

Low-delay Wireless Scheduling with Partial Channel-state Information

Aditya Gopalan

Department of Electrical Engineering
Technion-Israel Institute of Technology
Haifa 32000, Israel
Email: aditya@ee.technion.ac.il

Constantine Caramanis
and Sanjay Shakkottai

Department of Electrical and Computer Engineering
The University of Texas at Austin, USA
Email: {caramanis, shakkottai}@mail.utexas.edu

Abstract—We consider a server serving a time-slotted queued system of multiple packet-based flows, where not more than one flow can be serviced in a single time slot. The flows have exogenous packet arrivals and time-varying service rates. At each time, the server can observe instantaneous service rates for only a *subset* of flows (selected from a fixed collection of *observable* subsets) before scheduling a flow in the subset for service. We are interested in queue-length aware scheduling to keep the queues short. The limited availability of instantaneous service rate information requires the scheduler to make a careful choice of which subset of service rates to sample. We develop scheduling algorithms that use only partial service rate information from subsets of channels, and that minimize the likelihood of queue overflow in the system. Specifically, we present a new joint subset-sampling and scheduling algorithm called *Max-Exp* that uses only the current queue lengths to pick a subset of flows, and subsequently schedules a flow using the Exponential rule. When the collection of observable subsets is disjoint, we show that *Max-Exp* achieves the best exponential decay rate, among all scheduling algorithms using partial information, of the tail of the longest queue in the system. To accomplish this, we introduce novel analytical techniques for studying the performance of scheduling algorithms using partial state information, that are of independent interest. These include new sample-path large deviations results for processes obtained by nonrandom, predictable sampling of sequences of independent and identically distributed random variables, which show that scheduling with partial state information yields a rate function significantly different from the case of full information. As a special case, *Max-Exp* reduces to simply serving the flow with the longest queue when the observable subsets are singleton flows, i.e., when there is effectively no *a priori* channel-state information; thus, our results show that this greedy scheduling policy is large-deviations optimal.

I. INTRODUCTION

Next-generation wireless cellular systems such as LTE-Advanced [1] and WiMAX [2] promise high-speed packet-switched data services for a variety of applications, including file transfer, peer-to-peer sharing and real-time audio/video streaming. This demands effective scheduling in typical wireless environments with time-varying channels and limited resources, to guarantee high data rates to the users. Together with maximizing data rates or throughput, the scheduling algorithm at the cellular base station must keep packet delays in the system low, in order to support highly delay-sensitive applications like real-time video streaming.

There has been much recent work to develop wireless scheduling algorithms with optimal throughput and/or delay performance [3–7]. Such opportunistic scheduling algorithms utilize instantaneous wireless Channel State Information (CSI) from all users to make good scheduling decisions. However, in a practical situation with a large number of users in the network, channel state feedback resources could potentially be limited, i.e., it might be infeasible to acquire complete instantaneous CSI from all channels due to bandwidth and latency limitations. Instead, it might be possible to request CSI feedback from only a *subset of users* each time. Thus, it is important to develop algorithms that can schedule using only partial CSI rather than complete CSI, and at the same time afford the best possible delay performance.

Using partial CSI – from subsets of channels – entails a new dimension of opportunism in wireless scheduling. The scheduling algorithm needs to make a careful choice of which subsets to sample, together with how to use the sampled CSI for scheduling. Recently, natural extensions of complete-CSI scheduling algorithms to the partial-CSI setting have shown to have throughput-optimal properties [8], yet it is not clear how they perform in the sense of packet delays. The general structure of low-delay, partial-CSI scheduling algorithms remains unknown, i.e., how an algorithm should choose “good” subsets of channels, whether any additional backlog or statistical information is needed for picking subsets, and if so, how much, how users should be scheduled in the observed subset etc.

In this work, we develop algorithms for wireless scheduling that use only partial CSI, i.e., from subsets of channels, and that also enjoy high performance guarantees. We consider a wireless downlink where a base station schedules users using partial CSI from subsets of channels. Viewing the system queue lengths as a surrogate for packet delays, we seek scheduling strategies that can keep the longest queue in the system as short as possible, i.e., minimize the likelihood of overflow of the longest queue. We design a new scheduling algorithm, that we term *Max-Exp*, that obtains partial CSI relying on just current queue lengths and no other auxiliary information. Employing sample-path large deviations techniques, we show that when the observable channel subsets are disjoint, *Max-Exp* yields the best decay rate for the longest-

queue overflow probability, across all scheduling strategies which use subset-based CSI to schedule users. To the best of our knowledge, this is the first work that analyzes queue-overflow performance for scheduling with the information structure of partial CSI, and that provides a simple scheduling algorithm needing no extra statistical information which is actually rate-function optimal for buffer overflow.

From a technical standpoint, sample-path large deviations techniques have successfully been used to analyze wireless scheduling algorithms [5–7, 9]; yet, significant new analytical challenges emerge when studying the large deviations behavior of scheduling strategies that cannot access the full state of the system. A chief difference in this regard arises from the fact that when scheduling is carried out by observing the *complete* state/randomness of the system, large deviations occur depending on how the scheduler responds to atypical channel state behavior. In other words, a natural cause-effect relationship between the channel state process and scheduling actions is the basis for the analysis of large deviations performance. On the other hand, when *partial* channel state is *acquired* selectively by a scheduling algorithm, this cause-effect sequence is reversed – it is the algorithm that first decides what part of the channel state to sample; subsequently, this dynamic portion of the channel state can respond by behaving atypically. Viewed differently, the scheduling information structure no longer falls into an “experts” setting (all channel rates known in advance), but rather into a “bandit” setting (only chosen channel rates known), implying a fundamental change in the large deviations dynamics. Indeed, we are able to show that this difference results in a significantly different rate function than that encountered in the former complete-CSI case.

Also, the standard approach of analyzing queue overflow probability exponents using continuity of queue-length/delays as functions of the arrivals and channel processes [10, 11] becomes cumbersome due to the complex two-stage sampling and scheduling structure of scheduling with partial CSI. Thus, we are led to develop new sample-path large deviations results for processes with dynamically (and predictably) sampled randomness, which help to bound the resulting rate functions via connections to appropriate variational problems. We believe that these techniques and results are of independent interest as tools to analyze the behavior of scheduling policies that can only sample parts of the system state.

A. Related Work

For scheduling with complete CSI, there is a rich body of work on throughput-optimal scheduling algorithms, starting from the pioneering approach of Tassiulas et al. [3] to develop the Backpressure algorithm. A host of scheduling algorithms such as Max-Weight/Backpressure [3, 4], the Exponential rule [5, 12, 13] and the Log rule [6] have been developed for scheduling using full CSI. Many optimality results are now known for the delay/queue-length performance of the above full-CSI algorithms. These include expected queue length/delay bounds via Lyapunov function techniques

[14, 15], tail probability decay rates for queue lengths [5–7, 10, 11, 16, 17], heavy-traffic optimality [18] etc.

Throughput-maximizing scheduling has been studied with different forms of partial CSI, including infrequent channel state measurements [19], group/random-access based quantized channel state feedback [20, 21], optimal channel state probing with costs [22, 23], delayed CSI [24] and subset-based CSI [8]. However, to date, neither the structure nor performance results for queue overflow tails under scheduling with partial CSI are known.

B. Contributions

We describe a new scheduling algorithm – Max-Exp – for wireless scheduling with Channel State Information (CSI) restricted to a collection of *observable channel subsets*. The Max-Exp algorithm uses a suitable subset-selection strategy along with the well-known Exponential rule [12] to schedule users. We show the following:

- (a) We derive a lower bound on the rate function for overflow of the longest queue under the Max-Exp scheduling algorithm, using sample-path large deviations tools and their connection to variational optimal-control problems. A key technical contribution here is showing that the sample-path large deviations rate function, for algorithms that sample portions of system state, depends crucially on the sampling *frequencies* of these portions along with their individual rate functions. Conversely, we also show universal upper bounds on the queue overflow rate function of *any* scheduling policy that accesses partial, subset-based CSI. For this purpose, we introduce a novel martingale-based argument that, together with exponentially-twisted channel distributions, yields a universal upper bound on the buffer overflow probability exponent.
- (b) In the case where the collection of observable subsets available to the scheduler is disjoint, we prove that the lower bound on the large deviations buffer overflow rate function for Max-Exp matches the uniform upper bound on the rate function over all algorithms. This not only characterizes the exact buffer overflow exponent of the Max-Exp algorithm, but also shows rather surprisingly that *the simple Max-Exp strategy yields the optimal overflow exponent across all scheduling rules using partial CSI*. As a side consequence, this shows that for scheduling with singleton subsets of users, merely scheduling the user with the longest queue at each time slot – a greedy strategy when no CSI is available beforehand – is large-deviations rate function-optimal.

II. SYSTEM MODEL

This section describes the wireless system model we use along with its associated statistical assumptions. We consider a standard model of a wireless downlink system [4]: a time-slotted system of N users serviced by a single base station or server across N communication channels. In each time slot $k \in \{0, 1, 2, \dots\}$, the dynamics of the system are governed by three primary components:

- (a) **Arrivals:** An integer number of data packets $A_i(k)$ arrives to user i , $i = 1, \dots, N$. Packets get queued at their respective users if they are not immediately transmitted.
- (b) **Channel states:** The set of N channels assumes a random *channel state* $R(k)$, i.e. an N -tuple of integer *instantaneous service rates* (r_1, \dots, r_N) . At time slot k , let the instantaneous service rates be $(R_1(k), \dots, R_N(k))$.
- (c) **Scheduling:** One user $U(k) \in \{1, \dots, N\}$ is picked by a scheduling algorithm for service, and a number of packets not exceeding its instantaneous service rate is removed from its queue. Let $D_i(k)$ denote whether user i is scheduled in time slot k ($D_i(k) = 1$), or not ($D_i(k) = 0$). Then, user i 's queue length (denoted by $Q_i(\cdot)$) evolves as $Q_i(k+1) = [Q_i(k) + A_i(k) - D_i(k)R_i(k)]^+$, where $x^+ \equiv \max(x, 0)$.

We assume the following about the stochastics of the arrival and channel state processes:

Assumption 1 (Arrivals): Each user i 's arrival process $(A_i(k))_{k=0}^{\infty}$ is deterministic and equal to λ_i at all time slots. This is done merely for notational simplicity – any bounded, iid arrival process $(A_i(k))_{k=0}^{\infty}$ works, with the only modification being the large-deviations rate function of $A_i(k)$ added to all the rate function expressions in the paper.

Assumption 2 (Channel States): The joint channel states $R(k)$, $k = 0, 1, 2, \dots$ are independent and identically distributed across time, and take values from a finite set \mathcal{R} of integer N -tuples. Note that the channel states can have *any* joint distribution and can thus be correlated *across channels/users*.

Scheduling Model: Under scheduling with *partial channel state information*, at each time slot k , the scheduling algorithm makes *two* sequential choices to schedule a user:

- *Step 1:* Pick a *subset* $S(k)$ of the N channels, from a given collection \mathcal{O} of *observable subsets*. This choice can depend on all random variables in time slots up to and including k except the channel state $R(k)$.
- *Step 2:* Once the subset $S(k)$ of channels is chosen, the instantaneous service rates $(R_i(k))_{i \in S(k)}$ are revealed/available to the scheduling algorithm, and it chooses a user $U(k) \in S(k)$ for service, possibly depending on these service rates.

Note that at each time, the channel state information available to the scheduling algorithm is restricted to the chosen subset $S(k)$ of channels, as opposed to the full CSI case where all the service rates $(R_i(k))_{i=1}^N$ are available. For further detailed discussion about this scheduling model, we refer the reader to [8].

III. OBJECTIVE, ALGORITHMS AND MAIN RESULTS

Our focus is to design scheduling algorithms that reduce the likelihood of large queues in the system. Specifically, we seek to minimize the stationary probability (when it exists) that the longest queue in the system $\|Q(k)\|_{\infty} \triangleq \max_i Q_i(k)$ exceeds a threshold n . Alternatively, our goal is to maximize the *exponent* or *decay rate* of the exceedance probability

$$I \triangleq - \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} [\|Q(k)\|_{\infty} \geq n]$$

(when the limit exists), for scheduling algorithms that *observe only partial channel state* while scheduling. Note that for large n , $\mathbb{P}[\|Q(k)\|_{\infty} \geq n] \approx e^{-nI}$, so maximizing the exponent I gives smaller overflow probabilities. Also, it is well-known that packet delays are closely related to queue lengths [7], hence the reason for using I as our performance objective.

With this objective in mind, we introduce a new scheduling algorithm *Max-Exp* specified as follows:

Algorithm 1 Max-Exp

At each time slot k , breaking ties arbitrarily,

Step 1: Choose a subset $S(k)$, from the collection \mathcal{O} of observable subsets, such that

$$\sum_{i \in S(k)} \exp \left(\frac{Q_i(k)}{1 + \sqrt{Q(k)}} \right)$$

is maximized (here $\bar{Q}(k) \triangleq \frac{1}{N} \sum_{i=1}^N Q_i(k)$).

Step 2: Schedule a user $i \in S(k)$ such that $R_i(k) \exp \left(\frac{Q_i(k)}{1 + \sqrt{Q(k)}} \right)$ is maximized (the Exponential rule [12]).

By our probabilistic assumptions on the channel state process, Max-Exp makes the vector process of queue lengths at each time a discrete-time Markov chain. Following standard convention [4, 14, 15], we term the set of arrival rates $\lambda \equiv (\lambda_i)_{i=1}^N$ for which this Markov chain is positive-recurrent as the *throughput region* of Max-Exp. To not deviate from the main focus of this work, we state that when the observable subsets are disjoint, the throughput region of Max-Exp contains that of any other scheduling algorithm, i.e., Max-Exp is *throughput-optimal*. The proof of throughput-optimality is analogous to that of the Max-Sum-Queue scheduling algorithm [8]. Our main result states that Max-Exp yields the best (exponential) rate of decay of the tail of the longest queue over all strategies that use partial CSI from disjoint subsets:

Theorem 1 (Large Deviations Optimality of Max-Exp). *The following holds when the system's arrival rates λ lie in the interior of the throughput region of the Max-Exp scheduling algorithm:*

- (a) *Let \mathbb{P} denote the stationary probability distribution that the Max-Exp algorithm induces on the vector of queue lengths. Then, there exists $J_* > 0$ such that*

$$- \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} [\|Q(0)\|_{\infty} \geq n] \geq J_*.$$

- (b) *Let π be an arbitrary scheduling rule that induces a stationary distribution \mathbb{P}^{π} on the vector of queue lengths. If the system of observable subsets \mathcal{O} is disjoint, then*

$$- \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}^{\pi} [\|Q(0)\|_{\infty} \geq n] \leq J_*.$$

Thus, Max-Exp has the optimal large-deviations exponent (equal to J_) over all stabilizing scheduling policies with subset-based partial channel state information.*

Theorem 1 highlights the striking property that Max-Exp, using only current queue length information to sample channel subsets and the Exponential rule to schedule a sampled channel, yields the *fastest decay* of the buffer overflow probability across the whole spectrum of partial-CSI scheduling algorithms – including those that potentially use additional statistical information, traffic characteristics etc. The crucial scheduling step in Max-Exp is Step 1, which essentially samples the “right” channel subset depending on queue lengths. The result shows that queue length feedback is sufficient to guarantee good delay performance, provided suitable subsets of channel states are sampled as with the Max-Exp scheduling algorithm. We remark that the optimality of Max-Exp continues to hold even when all queue lengths are delayed by any bounded amount. We also remark that the restriction to disjoint observable subsets is necessary since otherwise, even throughput-optimality of Max-Exp-style scheduling rules does not hold [8]. This is briefly because the geometry of the throughput region is fundamentally different when subsets are disjoint, and its properties play a key role in the optimality proof for Max-Exp here.

En route to proving Theorem 1, we develop lower bounds for the large deviations exponents of partially and deterministically sampled iid processes, that are of independent interest. This results in a new rate function formulation in terms of variational optimization, that differs significantly from existing rate functions [5–7, 10, 11, 17] by explicitly incorporating partial channel state sampling behavior. Standard optimal control approaches for the full-CSI case cannot be applied to analyze partial-CSI scheduling algorithms – since only a portion of the channel state is revealed to the scheduler, the channel state process can cause large deviations by behaving atypically just in the revealed portion, and not jointly as a whole.

A related challenge arises in the process of finding universal upper bounds on the decay rate for arbitrary partial-CSI scheduling policies. Recent large-deviations work in full-CSI scheduling [5, 7] accomplishes this by calculating the “cost” of universal channel-state sample paths that cause buffer overflow under any scheduling algorithm; however, this procedure fails for algorithms actively sampling the channel state, since the cost of such sample paths intimately depends on the subset sampling behavior. To overcome this, we use a martingale-based argument in a novel way with the standard exponential tilting method to prove universal upper bounds on the exponent.

Observe that Max-Exp reduces to the following Max-Queue scheduling algorithm when the observable subsets are all the singleton users:

Algorithm 2 Max-Queue

At each time slot k , breaking ties arbitrarily,
Steps 1, 2: Schedule a user i such that $Q_i(k)$ is maximized.

Thus, an immediate corollary of Theorem 1 is the following optimality result for Max-Queue when the observable user sub-

sets are restricted to singletons, i.e., when there is effectively no CSI to use in scheduling:

Corollary 1 (Large Deviations Optimality of Max-Queue for singleton observable subsets). *If the system’s arrival rates λ lie in the interior of the throughput region of the Max-Queue scheduling algorithm, then Max-Queue has the optimal large-deviations exponent of the queue overflow probability over all stabilizing scheduling policies that can sample only individual channel states.*

Road map to Theorem 1: Though Theorem 1 for Max-Exp is our chief result, we prove it by first establishing the optimality result for Max-Queue (Corollary 1), and then extending the argument to the setting of general disjoint subsets. This is mainly because the essence of the optimality lies in the key subset selection step, and restricting attention to the case of singleton observable subsets allows us to concentrate on how subset selection influences the large deviations rate function of buffer overflow. Technically, another reason for this order of working is that Max-Queue can naturally be analyzed with the standard $O(n)$ fluid scaling, whereas showing the optimality property for Max-Exp requires using a more complex fluid limit framework at the $O(\sqrt{n})$ “local” fluid time-scale [5, 12].

IV. SAMPLE PATH LARGE DEVIATIONS FRAMEWORK

This section lays down preliminaries for the sample-path large deviations techniques we use to study overflow probabilities of wireless scheduling algorithms. Much of this framework is standard in recent large-deviations analyses of wireless systems [5–7], but we include it for completeness.

Fix an integer $T > 0$, and consider a sequence of (independent) queueing systems indexed by $n = 1, 2, \dots$, each with its own arrival and channel state processes, and evolving as described in Section II. Henceforth, we explicitly reference by the superscript (n) any quantity associated with the n th system. For any (possibly vector-valued) random process $X^{(n)}(k)$, $k = 0, 1, 2, \dots$ in the n th system, let us define its scaled (by $1/n$), shifted and piecewise linear version $x^{(n)}(\cdot)$ on the interval $[-T, 0]$ as follows: $x^{(n)}(t) \triangleq \frac{X^{(n)}(n(t+T))}{n}$ whenever $n(t+T)$ is an integer; $x^{(n)}(t)$ is linearly interpolated between these values for all other t .

For the n th queueing system, with k a nonnegative integer, let $F_i^{(n)}(k)$ be the total number of packets to queue i that arrived by time slot k , and $\hat{F}_i^{(n)}(k)$ be the number of packets that were served from queue i by time slot k . For a subset α of channels, let $C_\alpha^{(n)}(k)$ denote the total number of time slots before k when subset α was chosen by the scheduling algorithm. We also define its *sub-state* $R_\alpha^{(n)}(k)$ to be the vector of instantaneous service rates $R^{(n)}(k)$ restricted to the coordinates of α , i.e., $R_\alpha^{(n)}(k) = (R_i^{(n)}(k))_{i \in \alpha}$. Denote by $G_r^{\alpha, (n)}(k)$ the total number of time slots before time slot k when the subset α was picked and its sub-state was r ; and by $\hat{G}_{ri}^{\alpha, (n)}(k)$ the number of time slots before time k when subset α was picked, its observed sub-state was r and queue $i \in \alpha$ was ultimately scheduled for service. As

stated earlier, we denote by $Q_i^{(n)}(k)$ the length of queue i at time slot k , whose evolution is specified in Section II. Finally, we let $M^{(n)}(k)$ denote the vector-valued partial sums process corresponding to the sampled rates $R^{(n)}(k)\delta_{S(k)}$, i.e., $M^{(n)}(k) \triangleq \sum_{i=0}^k R^{(n)}(i)\delta_{S(i)}$.

Suppose a sequence of scaled processes $f_i^{(n)}(\cdot)$, $\hat{f}_i^{(n)}(\cdot)$, $c_\alpha^{(n)}(\cdot)$, $g_r^{\alpha,(n)}(\cdot)$, $\hat{g}_r^{\alpha,(n)}(\cdot)$, $q_i^{(n)}(\cdot)$ and $m^{(n)}(\cdot)$ converges uniformly (over $[-T, 0]$) to the corresponding “limit functions” $f_i(\cdot)$, $\hat{f}_i(\cdot)$, $c_\alpha(\cdot)$, $g_r^\alpha(\cdot)$, $\hat{g}_r^\alpha(\cdot)$, $q_i(\cdot)$ and $m(\cdot)$ on $[-T, 0]$. We call any such collection of joint limit functions, obtained via appropriately scaled pre-limit sample paths, a *Fluid Sample Path (FSP)* (we use the superscript T to emphasize the finite horizon $[-T, 0]$ if desired). We note that fluid sample paths inherit Lipschitz continuity (with the same Lipschitz constant) from their corresponding pre-limit processes indexed by n (when the pre-limits are Lipschitz-continuous), and are thus differentiable almost everywhere.

V. ANALYSIS: SINGLETON SUBSETS AND MAX-QUEUE

In this section, we treat the simpler setting where the disjoint observable subsets are all the singleton users in the system, i.e., $\mathcal{O} = \{\{i\} : 1 \leq i \leq N\}$. We use the subscript i to refer to subsets α . Thus, scheduling algorithms are essentially sampling algorithms – Step 2 of the algorithm is to schedule the lone user whose channel state is observed. In what follows, we describe the three key steps involved in showing that Max-Queue yields the optimal decay rate of buffer overflow probability.

A. Lower-bounding Max-Queue’s Decay Rate: Large Deviations for Sampled Processes

Consider the queueing system operating under an arbitrary *nonrandom* scheduling algorithm, i.e., the algorithm’s choice of a singleton user in the current time slot is a deterministic function of the entire history of observed users’ indices and channel states, and does not depend on the unobserved channel states in the past. Max-Queue with deterministic tie-breaking (e.g., pick the lowest-indexed queue when there are two or more longest queues) is an example of a nonrandom scheduling algorithm, since the current user chosen depends on accumulated queue lengths, which in turn depend directly on the channel rates obtained as a result of past scheduling choices.

The sequence of observed users and their channel states under a nonrandom scheduling algorithm is an outcome of sampling an iid vector-valued process (i.e., the full channel state) in a nonrandom and *predictable* (i.e., with sampling indices depending only on past observed history) manner. Our first key result (Proposition 2) essentially furnishes a lower bound for the deviation probability of the queue-length process (equivalently the cumulative process of observed channel states) in time slots $0, \dots, nT$, in terms of a novel sample-path large deviations rate function of the user selection and channel state paths.

Let us fix $T > 0$. For $q_0 \in \mathbb{R}^N$, let $\mathbb{P}_0^{n,T}$ be the probability measure of the n -th queueing system conditioned on starting

the system at $Q^{(n)}(0) = nq_0$ (i.e. $q^{(n)}(-T) = q_0$). If we denote by $\mathcal{C}_{\mathcal{L}}^+([-T, 0])$ the space of nonnegative \mathbb{R}^N -valued Lipschitz functions on $[-T, 0]$ equipped with the supremum norm, then we have:

Proposition 2 (Large Deviations for Sampled Processes). *Let Γ be a closed set of trajectories in $\mathcal{C}_{\mathcal{L}}^+([-T, 0])$. Then, under any nonrandom scheduling policy,*

$$\begin{aligned} & -\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_0^{n,T} \left[q^{(n)} \in \Gamma \right] \\ & \geq \inf_{(m^T, c^T, q^T)} \int_{-T}^0 \left[\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)} \right) \right] dt \quad (1) \\ & \text{subject to } (m^T, c^T, q^T) \text{ an FSP,} \\ & \quad q^T(-T) = 0, q^T \in \Gamma, \end{aligned}$$

with $\Lambda_i^*(\cdot)$ being the Legendre-Fenchel dual of $\Lambda_i(\lambda) = \log \mathbb{E}[e^{\lambda R_i(0)}]$, i.e., the Cramér rate function for the empirical mean of the marginal rate $(R_i(k))_k$.

Proposition 2 states that the “correct” sample-path large deviations rate function, for algorithms that can sample only singleton subsets of channels, is a combination of the standard rate functions Λ_i^* for the empirical means of individual channel rates *weighted by the corresponding channel selection frequencies* \dot{c}_i . Note the crucial dependence of the rate function on the subset selection process, captured by weighting Λ_i^* by \dot{c}_i in (1) – a significant departure from the rate function studied for the standard case of full channel state information where there is no pre-weighting by the algorithm-dependent factor \dot{c} [5, 7].

The proof of this result relies on the key fact that the sample-path trajectory of any nonrandom scheduling/sampling algorithm is completely determined by only the *sampled* user’s index and the observed channel state at all times, instead of the entire joint channel state process with unobserved channel states. Also, since only one component of the joint channel state is used at each instant, there is no loss of generality in assuming that all the channel state processes are independent with the original marginals. These two properties, together with exchangeability of the channel state process, allow us to derive a large deviations rate function for the random process of sampled channel states, which is further transformed to the rate function (1) as a function of empirical channel means and sampling frequencies. The proof is deferred to [25] due to space constraints.

Applying Proposition 2 with $\Gamma = \{q \in \mathcal{C}_{\mathcal{L}}^+([-T, 0]) : \|q(0)\|_\infty \geq 1\}$ gives a finite-horizon lower bound for the rate function of longest-queue overflow. For any FSP (m^T, c^T, q^T) feasible in the RHS of (1), we have

$$\int_{-T}^0 \left[\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)} \right) \right] dt \geq \inf_{t \in \mathcal{B}} \frac{\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)} \right)}{\frac{d}{dt} \|q(t)\|_\infty},$$

with \mathcal{B} denoting the (almost all) points in $[-T, 0]$ at which all

the relevant derivatives exist. Let us define

$$J_{**} \triangleq \inf_{\substack{T, (m^T, c^T, q^T) \\ 0 \leq t \leq T}} \frac{\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{m_i(t)}{c_i(t)} \right)}{\frac{d}{dt} \|q(t)\|_\infty},$$

with the infimum over all feasible FSPs, all regular points t , and *all finite horizons* T . Intuitively, we expect that the finite horizon probability distribution $\mathbb{P}_{q_0}^{n, T}$ tends to the stationary distribution \mathbb{P} ; this is borne out by the fact that minimizing the RHS of (1) across FSPs (m^T, c^T, q^T) over all finite horizons $T > 0$ yields the lower bound J_{**} on the *stationary* overflow probability exponent (see [25] for details):

Proposition 3 (Lower bound for Max-Queue’s Decay Rate).

$$-\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left[\|q^{(n)}(0)\|_\infty \geq 1 \right] \geq J_{**}. \quad (2)$$

Proposition 3 is thus a “cost per unit max-queue drift” lower bound on the decay rate of the queue overflow probability under Max-Queue.

B. Universal Large Deviations Upper Bound

We next derive a *uniform* upper bound for the stationary buffer overflow probability decay rate, over all singleton-CSI scheduling algorithms. A popular approach followed in recent work [5–7] to do this is by estimating the cost of “straight-line” joint channel state sample paths that universally cause buffer overflow. However, when only a *dynamically selected portion of the channel state* is visible to the scheduling algorithm, the cost (1) of such straight-line paths depends explicitly on the algorithm’s sampling behavior, so the standard approach fails.

For every i , let $\phi'_i \geq 0$ denote a “twisted” mean rate for channel i , and consider the quantity $\frac{\sum_i c'_i \Lambda_i^*(\phi'_i)}{[\max_i (\lambda_i - c'_i \phi'_i)]^+}$. Here, we assume that $\sum_i c'_i = 1$, and that the fraction is ∞ whenever the denominator is 0. Suppose a scheduling policy samples each channel i with frequency c'_i . Then, (a) the numerator of the above expression corresponds to the “instantaneous large deviations cost” of witnessing each channel i ’s mean rate be ϕ'_i (by (1)), while (b) the denominator can be interpreted as the average rate with which the longest queue grows when each channel i is sampled with a frequency c'_i . Maximizing the expression over all possible user sampling/scheduling frequencies $\{c'_i : \sum_i c'_i = 1, c'_i \geq 0\}$ induced by scheduling algorithms should thus give the highest possible large deviations cost for buffer overflow. This intuition is formalized in the following key result:

Proposition 4 (Universal Upper Bound on Decay Rate for any Algorithm). *Let π be a stabilizing scheduling policy for the arrival rate $\lambda = (\lambda_1, \dots, \lambda_N)$, and let \mathbb{P}^π be its associated stationary measure. For any $\phi'_i \in \mathbb{R}^+$, $i = 1, \dots, N$,*

$$\begin{aligned} & -\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}^\pi \left[\|q^{(n)}(0)\|_\infty \geq 1 \right] \\ & \leq \sup_{\substack{\sum_i c'_i = 1 \\ c'_i \geq 0}} \frac{\sum_i c'_i \Lambda_i^*(\phi'_i)}{[\max_i (\lambda_i - c'_i \phi'_i)]^+}. \end{aligned} \quad (3)$$

Each choice of the twisted means $(\phi'_i)_i$ above yields such an upper bound on the decay rate. Thus, the best possible upper bound is obtained by minimizing (3) over all choices $(\phi'_i)_i$.

According to Proposition 4, an upper bound on the buffer overflow rate function when scheduling with partial channel observability is the largest “weighted-cost per unit increase of the maximum queue,” over all possible frequencies of sampling subsets of channels. We emphasize that the maximization over the sampling frequencies c'_i , in (3), is a distinct feature that emerges while considering partial information algorithms, as opposed to the case where scheduling is performed with full joint CSI.

At the heart of the proof of Proposition 4 is a twisted measure construction where each channel’s marginal rate is ϕ'_i . Observing that the cumulative fluid service process $m(\cdot)$ is a submartingale under the twisted measure for any scheduling algorithm, the Doob-Meyer decomposition [26] allows us to express $m(\cdot)$ as the predictable algorithm-dependent component $\phi'_i c'_i(\cdot)$ plus a martingale noise component $\bar{m}(\cdot)$. This shows that with high probability, the service provided to each queue i is approximated by $\phi'_i c'_i(\cdot)$, i.e., we can effectively treat each channel i as having a deterministic fluid service rate of ϕ'_i . Analyzing this deterministic fluid system for overflow and translating the results back to the original probabilistic system gives us the result. Due to space limitations, we refer the reader to [25] for the full proof details.

C. Large Deviations Optimality of the Max-Queue Policy: Connecting the Upper and Lower Bounds

The final step in the proof of optimality of Max-Queue (Corollary 1) is carried out by showing that the lower bound for Max-Queue (2) in fact dominates the uniform upper bound (3) over all scheduling policies:

Proposition 5 (Matching Large Deviations Bounds, Max-Queue, Singleton Subsets). *There exist nonnegative $\hat{\phi}'_1, \dots, \hat{\phi}'_N$, with $\lambda \notin \mathcal{C}(\hat{\phi}'_1, \dots, \hat{\phi}'_N)$, such that*

$$\sup_{\substack{\sum_i c'_i = 1 \\ c'_i \geq 0}} \frac{\sum_i c'_i \Lambda_i^*(\hat{\phi}'_i)}{[\max_i (\lambda_i - c'_i \hat{\phi}'_i)]^+} \leq J_{**}.$$

The proof of this result involves solving the non-convex problem for the rate function lower bound given in Proposition 3, and relating the solution to a suitable uniform upper bound of the type prescribed by Proposition 4. It utilizes the convexity and lower-semicontinuity of the rate functions Λ_i^* , and is accomplished by considering the properties of the $(\phi'_i)_i$ which minimize the upper bound (3). Due to space constraints, we refer the reader to [25] for details.

VI. ANALYSIS: GENERAL SUBSETS AND MAX-EXP

Having shown that Max-Queue yields an optimal queue overflow exponent for scheduling using only single channel states, in this section we extend the result to the general setting of arbitrary disjoint subsets of observable channels and show that Max-Exp is optimal for the overflow exponent. To do

this, we follow the same key steps in obtaining the Max-Queue result – (a) prove lower bounds on the buffer overflow exponent for Max-Exp, (b) derive universal upper bounds on the buffer overflow exponent across all scheduling algorithms using subset channel state information, and (c) demonstrate that the upper and lower bounds match.

However, the approach to show optimality of the Max-Exp algorithm warrants more sophisticated analysis as compared to the case of Max-Queue. This is primarily due to the fact that the Max-Exp algorithm is not *scaling-invariant*, i.e., scaling all queue-lengths by a uniform constant changes the scheduling behavior. Intrinsically, Max-Exp operates at the $O(\sqrt{n})$ time-scale, i.e., when all the queue lengths are $O(n)$, a $O(\sqrt{n})$ change in them causes a shift in Max-Exp’s scheduling behavior. In other words, examining Max-Exp’s scheduling over $O(n)$ time slot intervals effectively “washes out” information about its actions, resulting in crude bounds. This sets Max-Exp apart from Max-Queue which is naturally coupled to the timescale of $O(n)$ time slots, and prevents us from using the standard $O(n)$ fluid scaling to analyze the fluid sample path behavior of Max-Exp.

Hence, our analysis for Max-Exp proceeds by looking at sample paths of the system’s processes over intervals of $O(\sqrt{n})$ time slots. For Step (a) above, analogous to Proposition 2, we establish a “refined” Mogulskii-type theorem for sample-path large deviations of predictably sampled processes over a sub- $O(n)$ timescale (a corresponding result for the full-CSI case was first proved in [5]). Next, we use the framework of *Local Fluid Sample Paths* (LFSPs, introduced in [12]) to obtain a lower bound on the decay exponent of Max-Exp’s overflow probability. LFSPs allow us to “magnify” the standard $O(n)$ fluid limit processes to examine events on the $O(\sqrt{n})$ “local fluid” timescale, and this helps us to match the lower and upper bounds for the decay exponent to establish the optimality of Max-Exp.

A. Lower Bounding Max-Exp’s Decay Rate: Refined-timescale Large Deviations for Sampled Processes and Local FSPs

Here, we extend the sampling-based large-deviations bound from Proposition 2 to hold over a finer-than- $O(n)$ timescale. The basic idea here is to lower-bound the large deviations cost from (1) by linearizing sample paths over the finer timescale. This expresses the intuitive notion that over the finer timescale, typical large deviations of random processes occur “locally along straight lines”.

The general approach for studying scheduling behavior on finer-than- $O(n)$ timescales [5] is to introduce a positive integer function $u(n)$, such that $u(n) \rightarrow \infty$ and $u(n)/n \rightarrow 0$ as $n \rightarrow \infty$. In our analysis, we will take $u(n) = \lceil \sqrt{n} \rceil$, the relevant timescale for the Max-Exp scheduling rule. For any non-decreasing, right-continuous-with-left-limits (RCLL) vector-valued function h on $[0, \infty)$, let $U^n h$ denote the continuous, piecewise-linearized version of h obtained from h as follows: we divide $[0, \infty)$ into contiguous subintervals of size $u(n)/n$ each, and linearize h within each subinterval. For $t \geq 0$, let $\theta^{(n)}(t)$ be the largest multiple of $u(n)/n$ not exceeding t .

Finally, for each observable subset α , let Λ_α^* be the Sanov rate function [27] for the empirical marginal distribution of the state of its channels $(R_i(1))_{i \in \alpha}$. Thus, the domain of Λ_α^* is the $|\mathcal{R}_\alpha|$ -dimensional simplex where \mathcal{R}_α is the set of all possible sub-states for subset α .

In order to track large-deviations costs over the $u(n)$ timescale, let us expand the definition of a Fluid Sample Path (FSP) to include the following *additional* functions:

- (a) A pre-limit *refined cost* function $\bar{J}_t^{(n)}$ defined over $[-T, 0]$ as: $\bar{J}_t^{(n)} \triangleq \int_{-T}^{\theta^{(n)}(t)} \left[\sum_{\alpha} (U^n c_\alpha^{(n)})'(u) \cdot \Lambda_\alpha^* \left(\frac{(U^n g^{\alpha, (n)})'(u)}{(U^n c_\alpha^{(n)})'(u)} \right) \right] du$, where $g^{\alpha, (n)} \equiv \left(g_r^{\alpha, (n)} \right)_{r \in \mathcal{R}_\alpha}$.
- (b) The uniform convergence $\bar{J}^{(n)} \rightarrow \bar{J}$ as $n \rightarrow \infty$, over $[-T, 0]$, to a non-negative, non-decreasing, Lipschitz-continuous *limit refined cost* function \bar{J} on $[-T, 0]$.

To simplify notation, for general Lipschitz-continuous functions $(c_\alpha)_\alpha$ and $(g_r^\alpha)_{\alpha r}$ of the appropriate vector dimension on $[-T, 0]$, we denote

$$J_t \equiv J_{(c, g)}(t) \triangleq \int_{-T}^t \left[\sum_{\alpha} \dot{c}_\alpha(u) \Lambda_\alpha^* \left(\frac{\dot{g}^\alpha(u)}{\dot{c}_\alpha(u)} \right) \right] du. \quad (4)$$

The following key finite-horizon result strengthens Proposition 2. It states that for any nonrandom scheduling algorithm, the sample path large deviations rate function for the queue length process is lower-bounded by the minimum *refined cost* over valid fluid sample paths.

Proposition 6 (Refined-time-scale Large Deviations for Sampled Processes). *Let Γ be a closed set of trajectories in $C_{\mathcal{L}}^+([-T, 0])$. Then, under a nonrandom scheduling policy,*

$$-\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_0^{n, T} \left[q^{(n)} \in \Gamma \right] \geq \inf \{ \bar{J}_0 : (q, \bar{J}) \text{ FSP on } [-T, 0], q \in \Gamma \}. \quad (5)$$

The proof of this proposition combines ideas from the large deviations of sampled processes (Proposition 2) and the refined Mogulskii theorem shown by Stolyar [5] to establish the above bound. Due to space constraints, we refer the reader to [25] for the proof.

Similar to extending the result of Proposition 2 to the stationary queue length distribution, minimizing the RHS of (5) across FSPs over all finite time horizons $T > 0$ yields a lower bound for the large deviations rate of the *stationary* queue overflow probability. This uses standard tools (see, for instance, [5, 7]), and we omit the proof for brevity.

Local Fluid Sample Paths: Analyzing the variational problem on the right-hand side of (6) (minimized across all finite horizons T) demands a close look at the derivatives of FSPs under the Max-Exp scheduling algorithm. At the same time, since the Max-Exp rule naturally operates at the $O(\sqrt{n})$ time-scale, derivative information is typically “washed out” of

the standard $O(n)$ -scaled fluid sample paths. This motivates us to define and use Local Fluid Sample Paths (LFSPs) [5, 12] with a $O(\sqrt{n})$ -type scaling, under which information about scheduling choices and drifts can be clearly understood with regard to the Max-Exp scheduling rule. Our goal is to show, by “magnifying” such a feasible FSP about some $\tau \in [-T, 0]$ and taking “local” fluid limits at τ , that the unit large-deviations cost of raising the maximum queue in the associated LFSP is roughly \bar{J}_0 . Thus, a further lower bound on the Max-Exp decay rate is the least large-deviations cost per unit increase of the maximum queue over all feasible LFSPs.

The actual LFSP construction exactly parallels the one followed by Stolyar [5], and proceeds by taking the fluid-scaled functions and magnifying space and time by a factor of \sqrt{n} . For instance, for the function $f_i^{(n)}$ with fixed $\tau \in [-T, 0]$ and $S > 0$, we let $\diamond f_i^{(n)}(s) \triangleq \sqrt{n} \left[f_i^{(n)}(\tau + s\sqrt{n}) - f_i^{(n)}(\tau) \right]$ for $s \in [0, S]$, and take a uniform limit of such rescaled, centered functions, which is the LFSP $\diamond f_i(\cdot)$ over $[0, S]$. We omit the details of the LFSP construction and refer the reader to [25] instead.

With this framework of LFSPs set up, consider a feasible FSP (q, \bar{J}) on $[-T, 0]$ for the right-hand side of (5), for which $q(-T) = 0$, $q(0) \geq 1$ and whose refined cost is \bar{J}_0 . Fix also an arbitrary $\epsilon > 0$. Then, there must exist a time point $\tau \in (-T, 0)$ such that $q_*(\tau) > 0$, $q'_*(\tau) > 0$, $\bar{J}'(\tau) > 0$, and $\frac{\bar{J}'(\tau)}{q'_*(\tau)} < \bar{J}_0 + \epsilon$. Continuing further using techniques similar to those in [5], we can show that given an arbitrary $S > 0$ (and $\epsilon > 0$), we can construct/find an LFSP at τ , together with a constant $\epsilon_1 > 0$ such that

$$\diamond q_*(S) - \diamond q_*(0) \geq \epsilon_1 S, \quad \text{and} \quad (6)$$

$$\frac{J_{(\diamond c, \diamond g)}(S) - J_{(\diamond c, \diamond g)}(0)}{\diamond q_*(S) - \diamond q_*(0)} \leq \bar{J}_t + \epsilon, \quad (7)$$

i.e., we can approximate the cost of the original FSP arbitrarily well with the “unit cost of raising $\diamond q_*$ ” of a suitably constructed LFSP.

To complete the lower-bound on the queue overflow exponent of the Max-Exp rule, consider for a general LFSP the following **potential function** of its queue state: $\Psi(\diamond q) \triangleq \max_{\alpha \in \mathcal{O}} \Psi_{\alpha}(\diamond q) \equiv \max_{\alpha \in \mathcal{O}} \sum_{i \in \alpha} e^{\diamond q_i + b_i}$, together with its logarithm $\Phi(\diamond q) \triangleq \log \Psi(\diamond q) = \max_{\alpha} \log \Psi_{\alpha}(\diamond q)$.

We state as a fact that the function $\Phi(\diamond q)$ uniformly approximates $\diamond q_* \equiv \|\diamond q\|_{\infty}$, in the sense that $|\Phi(\diamond q) - \diamond q_*| \leq \Delta$ for some fixed $\Delta > 0$. Now, consider the feasible FSP (q, \bar{J}) on $[-T, 0]$ for the right-hand side of (5) as before. Using the above approximation fact and (6), (7), an LFSP can be constructed on some interval $[0, S]$ such that for a suitable $\epsilon_1 > 0$,

$$\Phi(\diamond q(S)) - \Phi(\diamond q(0)) \geq (\epsilon_1/2)S, \quad (8)$$

$$\frac{J_{(\diamond c, \diamond g)}(S) - J_{(\diamond c, \diamond g)}(0)}{\Phi(\diamond q(S)) - \Phi(\diamond q(0))} \leq \bar{J}(t) + 2\epsilon. \quad (9)$$

We now turn to the LHS of (9) – modulo the arbitrarily small $\epsilon > 0$, we have seen that it is a lower bound on the original

FSP cost \bar{J}_0 . We can write

$$\begin{aligned} \frac{J_{(\diamond c, \diamond g)}(S) - J_{(\diamond c, \diamond g)}(0)}{\Phi(\diamond q(S)) - \Phi(\diamond q(0))} &= \frac{\int_0^S \frac{d}{ds} J_{(\diamond c, \diamond g)}(s) ds}{\int_0^S \frac{d}{ds} \Phi(\diamond q(s)) ds} \\ &\geq \inf_{s \in [0, S]} \frac{\frac{d}{ds} J_{(\diamond c, \diamond g)}(s)}{\frac{d}{ds} \Phi(\diamond q(s))} \stackrel{(a)}{\geq} \inf_{s \in [0, S]} \frac{\sum_{\alpha} \diamond \dot{c}_{\alpha}(s) \Lambda_{\alpha}^* \left(\frac{\diamond \dot{g}^{\alpha}(s)}{\diamond \dot{c}_{\alpha}(s)} \right)}{\frac{d}{ds} \Phi(\diamond q(s))} \\ &= \inf_{s \in [0, S]} \frac{\sum_{\alpha} \diamond \dot{c}_{\alpha}(s) \Lambda_{\alpha}^* (\diamond \dot{g}^{\alpha}(s) / \diamond \dot{c}_{\alpha}(s))}{\frac{d}{ds} \Phi(\diamond q(s))}, \end{aligned} \quad (10)$$

where (a) follows from the definition (4). As a consequence of the above inequality, we can finally give the following large-deviations lower bound:

Proposition 7 (Lower Bound for Max-Exp’s Decay Rate). *If \mathbb{P} denotes the stationary measure induced by the Max-Exp policy, then*

$$\begin{aligned} -\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left[\|q^{(n)}(0)\|_{\infty} \geq 1 \right] \\ \geq \inf_{s \in [0, S]} \frac{\sum_{\alpha} \diamond \dot{c}_{\alpha}(s) \Lambda_{\alpha}^* (\phi'_{\alpha}(s))}{\frac{d}{ds} \Phi(\diamond q(s))} \end{aligned} \quad (11)$$

for any valid Local Fluid Sample Path (LFSP).

Letting J_* denote the infimum on the RHS of (11) over all valid LFSPs, a further lower bound on the buffer overflow exponent of Max-Exp is thus J_* .

B. Universal Large Deviations Upper Bound

The next key step in our program to show Max-Exp’s tail optimality is to exhibit a *uniform* upper bound on the decay rate of the stationary queue-overflow probability across all stabilizing scheduling algorithms. Following a similar route as for Proposition 4, we accomplish this by using twisted marginal probability distributions for the subset channel states, and the local/subset-based throughput regions that they induce.

Recall that for an observable subset α , \mathcal{R}_{α} denotes the (finite) set of all possible (joint) sub-states that can be observed channels in α . We use Π_{α} to denote the $|\mathcal{R}_{\alpha}|$ -valued simplex, i.e., the set of all probability measures on the sub-states of α . Any distribution $\phi'_{\alpha} \in \Pi_{\alpha}$ induces a *subset throughput region* $V_{\phi'_{\alpha}}$, which represents all the long-term average service rates that can be sustained to users in α when the sub-states are distributed as ϕ'_{α} (see also [4, 8]). The uniform large-deviations upper-bound can now be stated for any stabilizing scheduling policy π :

Proposition 8 (Universal Upper Bound on Decay Rate for any Algorithm). *Let π be a stabilizing scheduling policy for arrival rates $\lambda = (\lambda_1, \dots, \lambda_n)$, and let \mathbb{P}^{π} be the associated stationary measure. Let distributions $\phi'_{\alpha} \in \Pi_{\alpha}$ be fixed for every α such that $\lambda \notin \mathcal{CH}((V_{\phi'_{\alpha}})_{\alpha})$. Then,*

$$\begin{aligned} -\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}^{\pi} \left[\|q^{(n)}(0)\|_{\infty} \geq 1 \right] \\ \leq \sup_{\substack{\sum_{\alpha} c'_{\alpha} = 1 \\ c'_{\alpha} \geq 0}} \left[\frac{\sum_{\alpha} c'_{\alpha} \Lambda_{\alpha}^* (\phi'_{\alpha})}{\max_{\alpha, v_{\alpha} \in V_{\phi'_{\alpha}}} \max_{i \in \alpha} (\lambda_i - c'_{\alpha} v_{\alpha, i})} \right]. \end{aligned} \quad (12)$$

The proof of this result uses ideas similar to that of Proposition 4, but is more sophisticated due to dealing with subsets of channels and their associated multidimensional rate regions. It employs a key concentration result for vector-valued martingales and exploits the convexity of the subset rate regions. Due to space constraints, we skip the details and refer the reader to [25].

C. Large Deviations Optimality of the Max-Exp Policy: Connecting the Upper and Lower Bounds

The crucial final step in proving tail optimality for Max-Exp (Theorem 1) is to show that the lower bound J_* on its decay exponent in fact matches the uniform upper bound (12) on the decay exponent of any stabilizing scheduling policy, along the lines of Proposition 4 for Max-Queue.

Proposition 9 (Matching Large Deviations Bounds, Max-Exp, Arbitrary Disjoint Subsets). *Let π be any stabilizing scheduling policy for arrival rates $\lambda = (\lambda_1, \dots, \lambda_n)$, and let \mathbb{P}^π be the associated stationary measure. Then,*

$$-\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}^\pi \left[\|q^{(n)}(0)\|_\infty \geq 1 \right] \leq J_*,$$

i.e., Max-Exp has the optimal large-deviations exponent (equal to J_) over all stabilizing scheduling policies with subset-based partial channel state information.*

To show this result for an arbitrary collection of disjoint observable subsets, we use the key idea that each subset α can be viewed as a fictitious “queue” holding $\Psi_\alpha(s)$ packets at each time s . This brings the situation in correspondence with the singleton queues case of Max-Queue, and together with convexity and lower-semicontinuity considerations and the steps followed to prove Proposition 8, the above Proposition can be established. Due to space limitations, we defer the full proof to [25].

VII. ACKNOWLEDGMENTS

This work was partially supported by the NSF grants CNS-0325788, CNS-0347400, CNS-0519535, CNS-0721532, EFRI-0735905 and by the DARPA ITMANET program.

REFERENCES

- [1] S. Parkvall and D. Astely, “The evolution of LTE towards IMT-advanced,” *Journal of Communications*, vol. 4, no. 3, 2009.
- [2] J. G. Andrews, A. Ghosh, and R. Muhamed, *Fundamentals of WiMAX: Understanding Broadband Wireless Networking*. Prentice Hall, 2007.
- [3] L. Tassiulas and A. Ephremides, “Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks,” *IEEE Transactions on Automatic Control*, vol. 4, pp. 1936–1948, December 1992.
- [4] M. Andrews, K. Kumaran, K. Ramanan, A. L. Stolyar, R. Vijayakumar, and P. Whiting, “Scheduling in a queueing system with asynchronously varying service rates,” *Probability in Engineering and Information Sciences*, vol. 14, pp. 191–217, 2004.
- [5] A. Stolyar, “Large deviations of queues sharing a randomly time-varying server,” *Queueing Systems*, pp. 1–35, 2008.
- [6] B. Sadiq and G. de Veciana, “Optimality and large deviations for multi-user wireless systems under pseudo-log opportunistic scheduling,” in *Proc. 46th Annual Allerton Conference on Communication Control and Computing*, pp. 1–8, 2008.
- [7] V. J. Venkataramanan and X. Lin, “On the queue-overflow probability of wireless systems: A new approach combining large deviations with Lyapunov functions,” *submitted to IEEE Trans. Information Theory*, 2009. <http://cobweb.ecn.purdue.edu/~linx/papMr/it09-ldp-updated2.pdf>.
- [8] A. Gopalan, C. Caramanis, and S. Shakkottai, “On wireless scheduling with partial channel state information,” *IEEE Trans. on Information Theory*, Jan. 2012. Previous version in the Proc. of the Allerton Conference on Communication, Control and Computing, 2007.
- [9] D. Bertsimas, I. Paschalidis, and J. Tsitsiklis, “Asymptotic buffer overflow probabilities in multiclass multiplexers: an optimal control approach,” *Automatic Control, IEEE Transactions on*, vol. 43, pp. 315–335, Mar 1998.
- [10] S. Shakkottai, “Effective capacity and QoS for wireless scheduling,” *IEEE Transactions on Automatic Control*, February 2008.
- [11] L. Ying, G. Dullerud, and R. Srikant, “A large deviations analysis of scheduling in wireless networks,” *IEEE Transactions on Information Theory*, vol. 52, November 2006.
- [12] S. Shakkottai and A. Stolyar, “Scheduling for multiple flows sharing a time-varying channel: The exponential rule,” *American Mathematical Society Translations, Series 2, A volume in memory of F. Karpelevich, Yu. M. Suhov, Editor*, vol. 207, 2002.
- [13] S. Shakkottai, R. Srikant, and A. L. Stolyar, “Pathwise optimality of the exponential scheduling rule for wireless channels,” *Advances in Applied Probability*, vol. 36, no. 4, pp. 1021–1045, 2004.
- [14] M. Neely, E. Modiano, and C. Rohrs, “Dynamic power allocation and routing for time-varying wireless networks,” *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 1, pp. 89–103, 2005.
- [15] A. Eryilmaz, R. Srikant, and J. R. Perkins, “Stable scheduling policies for fading wireless channels,” *IEEE/ACM Transactions on Networking*, vol. 13, pp. 411–424, 2005.
- [16] A. Stolyar and K. Ramanan, “Largest weighted delay first scheduling: Large deviations and optimality,” *Annals of Applied Probability*, pp. 1–48, 2001.
- [17] V. J. Venkataramanan and X. Lin, “Structural properties of LDP for queue-length based wireless scheduling algorithms,” in *Proc. 45th Annual Allerton Conference on Communication Control and Computing*, 2007.
- [18] A. Stolyar, “Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic,” *Annals of Applied Probability*, vol. 14, No.1, pp. 1–53, 2004.
- [19] K. Kar, X. Luo, and S. Sarkar, “Throughput-optimal scheduling in multi-channel access point networks under infrequent channel measurements,” *IEEE Transactions on Wireless Communications*, vol. 7, no. 7, pp. 2619–2629, 2008.
- [20] S. Patil and G. de Veciana, “Reducing feedback for opportunistic scheduling in wireless systems,” *IEEE Trans. on Wireless Comm.*, vol. 12, pp. 4227–38, December 2007.
- [21] T. Tang and R. W. Heath, “Opportunistic feedback for downlink multi-user diversity,” *IEEE Communication Letters*, vol. 9, pp. 948–950, 2005.
- [22] P. Chaporkar, A. Proutiere, H. Asnani, and A. Karandikar, “Scheduling with limited information in wireless systems,” in *Proceedings of the Tenth ACM International Symposium on Mobile Ad-hoc Networking and Computing*, pp. 75–84, 2009.
- [23] N. Chang and M. Liu, “Optimal channel probing and transmission scheduling for opportunistic spectrum access,” in *Proc. ACM International Conference on Mobile Computing and Networking (MobiCom)*, (Montreal, Canada), September 2007.
- [24] L. Ying and S. Shakkottai, “On throughput optimality with delayed network-state information,” in *Proc. of the Information Theory and Applications Workshop*, pp. 339–344, January 2008.
- [25] A. Gopalan, C. Caramanis, and S. Shakkottai, “Wireless scheduling with partial channel-state information: Large deviations and optimality,” tech. rep., The University of Texas at Austin, July 2011.
- [26] R. Durrett, *Probability: Theory and Examples*. Belmont, CA: Brooks/Cole - Thomson Learning, 2005.
- [27] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*. Jones and Bartlett, 1993.