \right) \\ & \geq \frac{b+1}{M} \left[\min \left(\tau(1-\rho), \delta \log \frac{1}{1-q}, \frac{2\delta H(\frac{q}{2}|q)}{q} \right) - \epsilon' \right] > 0, \end{aligned}$$

and completing the proof since the last inequality holds for all $\epsilon' \in (0, c)$ for some $c > 0$, and for every constant $\rho > 0$.

REFERENCES

- [1] 3GPP TR 25.913. Requirements for Evolved UTRA (E-UTRA) and Evolved UTRAN (E-UTRAN). March 2006.
- [2] S. Bodas, S. Shakkottai, L. Ying, and R. Srikant. Scheduling in Multi-Channel Wireless Networks: Rate Function Optimality in the Small-Buffer Regime. In *Proc. SIGMETRICS/Performance Conf.*, Jun. 2009.
- [3] S. Bodas, S. Shakkottai, L. Ying, and R. Srikant. Low-complexity Scheduling Algorithms for Multi-channel Downlink Wireless Networks. In *Proc. IEEE Infocom*, Mar. 2010.
- [4] Pierre Brémaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer Science + Business Media, LLC, 2008.
- [5] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Springer-Verlag New York, Inc., second edition, 1998.
- [6] WiMax Forum. Mobile WiMAX Part I: A technical overview and performance evaluation. March 2006. White Paper.
- [7] S. Kittipiyakul and T. Javidi. Delay-Optimal Server Allocation in Multiqueue Multiserver Systems with Time-Varying Connectivities. *IEEE Trans. Inform. Theory*, 55(5):2319 – 2333, May 2009.
- [8] S.P. Meyn. Stability and asymptotic optimality of generalized maxweight policies. *SIAM J. Control and Optimization*, 2008. to appear.
- [9] S. Shakkottai, R. Srikant, and A. Stolyar. Pathwise optimality of the exponential scheduling rule for wireless channels. *Ann. Appl. Prob.*, 36(4):1021–1045, December 2004.
- [10] A. Stolyar. MaxWeight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *Ann. Appl. Prob.*, 14(1), 2004.
- [11] A. Stolyar. Large deviations of queues sharing a randomly time-varying server. *Queueing Systems*, 59:1–35, 2008.
- [12] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automat. Contr.*, 4:1936–1948, December 1992.
- [13] V. J. Venkataramanan and X. Lin. Structural properties of LDP for queue-length based wireless scheduling algorithms. In *Proc. Ann. Allerton Conf. Communication, Control and Computing*, Monticello, Illinois, September 2007.
- [14] L. Ying, R. Srikant, A. Eryilmaz, and G. Dullerud. A large deviations analysis of scheduling in wireless networks. *IEEE Trans. Inform. Theory*, 52(11):5088–5098, November 2006.
- [15] S. Bodas, S. Shakkottai, L. Ying, and R. Srikant. Scheduling Algorithms for Wireless Downlink Networks: Tech Report. <http://users.ece.utexas.edu/~shakkott/infocom11-long-BodShaYinSri.pdf>, 2010.