Wireless Scheduling with Partial Information: Large Deviations and Optimality

Aditya Gopalan, Constantine Caramanis, Sanjay Shakkottai

July 29, 2011

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.

1 Introduction

Next-generation wireless cellular networks, based on standards such as 3GPP-LTE [1], promise high-speed packet-switched data services for a variety of applications, including file transfer, peer-to-peer sharing and realtime audio/video streaming. With a large number of mobile users in a cellular network, the time-varying nature of the wireless medium along with interference effects makes wireless resources scarce. This necessitates effective data scheduling on the downlink of such a network to maximize data rates to users. In addition to data rates or throughput, to support highly delay sensitive applications like video streaming, the scheduling algorithm at the cellular base station must ensure that packet delays in the system are kept low. There has been much recent work to develop scheduling algorithms with optimal throughput and good delay performance. Such algorithms operate opportunistically and utilize instantaneous wireless Channel State Information

(CSI) from the entire system to make good scheduling decisions. However, in a practical situation with many users in the network, channel state feedback resources are typically limited, i.e., it may be infeasible to acquire complete instantaneous CSI from all channels, and instead request CSI feedback from only a subset of users at each time. Thus, it is important to design algorithms that can schedule with partial CSI as opposed to full CSI, and that perform well in the sense of delay.

Using partial CSI from subsets of channels, however, requires a significantly different manner of opportunism in wireless scheduling. The scheduling algorithm now needs to make a careful choice of which instantaneous channel states to sample, before using the CSI to schedule data to users. In such a setting, though natural extensions of complete-CSI scheduling algorithms exist with throughput-optimal properties [2], it is not clear what their delay performance is. The structure of delay-minimizing algorithms is also unknown, i.e., how a algorithm for delay should choose subsets of channels to acquire instantaneous CSI, whether any additional statistical information is needed for picking subsets, how users should be scheduled in the observed subset etc.

In this work, we model a time-slotted wireless downlink, in which a base station schedules data to a user chosen from a fixed population of users using partial CSI from subsets of users. 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 the longest queue overflowing past a threshold. We present a new scheduling algorithm Max-Exp that obtains partial CSI relying on just current accumulated queue lengths and no other auxiliary information. Employing sample-path large deviations techniques, we show that 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 that 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 viewpoint, though sophisticated large deviations techniques exist to analyze algorithms for wireless scheduling, studying the queue overflow performance of scheduling strategies that do not know the complete state of the system presents significantly new analytical challenges. The fact that the partial CSI acquired from users depends primarily on the subset dynamically sampled by the scheduling algorithm gives rise to a large deviations behaviour (or rate function) significantly different than the ones known for traditional settings with complete CSI. Also, standard approaches of analyzing large-deviations properties using continuity of queue-length/delays with respect to arrivals and channel processes are difficult to apply here due to the complex two-step structure of scheduling algorithms using 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 relevant rate functions by connecting them to suitable variational problems. We believe these techniques are of independent interest as tools to analyze the behaviour of various types of scheduling policies that use information about the state of a system by actively deciding to sample a portion of it. Finally, solving the complex variational problems that result in our setting to establish optimality across all possible scheduling rules is another key technical contribution in this work.

1.1 Related Work

For systems with complete CSI, there is a rich body of work on throughput-optimal scheduling algorithms, starting from and based on the pioneering approach of Tassiulas et al [3] in developing the Backpressure algorithm. Since then, a host of scheduling algorithms such as Max-Weight/Backpressure [3, 4], the Exponential rule [5–7] and the Log rule [8] have been developed for scheduling using full CSI, and their throughput and delay performance well-studied. A number of optimality results have been shown in the sense of delay/queue-length for many of the aforementioned full-CSI algorithms in various flavours. These include expected queue length/delay bounds via Lyapunov function techniques [9, 10], tail probability decay rates for queue lengths [11, 12, 8, 13, 14, 7, 15], heavy-traffic analysis [16] etc.

Throughput-maximizing scheduling has also been studied for systems with different kinds of partial CSI, such as subset-based CSI [2], delayed CSI [17], infrequent CSI [17] etc. However, to date, there are no known studies for the delay performance of scheduling algorithms when only partial CSI is available. Additionally, the structure of delay-optimizing scheduling algorithms has remained open in the wireless scheduling literature.

1.2 Main Contributions

Our main contributions in this work are as follows:

1. We present a new scheduling algorithm called Max-Exp for scheduling over a wireless downlink when channel state information is restricted to a collection of observable subsets of users. Max-Exp picks, at each time, a subset of channel states to observe depending on an appropriate exponentiated sum of the subset's queue lengths. Having done that, it uses the well-known Exponential rule [5] to schedule a user from the observed subset. Thus, Max-Exp does not use any additional information (e.g. data or channel statistics etc.) other than user queue lengths to dynamically pick subsets, and only the instantaneously observed subset channel states to schedule a user. Also, Max-Exp reduces to the Exponential rule in the case of full CSI, and to the Max-Queue rule (i.e. schedule the user with the largest queue) when the observable subsets are all the single users in the system, and is thus naturally generalizes both these strategies.

2. 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. In this regard, a key technical contribution is to develop large deviations properties for processes obtained by predictably sampling iid sequences. This analysis differs from existing large-deviations approaches for the case of full CSI – such approaches cannot be applied to analyze algorithms using partial CSI as the randomness of the channel state is selectively revealed (i.e. the algorithm gets to "see" the desired subset's channel state) depending on the algorithm's decision at that time. This means that the subset-selection behaviour of the scheduling algorithm can influence the rate function in a crucial way, resulting in a significantly different and more complex variational problem for buffer overflow than the ones studied for the case of full CSI.

3. In the opposite sense, we obtain a uniform upper bound on the buffer overflow rate function of *any* scheduling policy that accesses partial (subset-based) CSI. To establish this, we follow the approach of using "exponentially twisted" channel state distributions. Here again, a technical challenge arises due to the fact that for an arbitrary scheduling algorithm, the large-deviations "cost" of buffer overflow depends crucially on its subset sampling behaviour – different scheduling algorithms could sample subsets with vastly differing frequencies resulting in potentially different costs to twist channel state distributions of subsets, and hence different rate functions. We develop a novel martingale-based technique to quantify this effect and derive a universal upper bound on the buffer overflow exponent.

4. Assuming that 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 exponent across all scheduling rules using partial CSI*. As a corollary, it follows 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. Again, technically, showing that the lower and upper bounds match involves analyzing a complex and non-convex variational problem arising from the rate function for predictably sampled random processes, and is another key contribution of this work.

1.3 Organization

This document is organized as follows: we first describe the model of the wireless system that we use to study scheduling strategies and their performance (Section 2), and state our main results and discuss their significance (Section 3). Then we set up large-buffer scalings and preliminaries that we need to analyze wireless scheduling performance (Section 4), and describe the derivations of our main results for two wireless scheduling algorithms in Sections 5 and 6.

2 System Model

This section is devoted to describing the wireless system model we use, and the associated statistical assumptions. We consider a model of a wireless downlink system that is standard and widely used in existing literature [4, 2, 9, 10]: a time-slotted system of N users serviced by a single base station or server across N communication channels. In each time slot $k \in \{0, 1, 2, \ldots\}$, the dynamics of the system are governed by three primary components as follows:

1. **Arrivals:** An integer number of data packets $A_i(k)$ arrives to user i, $i = 1, \ldots, N$. Packets get queued at their respective users if they are not immediately transmitted.

2. **Channel states:** The set of N channels assumes a random *channel state* $R(k)$, i.e. an N-tuple of integer *instantaneous service rates* $(r_1, \ldots, r_N)$. At time slot k, let the instantaneous service rates be $(R_1(k), \ldots, R_N(k))$.

3. **Scheduling:** One user $U(k) \in \{1, \ldots, N\}$ 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

$$Q_i(k+1) = [Q_i(k) + A_i(k) - D_i(k)R_i(k)]^+, \tag{1}$$

where $x^+ \equiv \max(x, 0)$.

We make the following assumptions about the probabilistic structure of the arrival and channel state processes, and the information structure that is used to schedule users:

1. **Arrivals:** For the sake of notational simplicity, we assume that each user i's arrival process $(A_i(k))_{t=0}^{\infty}$ is deterministic and equal to λ_i at all time slots.

2. **Channel states:** The (joint) channel states $R(k)$, $k = 0, \ldots$ are independent and identically distributed, taking values in a finite set $\mathcal{R}$ of integer N-tuples.

3. **Scheduling:** We study scheduling with *partial channel state information.* In this framework, at each time slot k, the scheduling algorithm must make *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.

   Thus, the instantaneous channel state information available to the scheduling algorithm is limited at each time to a subset of channels which the algorithm can specify, as opposed to the case where the entire joint channel state is revealed to the algorithm prior to scheduling.

3 Main Results

As discussed in the introduction, our focus in this work is on designing wireless scheduling algorithms that reduce the likelihood of large queues in the system. That is, we seek to minimize the stationary probability (if it exists) that the longest queue in the system $||Q(k)||_\infty \stackrel{\triangle}{=} \max_i Q_i(k)$ exceeds a given threshold n. Characterizing the stationary queue length distribution in such systems is analytically difficult, so towards this end, we focus on the rate of decay of the exceedance probability as $n \to \infty$ (widely known as the *large buffer* scaling regime). A series of prior work [11, 12, 8, 13, 14, 7, 15] has established large deviation principles in various forms for the steady-state queue lengths under scheduling when the complete state of the system is known, i.e. shown that the large buffer overflow probability decays to zero exponentially fast. In the same spirit, we concern ourselves with the (exponential) rate of decay

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

(when the limit exists), for scheduling algorithms that *can observe only partial channel state* prior to scheduling. The usefulness of this exponent stems from the approximation

$$\mathbb{P}[||Q(k)||_\infty \geq n] \approx e^{-nI(\lambda)} \qquad \text{(large } n\text{)}$$

that it affords for the actual overflow probabilities. Also, as pointed out in [14], for the case of deterministic arrivals (treated here), the tail of the probability distribution for the delay $D_i(k)$ until transmission experienced by a packet for user i arriving at time k is directly related to that of its queue length by

$$\mathbb{P}[D_i(k) \geq d_i] = \mathbb{P}[Q_i(k) \geq \lambda_i(d_i - 1)];$$

thus, insights can be gained about the exponents of packet delays from those of the system queue lengths.

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

- Scheduling policy **Max-Exp**: At time slot k,
  1. Choose the subset $S(k)$, from the collection $\mathcal{O}$ of observable subsets, such that
  $$\sum_{i \in S(k)} \exp\left(\frac{Q_i(k)}{1 + \sqrt{\overline{Q}(k)}}\right)$$
  is maximized (here $\overline{Q}(k) \triangleq \frac{1}{N}\sum_{i=1}^{N} Q_i(k)$).
  2. Pick a user i, from $S(k)$, such that $R_i(k) \exp\left(\frac{Q_i(k)}{1+\sqrt{\overline{Q}(k)}}\right)$ is maximized.

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. 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. In order not to deviate from the main focus of this work, we state without proof that the throughput region of Max-Exp is a superset of that of any other scheduling algorithm, i.e., Max-Exp is *throughput-optimal*. With this we are in a position to state the main result of this work – 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 the Max-Exp algorithm for general disjoint observable subsets)**.** *The following holds when the system's arrival rates λ are within the throughput region of the Max-Exp scheduling algorithm:*

1. *Let $\mathbb{P}$ denote the stationary probability distribution that the Max-Exp algorithm induces on the vector of queue lengths. Then, there exists $J_* > 0$ with*
$$-\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left[||Q(0)||_\infty \geq n\right] \geq J_*.$$
2. *Let π be an arbitrary scheduling rule that induces a stationary distribution $\mathbb{P}^\pi$ on the vector of queue lengths. If the system of observable subsets $\mathcal{O}$ is disjoint, then*
$$-\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}^\pi\left[||Q(0)||_\infty \geq n\right] \leq J_*,$$
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.

Since Max-Exp reduces to the Max-Queue scheduling algorithm (i.e., scheduling the user with the longest queue at each time) when the observable subsets are all the singleton users, an immediate corollary is the following optimality result for Max-Queue when subsets are restricted to singletons (i.e., when there is effectively no channel state information to use for scheduling):

Corollary 2 (Large Deviations Optimality of the Max-Queue algorithm for singleton observable subsets)**.** *If the observable subsets are all the singleton subsets of users, and the system's arrival rates λ are within 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 with singleton subset-based partial channel state information.*

Though Theorem 1 is our chief result, our roadmap towards proving it in the subsequent sections actually involves showing the result of Corollary 2 first, and then generalizing the argument to the setting of general disjoint subsets. This is because the crux of many of the large deviations techniques we develop in this work lies in the analysis of the simpler singleton subset setting. Another reason for this order of work is that technically, Max-Queue can be analyzed with a standard $O(n)$ fluid scaling, whereas showing the optimality property for Max-Exp in general requires using a more complex refined fluid limit framework at the $O(\sqrt{n})$ "local" timescale as carried out in [5, 7].

4 Large Deviations Framework

This section lays down technical preliminaries for the sample-path large deviations techniques that we use to study the large-buffer overflow probabilities of wireless scheduling algorithms. These constructions are standard in the large-deviations analysis of wireless systems [14, 7, 8].

Throughout this work, we denote by $(\Omega, \mathcal{F}, \mathbb{P})$ a common probability space that supports all defined random variables and processes. To relate the problem of finding tail decay rates of queue lengths to a sample-path large deviations setting, fix an integer $T > 0$, and consider a sequence of (independent) queueing systems indexed by $n = 1, 2, \ldots$, each with its own arrival and channel state processes and evolving as described in Section 2. Henceforth, we explicitly reference by the superscript (n) any quantity associated with the nth system. For any (possibly vector-valued) random process $X^{(n)}(k)$, $k = 0, 1, 2, \ldots$ in the nth system, let us define its scaled, shifted and piecewise linear version $x^{(n)}(\cdot)$ on the interval $[-T, 0]$ by

$$x^{(n)}(t) = \begin{cases} \frac{X^{(n)}(n(t+T))}{n} & \text{if } n(t+T) \text{ is an integer;} \\ \frac{X^{(n)}(\lfloor n(t+T) \rfloor)}{n} + \frac{X^{(n)}(\lceil n(t+T) \rceil) - X^{(n)}(\lfloor n(t+T) \rfloor)}{n(n(t+T) - \lfloor n(t+T) \rfloor)} & \text{otherwise.} \end{cases}$$

In other words, we transform the discrete-time process $X^{(n)}(\cdot)$ on $0, 1, 2, \ldots, nT$ to the piecewise linear and continuous process $x^{(n)}(\cdot)$ on $[-T, 0]$ by compressing time by a factor of n, shifting time by T, scaling space by $\frac{1}{n}$ and finally linearly interpolating between the discrete points.

We define the following (discrete) random processes associated to the nth 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. Let $C_\alpha^{(n)}(k)$ denote the total number of time slots before k when subset α was chosen by the scheduling algorithm. For a subset α of channels, let its *sub-state* $R_\alpha^{(n)}(k)$ 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 governed by (1). 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., $M^{(n)}(k) \triangleq \sum_{i=0}^{k} R^{(n)}(i)\delta_{S(i)}$.

Our assumptions on deterministic arrivals and a finite support for the channel state vector $R(k)$ in each time slot k imply that the scaled processes $f_i^{(n)}(\cdot)$, $\hat{f}_i^{(n)}(\cdot)$, $c_\alpha^{(n)}(\cdot)$, $g_r^{\alpha,(n)}(\cdot)$, $\hat{g}_{ri}^{\alpha,(n)}(\cdot)$, $q_i^{(n)}(\cdot)$ and $m^{(n)}(\cdot)$, on the interval $[-T, 0]$, are uniformly Lipschitz-continuous (with Lipschitz constant L, say) for all n. By the Arzelà-Ascoli theorem, as $n \to \infty$, there must exist a subsequence along which these processes converge uniformly on $[-T, 0]$ to corresponding "limit functions" $f_i(\cdot)$, $\hat{f}_i(\cdot)$, $c_\alpha(\cdot)$, $g_r^\alpha(\cdot)$, $\hat{g}_{ri}^\alpha(\cdot)$, $q_i(\cdot)$ and $m(\cdot)$ defined on $[-T, 0]$. We call any such collection of joint limit functions (i.e., obtained via appropriately scaled prelimit sample paths) a *Fluid Sample Path (FSP)* (and will 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 L) from their corresponding prelimit processes indexed by n, and are thus differentiable Lebesgue-almost everywhere.

5 Singleton Subsets – Max-Queue Tail Optimality

Our approach towards showing our main result – the asymptotic optimality of the Max-Exp algorithm for general disjoint observable subsets – is to first treat the simpler setting where the (disjoint) observable subsets are all the singleton users in the system. This section is devoted to establishing the buffer-overflow optimality of the Max-Queue algorithm. Thus, our standing assumption throughout this section is that the observable subsets are just all the singleton users, i.e. $\mathcal{O} = \{\{i\} : 1 \le i \le N\}$. We use the subscript i to refer to subsets α. Thus, scheduling policies are essentially sampling policies, and Step 2 in scheduling is trivial (schedule the user whose channel state is observed). Consider the Max-Queue (MQ) scheduling/sampling policy, under which the user with the largest queue length (with say a fixed tie-breaking order) is observed and scheduled. In this section, the goal is to show that Max-Queue yields an optimal buffer-overflow exponent, using the following outline:

1. *(Large Deviations Properties For Sampling on Finite Intervals)* Show a lower bound on the overflow exponent of MQ for a finite horizon starting from empty queues, using sample path large deviations for sampled processes and connections to associated variational problems

2. *(Large Deviations Properties for the Stationary Queue Distribution)* Extend the finite-horizon result to a lower bound over an "infinite horizon", i.e., on the stationary queue overflow probability for MQ, using techniques from Friedlin-Wentzell theory [14, 7]

3. *("Straight-Line" Large Deviations Bounds for all Policies)* Provide a way of upper-bounding the queue overflow exponent uniformly for *any* scheduling policy overflow exponent by "twist-

ing" the subset marginal channel state distributions, constructing "straight-line" overflow events and calculating their probabilities

4. *(Optimality of Max-Queue's rate function)* Show that the lower bound on the rate function for the stationary queue overflow probability for MQ actually acts as a uniform upper bound on the rate function for overflow for *any* scheduling algorithm

5.1 Large Deviations for Sampled Processes

Throughout this subsection, we consider the queueing system operating under an arbitrary scheduling policy that is *deterministic*, i.e. does not use additional randomness apart from that of the channel states. Fix $T > 0$. For $q_0 \in \mathbb{R}^N$, let $\mathbb{P}_{q_0}^{n,T}$ be the probability measure of the n-th queueing system conditioned on starting the system at $Q^{(n)}(0) = nq_0$ (i.e. $q^{(n)}(-T) = q_0$). 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.

Our first important result is a sample-path large deviations lower bound on the queue length process over a finite horizon of time slots $0, 1, \ldots, nT$.

Proposition 1. *Let $\mathcal{Q}$ be a closed set in $\mathbb{R}^N$, and let Γ be a closed set of trajectories in $\mathcal{C}_{\mathcal{L}}^{+}([-T, 0])$. Then, under any deterministic scheduling policy,*

$$
\begin{aligned}
& -\limsup_{n\to\infty} \frac{1}{n} \log \sup_{q_0 \in \mathcal{Q}} \mathbb{P}_{q_0}^{n,T}\left[q^{(n)} \in \Gamma\right] \\
& \geq \inf_{(m^T, c^T, q^T)} \quad \int_{-T}^{0} \left[\sum_{i=1}^{N} \dot{c}_i(t) \Lambda_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)}\right)\right] dt \qquad (2) \\
& \quad \text{subject to} \quad (m^T, c^T, q^T) \text{ an FSP}, \\
& \qquad\qquad\qquad q^T(-T) \in \mathcal{Q}, q^T \in \Gamma,
\end{aligned}
$$

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

Proposition 1 essentially provides the sample-path large deviations rate function for the queue-length process, in terms of the sample paths of the subset/user selection and channel state processes. It informs us that the "right" sample-path large deviations rate function to consider – for algorithms that sample subsets of channels and schedule using that information – is a combination of conventional Mogulskii-type rate functions for the subsets (Λ_i^*) weighted by their corresponding fluid selection frequencies ($\dot{c}_i$). Note the crucial and novel dependence of the rate function on the subset selection process (captured by weighting Λ_i^* by $\dot{c}_i$) – 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 policy-dependent factor $\dot{c}$ (i.e., standard Mogulskii-type rate functions [14, 7]). The proof of the proposition is deferred to Appendix A, and relies on the key fact that the sample-path behaviour of any deterministic scheduling/sampling algorithm is completely specified by the sampled user's index and the observed channel state at all times, and not necessarily on the entire joint channel state process including unobserved channel states.

5.2 Large Deviations Lower Bound for Max-Queue

Having established a lower bound for the large deviations rate function for the probability of queue overflow for a finite horizon T conditioned on a fixed starting state (Proposition 1), in this section we proceed to extend this to a lower bound for the queue overflow rate function for the stationary distribution under Max-Queue. Recall that a unique stationary distribution exists because Max-Queue makes the irreducible and aperiodic system state Markov chain positive recurrent. Intuitively, we expect that the finite horizon probability distribution $\mathbb{P}_{q_0}^{n,T}$ tends to the stationary distribution $\mathbb{P}$, and hence seek to show that further minimizing the right hand side of (2) over all finite horizons $T > 0$ yields a lower bound on this stationary overflow probability.

Such a procedure to extend finite horizon bounds to bounds on the stationary probabilities has been developed in [14, 7], and uses techniques from Friedlin-Wentzell large-deviations theory. A similar approach works in our case, and for the sake of clarity we show only the crucial properties for our model that are needed to obtain the result.

The key result we show in this section is:

Theorem 3. *Let $\mathbb{P}$ denote the stationary probability distribution of the system state under Max-Queue. Then,*

$$\begin{aligned}
&-\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left[||q^{(n)}(0)||_\infty \geq 1\right] \\
&\geq \inf_{T,(m^T,c^T,q^T)} \quad \int_{-T}^{0} \left[\sum_{i=1}^{N} \dot{c}_i(t)\Lambda_i^*\left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)}\right)\right] dt \\
&\quad \text{subject to} \quad (m^T, c^T, q^T) \text{ an FSP in } [-T, 0], \\
&\qquad\qquad\qquad q^T(-T) = 0, ||q^T(0)||_\infty \geq 1, \\
&\qquad\qquad\qquad T \geq 0.
\end{aligned} \tag{3}$$

See Appendix B for the proof of the theorem, which is carried out in a fashion analogous to that of Theorem 8.4 in [7]. For any fluid sample path (m^T, c^T, q^T) feasible for the conclusion (3) of Theorem 3, we have that

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

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

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

with the infimum occurring over all feasible FSPs and at regular points t. We record a further lower bound on the decay rate of the *stationary* queue overflow probability under Max-Queue as

Proposition 2 (Lower bound on rate function under Max-Queue)**.** *When λ is in the interior of the Max-Queue throughput region, if $\mathbb{P}$ denotes the stationary measure of Max-Queue, then*

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

5.3 Large Deviations Uniform Upper Bound across all Policies

This section is devoted to exhibiting a uniform upper bound on the decay rate of the probability of buffer overflow, for *any* scheduling policy scheduling by observing a single channel state at each time. This bound helps us to eventually prove the large-deviations optimality of Max-Queue across all scheduling algorithms.

We present our main result here that develops a uniform upper bound on the large deviations rate function over *all* stabilizing scheduling policies. Note that by a stabilizing scheduling policy π, we mean a scheduling rule that, using the state of the queues to pick a user/channel to schedule (possibly with randomization), makes the discrete time Markov chain of queue lengths aperiodic, irreducible and positive recurrent.

Theorem 4 (Upper bound on rate function under any policy)**.** *Let π be a stabilizing scheduling policy for the arrival rate $\lambda = (\lambda_1, \ldots, \lambda_N)$, and let $\mathbb{P}^\pi$ be its associated stationary measure. For any $\phi'_i \in \mathbb{R}^+$, $i = 1, \ldots, N$,*

$$-\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}^\pi \left[||q^{(n)}(0)||_\infty \geq 1 \right] \leq \sup_{\substack{\sum_i c'_i = 1 \\ c'_i \geq 0}} \frac{\sum_i c'_i \Lambda^*_i(\phi'_i)}{\max_i (\lambda_i - c'_i \phi'_i) \vee 0}. \tag{4}$$

Theorem 4 is a key ingredient towards our optimality result for Max-Queue, and states that 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. The weighted cost here is the same sample-path large deviations rate function introduced in Theorem 1, and this upper bound is valid for any "twisted" distributions ϕ'_i that all the subsets' channel states can assume. The proof of Theorem 4 requires a novel martingale technique to establish the upper bound, and we refer the reader to Appendix C for the full details of the proof.

5.4 Large Deviations Optimality of The Max-Queue Policy

In this section, we finally establish that the simple Max-Queue scheduling policy yields the optimum decay rate of the tail of the maximum queue length over *all* stabilizing scheduling policies. This is accomplished by using the lower bound for Max-Queue and uniform upper bound for all scheduling policies on the large deviations rate function from Proposition 2 and Theorem 4 respectively. The result we show is:

Theorem 5 (Large-Deviations Optimality of Max-Queue for singleton channel subsets)**.** *There exist $\hat{\phi}'_1, \ldots, \hat{\phi}'_N$, with $\lambda \notin \mathcal{C}(\hat{\phi}'_1, \ldots, \hat{\phi}'_N)$, such that*

$$\sup_{\substack{\sum_i c'_i = 1 \\ c'_i \geq 0}} \frac{\sum_i c'_i \Lambda^*_i(\hat{\phi}'_i)}{\max_i \left(\lambda_i - c'_i \hat{\phi}'_i \right)} \leq J_*,$$

i.e., Max-Queue has the optimal large-deviations exponent (equal to J_*) over all scheduling policies.

Showing the above result involves solving the complex variational problem for the rate function lower bound given by Proposition 2, and relating the solution to a suitable uniform upper bound of the type prescribed by Theorem 4. We defer the technical details of the proof to Appendix D.

6 General Subsets – Max-Exp Tail Optimality

Having shown the optimality of the Max-Queue algorithm for scheduling with singleton observable channel subsets, we turn to establishing the optimality property for the Max-Exp algorithm when scheduling with an arbitrary collection of mutually disjoint observable subsets of channels. We show, in this section, that the essential techniques and arguments for the optimality of Max-Queue can be adapted to the general setting in order to prove that Max-Exp yields the best rate function for buffer overflow across all scheduling policies that sample channel states from subsets and schedule users.

We start by recalling the definition of the Max-Exp scheduling rule:

Scheduling policy **Max-Exp**: At time slot k,

1. Pick a subset α, from the collection $\mathcal{O}$ of observable subsets, such that

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

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

2. Pick a user i, from α, such that $R_i(k) \exp\left(\frac{Q_i(k)}{1+\sqrt{\overline{Q}(k)}}\right)$ is maximized.

To show the large-deviations queue-overflow optimality of Max-Exp, we will follow the same plan used in the previous section to demonstrate the optimality of Max-Queue, viz.

1. Establish a *lower bound* on the overflow exponent of Max-Exp using sampling-based large deviations arguments, over a finite horizon, in terms of the solution to an associated variational problem

2. Extend the finite horizon result to a lower bound on the stationary queue overflow probability using standard Friedlin-Wentzell theory [14, 7]

3. Show, straight-line large-deviations *upper bounds*, that hold uniformly over all scheduling policies, i.e. all subset-selection and user scheduling strategies

4. Show that the lower bound on the rate function for the stationary queue overflow probability for Max-Exp actually acts as a uniform upper bound on overflow across *all* scheduling algorithms, thereby showing the optimality of Max-Exp

We remark that the Max-Exp rule is not *scaling-invariant* (i.e., scaling all queue-lengths by a uniform constant changes the scheduling behaviour), and that it naturally operates at the $O(\sqrt{n})$ time-scale (i.e., with all the queue lengths $O(n)$, a $O(\sqrt{n})$ change in them causes a shift in scheduling behaviour), unlike Max-Queue which is naturally coupled to the $O(n)$ timescale. This prevents us from using the standard fluid scaling to analyze the fluid sample path behaviour of Max-Exp. Instead, in Step 1 above, we will need to first establish a “refined” Mogulskii-type theorem for sample-path large deviations of predictably sampled processes over a $O(\sqrt{n})$ timescale, and next use the notion of *Local Fluid Sample Paths* (LFSPs), first introduced in [7], that allows us to “magnify” fluid limits at a point in time and prove the large deviations lower bounds in steps 1 and 2 above.

6.1 Sampling-based Finite-horizon LD Lower Bound

Our aim in this section is to extend the sampling-based large-deviations bound of Proposition 1 to hold over a finer-than-$O(n)$ timescale. Towards this, as used in [7], let us consider a positive integer function $u(n) \to \infty$ as $n \to \infty$ with $u(n)/n \to 0$ (e.g. $u(n) = \lceil \sqrt{n} \rceil$). For any non-decreasing RCLL vector-valued function h defined on $[0, \infty)$, denote by $U^n h$ the continous, 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. Also, for $t \geq 0$, let $\theta^{(n)}(t)$ be the largest multiple of $u(n)/n$ not exceeding t.

For each observable subset α, let Λ^*_α be the Sanov rate function for the marginal empirical distribution of its joint sub-state [18], thus the domain of Λ^*_α is the $|\mathcal{R}_\alpha|$-dimensional simplex where $\mathcal{R}_\alpha$ is the set of all possible sub-states for subset α.

In all that follows, we expand the definition of a Fluid Sample Path (FSP) to include-

1. A prelimit “refined cost” function $\bar{J}^{(n)}(\cdot)$, defined over $[-T, 0]$ as
$$\bar{J}^{(n)}(t) \triangleq \int_{-T}^{\theta^{(n)}(t)} \left[\sum_\alpha (U^n c_\alpha^{(n)})'(u) \cdot \Lambda^*_\alpha \left(\frac{(U^n g^{\alpha,(n)})'(u)}{(U^n c_\alpha^{(n)})'(u)} \right) \right] du,$$
where $g^{\alpha,(n)} \equiv \left(g_r^{\alpha,(n)} \right)_{r \in \mathcal{R}_\alpha}$.

2. The uniform convergence (in $[-T, 0]$)
$$\bar{J}^{(n)} \to \bar{J}$$
to some non-negative non-decreasing Lipschitz-continuous “refined cost” function $\bar{J}$, as $n \to \infty$.

Let us also define, for general Lipschitz-continuous functions $(c_\alpha)_\alpha$ and $(g_r^\alpha)_{\alpha r}$ of the appropriate vector dimension on $[-T, 0]$,
$$J(t) \equiv J_{(c,g)}(t) \triangleq \int_{-T}^{t} \left[\sum_\alpha \dot{c}_\alpha(u)\; \Lambda^*_\alpha \left(\frac{\dot{g}^\alpha(u)}{\dot{c}_\alpha(u)} \right) \right] du. \tag{5}$$

We can now state our refined-timescale sample-path large deviations lower bound:

Proposition 3. *Let Γ be a closed set of trajectories in $\mathcal{C}^+_{\mathcal{L}}([-T, 0])$. Then, under a deterministic scheduling policy,*
$$\begin{aligned} &-\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}_0^{n,T} \left[q^{(n)} \in \Gamma \right] \\ &\geq \inf\{\bar{J}_t : (q, \bar{J}) \text{ } FSP \text{ } on \text{ } [-T, 0], q \in \Gamma, t \in [-T, 0]\}. \end{aligned} \tag{6}$$

6.2 Extending The Lower Bound to The Stationary Queue Distribution

As with the approach followed to extend the result of Proposition 1 to the stationary measure under Max-Queue (i.e., to Theorem 3), we can use standard Friedlin-Wentzell-type techniques to extend Proposition 3 to a large-deviations lower bound [14, 7] for the stationary measure under the Max-Exp scheduling policy. Note that this requires showing that Max-Exp is throughput-optimal – a fact whose proof we omit for brevity, but which results from a straightforward modification of the proof of throughput-optimality of the Max-Sum Queue scheduling algorithm (see [19] for details).

Theorem 6. *Let $\mathbb{P}$ denote the stationary measure induced by the Max-Exp policy. Then,*

$$\begin{aligned} &- \limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left[||q^{(n)}(0)||_\infty \geq 1 \right] \\ &\geq \inf_{T\geq 0} \inf\{\bar{J}_t : (q, \bar{J}) \text{ FSP on } [-T,0], q(-T) = 0, q(t) \geq 1, t \in [-T,0]\}. \end{aligned} \tag{7}$$

6.3 Straight-line Uniform LD Upper Bounds over all policies

In this section, we establish a crucial *upper* bound on decay rate of the stationary queue-overflow probability *uniformly* for any stabilizing scheduling policy, along the lines of Theorem 4. This is stated and carried out in terms of "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, $\mathcal{R}_\alpha$ 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, 2]). The uniform large-deviations upper-bound can now be stated for any stabilizing scheduling policy π:

Theorem 7. *Let π be a stabilizing scheduling policy for arrival rates $\lambda = (\lambda_1, \ldots, \lambda_n)$, and let $\mathbb{P}^\pi$ 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,*

$$- \liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}^\pi \left[||q^{(n)}(0)||_\infty \geq 1 \right] \leq \sup_{\substack{\sum_\alpha c'_\alpha = 1 \\ c'_\alpha \geq 0}} \left[\frac{\sum_\alpha c'_\alpha \Lambda^*_\alpha(\phi'_\alpha)}{\max_{\alpha, v_\alpha \in V_{\phi'_\alpha}} \max_{i\in\alpha}(\lambda_i - c'_\alpha v_{\alpha,i})} \right]. \tag{8}$$

6.4 Showing Max-Exp's Overflow Exponent is Optimal

Finally, in this section, we establish that the large-deviations buffer overflow exponent for the Max-Exp scheduling algorithm is in fact optimal over all stabilizing scheduling rules. For this, we leverage the large-deviations lower bound for the Max-Exp scheduling algorithm (Theorem 6) and show that it is actually a uniform upper bound over all scheduling rules as prescribed by Theorem 7.

Roadmap:

1. Consider a feasible FSP $(q, \bar{J})$ on $[-T, 0]$ for Theorem 6, i.e., $q(-T) = 0$, $q(t) = 1$ for some $t \in [-T, 0]$. We will show, by "magnifying" the FSP about some $\tau \in [-T, 0]$ and taking

"local" fluid limits, that the unit large-deviations cost of raising the maximum queue in the associated Local Fluid Sample Path (LFSP) [20, 7] at τ is also $\bar{J}(T)$.

2. Thus, a further lower bound on the Max-Exp rate function is the least large-deviations cost per unit increase of maximum queue, over feasible local fluid sample paths – call it J_*.

3. In the context of Theorem 7, we exhibit suitable twisted subset distributions $\phi'_\alpha \in \Pi_\alpha \ \forall \alpha$ such that the RHS of (8) is at most J_*, proving the claimed result.

6.4.1 From Low Cost FSPs to Low Cost Local FSPs

The variational problem on the right-hand side of (7) necessitates a closer look at the derivatives of fluid sample paths under the Max-Exp scheduling algorithm. At the same time, since the Max-Exp rule naturally operates at the $O(\sqrt{n})$ timescale, derivative information typically is "washed out" of the standard $O(n)$-scaled fluid sample paths. This motivates us to define and use Local Fluid Sample Paths (LFSPs) with a $O(\sqrt{n})$-type scaling, in which information about scheduling choices and drifts can be clearly understood with regard to the Max-Exp scheduling rule.

The formal LFSP construction is along the lines of that used in [20, 7], and is as follows. Consider a standard fluid sample path on $[-T, 0]$ (along with its prelimit functions) and call it ψ. Let us introduce the "recentered" queue lengths

$$\tilde{Q}_i^{(n)}(t) \triangleq Q_i^{(n)}(t) - b_i \sqrt{\bar{Q}^{(n)}(t)},$$

where b_i, $i = 1, \ldots, N$ are such that for every (disjoint) observable subset α, the vector $(e^{b_i})_{i \in \alpha}$ is an outer normal to the subset rate region V_α (under the natural marginal distribution of the sub-state $R_\alpha(1)$) at some point $v^*_\alpha \in V_\alpha$ such that $v^*_\alpha > \lambda|_\alpha$. The fluid-scaled version of $\tilde{Q}_i^{(n)}$ is

$$\tilde{q}_i^{(n)}(t) = q_i^{(n)}(t) - \frac{b_i}{\sqrt{n}} \sqrt{\bar{q}^{(n)}(t)},$$

so we have the uniform convergence

$$\tilde{q}_i^{(n)} \to q_i,$$

and

$$\tilde{q}_*^{(n)} \triangleq \max_i \tilde{q}_i^{(n)} \to q_* \triangleq \max_i q_i.$$

Let $\tau \in [-T, 0]$ be fixed, such that $q_*(\tau) > 0$. Also, fix $S > 0$ and set $\sigma_n \triangleq \frac{1}{\sqrt{n}} \sqrt{\bar{q}^{(n)}(\tau)}$. Suppose we pick, for each n, a sequence of time intervals $[t_1^{(n)}, t_2^{(n)}] \subseteq [-T, 0]$, such that $t_2^{(n)} - t_1^{(n)} = S\sigma_n$ and $t_1^{(n)} \to \tau$ as $n \to \infty$. Then, for each n and $s \in [0, S]$, consider the following "centered" and “rescaled” functions-

$$
\begin{aligned}
{}_\diamond q_i^{(n)}(s) &\triangleq \frac{1}{\sigma_n}[\tilde{q}_i^{(n)}(t_1^{(n)}+\sigma_n s)-\tilde{q}_*^{(n)}(t_1^{(n)})], \quad i=1,\ldots,N,\\
{}_\diamond q_*^{(n)}(s) &\triangleq \max_i {}_\diamond q_i^{(n)}(s) = \frac{1}{\sigma_n}[\tilde{q}_*^{(n)}(t_1^{(n)}+\sigma_n s)-\tilde{q}_*^{(n)}(t_1^{(n)})],\\
{}_\diamond f_i^{(n)}(s) &\triangleq \frac{1}{\sigma_n}[f_i^{(n)}(t_1^{(n)}+\sigma_n s)-f_i^{(n)}(t_1^{(n)})], \quad i=1,\ldots,N,\\
{}_\diamond \hat{f}_i^{(n)}(s) &\triangleq \frac{1}{\sigma_n}[\hat{f}_i^{(n)}(t_1^{(n)}+\sigma_n s)-\hat{f}_i^{(n)}(t_1^{(n)})], \quad i=1,\ldots,N,\\
{}_\diamond c_\alpha^{(n)}(s) &\triangleq \frac{1}{\sigma_n}[c_\alpha^{(n)}(t_1^{(n)}+\sigma_n s)-c_\alpha^{(n)}(t_1^{(n)})], \quad \alpha\in\mathcal{O},\\
{}_\diamond g_r^{\alpha,(n)}(s) &\triangleq \frac{1}{\sigma_n}[g_r^{\alpha,(n)}(t_1^{(n)}+\sigma_n s)-g_r^{\alpha,(n)}(t_1^{(n)})], \quad \alpha\in\mathcal{O}, r\in\mathcal{R}_\alpha,\\
{}_\diamond \hat{g}_{ri}^{\alpha,(n)}(s) &\triangleq \frac{1}{\sigma_n}[\hat{g}_{ri}^{\alpha,(n)}(t_1^{(n)}+\sigma_n s)-\hat{g}_{ri}^{\alpha,(n)}(t_1^{(n)})], \quad \alpha\in\mathcal{O}, r\in\mathcal{R}_\alpha, i=1,\ldots,N,\\
{}_\diamond m^{(n)}(s) &\triangleq \frac{1}{\sigma_n}[m^{(n)}(t_1^{(n)}+\sigma_n s)-m^{(n)}(t_1^{(n)})].
\end{aligned}
$$

It follows that we can choose a subsequence of n along which the following uniform convergence to Lipschitz functions holds on $[0,S]$ [7]:

$$
\begin{aligned}
&\left({}_\diamond q^{(n)}, {}_\diamond q_*^{(n)}, {}_\diamond f^{(n)}, {}_\diamond \hat{f}^{(n)}, ({}_\diamond c_\alpha^{(n)})_\alpha, ({}_\diamond g_r^{\alpha,(n)})_{\alpha r}, ({}_\diamond \hat{g}_{ri}^{\alpha,(n)})_{\alpha r,i}, {}_\diamond m^{(n)}\right) \to\\
&\quad\left({}_\diamond q, {}_\diamond q_*, {}_\diamond f, {}_\diamond \hat{f}, ({}_\diamond g_r^{\alpha})_{\alpha r}, ({}_\diamond \hat{g}_{ri}^{\alpha})_{\alpha r,i}, {}_\diamond m\right). \qquad (9)
\end{aligned}
$$

Note that each ${}_\diamond q_i$ can be either finite Lipschitz or $-\infty$; we appropriately extend the definition of uniform convergence in the latter case. We can also observe that

$$
\liminf_{n\to\infty}\frac{1}{\sigma_n}[\bar{J}^{(n)}(t_2^{(n)})-\bar{J}^{(n)}(t_1^{(n)})] \ge J_{({}_\diamond c, {}_\diamond g)}(S).
$$

The tuple on the right-hand side of (9) above is what we call a *Local Fluid Sample Path* at (scaled) time τ. The following lemma, analogous to Lemma 9.1 in [7], is crucial to understand the local timescale behaviour of the Max-Exp scheduling algorithm:

Lemma 1. *For any LFSP over an interval* $[0,S]$*, for (Lebesgue) almost all* $s\in[0,S]$*, the following derivatives exist and are finite:*

$$
\begin{aligned}
&{}_\diamond \dot{q}(s), \quad {}_\diamond \dot{q}_*(s), \quad {}_\diamond \dot{c}_\alpha(s), \quad {}_\diamond \dot{g}_r^{\alpha}(s), \quad {}_\diamond \dot{\hat{g}}_{ri}^{\alpha}(s), \quad {}_\diamond m(s),\\
&\lambda(s) \triangleq \frac{d}{ds}{}_\diamond f(s), \quad \phi'_{\alpha r}(s) \triangleq \frac{{}_\diamond \dot{g}_r^{\alpha}(s)}{{}_\diamond \dot{c}_\alpha(s)}, \quad v(s) \triangleq \frac{d}{ds}{}_\diamond \hat{f}(s).
\end{aligned}
$$

Moreover, the following relations hold:

$$\lambda(s) = \lambda \quad \forall s \in [0, S],$$
$$_\diamond\dot{q}(s) = \lambda(s) - v(s),$$
$$_\diamond q_*(s) = \max_i {_\diamond q_i(s)},$$
$$_\diamond\dot{q}_*(s) = {_\diamond\dot{q}_i(s)} \quad \textit{for each } i \textit{ such that } {_\diamond q_i(s)} = {_\diamond q_*(s)},$$
$$v_i(s) = \sum_{r \in \mathcal{R}_\alpha} {_\diamond\dot{\hat{g}}^\alpha_{ri}(s)} \mu^\alpha_{ri} \quad \textit{for each } i \in \alpha, \alpha \in \mathcal{O},$$
$$\sum_{i \in \alpha} {_\diamond\dot{\hat{g}}^\alpha_{ri}(s)} = {_\diamond\dot{g}^\alpha_r(s)},$$
$$\sum_{r \in \mathcal{R}_\alpha} {_\diamond\dot{g}^\alpha_r(s)} = {_\diamond\dot{c}_\alpha(s)},$$
$$\sum_{\alpha \in \mathcal{O}} {_\diamond c_\alpha(s)} = 1,$$
$$\forall \alpha \in \mathcal{O} \quad \frac{v_\alpha(s)}{_\diamond\dot{c}_\alpha(s)} = \arg \max_{\eta_\alpha \in V_{\phi'_\alpha}(s)} \left\langle e^{_\diamond q(s)+b}, \eta_\alpha \right\rangle_\alpha,$$
$$\sum_{i \in \beta} e^{_\diamond q_i(s)+b_i} < \max_{\alpha \in \mathcal{O}} \sum_{i \in \alpha} e^{_\diamond q_i(s)+b_i} \Rightarrow {_\diamond\dot{c}_\beta(s)} = 0 \quad \textit{for each } \beta \in \mathcal{O},$$
$$\frac{d}{du} \sum_{i \in \beta} e^{_\diamond q_i(u)+b_i} \Bigg|_{u=s} = \frac{d}{du} \sum_{i \in \gamma} e^{_\diamond q_i(u)+b_i} \Bigg|_{u=s} \quad \textit{for } \beta, \gamma \in \arg \max_{\alpha \in \mathcal{O}} \sum_{i \in \alpha} e^{_\diamond q_i(s)+b_i}.$$

(Note: disjointness of subsets enters implicitly here)

With this framework of LFSPs set up, we can resume the main development from Theorem 6. Consider a feasible FSP $(q, \bar{J})$ on $[-T, 0]$ for the right-hand side of (7) (i.e., for which $q(-T) = 0$ and $q(t) = 1$ for some $t \in [-T, 0]$) whose refined cost is $\bar{J}(t)$. Fix also an arbitrary $\epsilon > 0$. Then, there must exist a time point $\tau \in (-T, t)$ such that $q_*(\tau) > 0$, $q'_*(\tau) > 0$, $\bar{J}'(\tau) > 0$, and

$$\frac{\bar{J}'(\tau)}{q'_*(\tau))} < \bar{J}(t) + \epsilon.$$

Continuing further using techniques similar to that in [7], 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

$$_\diamond q_*(S) - {_\diamond q_*(0)} \geq \epsilon_1 S, \quad \textit{and} \tag{10}$$
$$\frac{J_{(_\diamond c, _\diamond g)}(S) - J_{(_\diamond c, _\diamond g)}(0)}{_\diamond q_*(S) - {_\diamond q_*(0)}} \leq \bar{J}(t) + \epsilon, \tag{11}$$

i.e., we are able to approximate the cost of FSPs arbitrarily well with the "unit cost of raising $_\diamond q_*$" of suitably constructed LFSPs.

6.4.2 Further Lower Bound on Rate Function in Terms of LFSP Costs

We use the techniques of the previous section to further lower-bound the queue overflow exponent of the Max-Exp rule. For a general LFSP, we introduce the following "potential function" of its queue state:

$$\Psi(\diamond q) \stackrel{\triangle}{=} \max_{\alpha\in\mathcal{O}} \Psi_\alpha(\diamond q) \equiv \max_{\alpha\in\mathcal{O}} \sum_{i\in\alpha} e^{\diamond q_i + b_i},$$

together with its logarithm

$$\Phi(\diamond q) \stackrel{\triangle}{=} \log\Psi(\diamond q) = \max_{\alpha} \log\Psi_\alpha(\diamond q).$$

Fact: 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 an FSP feasible for the infimum (7) in Theorem 6. By combining the above fact with the conclusions of the previous section (i.e. the properties (10) and (11)), we have that for an arbitrarily small $\epsilon > 0$, an LFSP can be constructed on some interval $[0, S]$ so that the following properties hold for a suitable $\epsilon_1 > 0$:

$$\Phi(\diamond q(S)) - \Phi(\diamond q(0)) \geq (\epsilon_1/2)S, \tag{12}$$

$$\frac{J_{(\diamond c,\diamond g)}(S) - J_{(\diamond c,\diamond g)}(0)}{\Phi(\diamond q(S)) - \Phi(\diamond q(0))} \leq \bar{J}(t) + 2\epsilon. \tag{13}$$

In the sequel, we will concentrate on the LHS of (13) (modulo an arbitrarily small $\epsilon > 0$, it is a lower bound on the original FSP cost $\bar{J}(t)$). We can write

$$\begin{aligned}
\frac{J_{(\diamond c,\diamond g)}(S) - J_{(\diamond c,\diamond g)}(0)}{\Phi(\diamond q(S)) - \Phi(\diamond q(0))} &= \frac{\int_0^S \frac{d}{ds} J_{(\diamond c,\diamond g)}(s)ds}{\int_0^S \frac{d}{ds}\Phi(\diamond q(s))ds} \geq \inf_{s\in[0,S]} \frac{\frac{d}{ds} J_{(\diamond c,\diamond g)}(s)}{\frac{d}{ds}\Phi(\diamond q(s))} \\
&= \inf_{s\in[0,S]} \frac{\sum_\alpha \diamond\dot{c}_\alpha(s)\ \Lambda^*_\alpha\left(\frac{\diamond\dot{g}^\alpha(s)}{\diamond\dot{c}_\alpha(s)}\right)}{\frac{d}{ds}\Phi(\diamond q(s))} \quad \text{(definition (5))} \\
&= \inf_{s\in[0,S]} \frac{\sum_\alpha \diamond\dot{c}_\alpha(s)\ \Lambda^*_\alpha\left(\phi'_\alpha(s)\right)}{\frac{d}{ds}\Phi(\diamond q(s))} \quad \text{(Lemma 1)}.
\end{aligned}$$

Note that we have slightly abused notation to indicate that the infimum above is, in fact, over the (Lebesgue-a.e.) *regular* points $s \in [0, S]$. As a consequence of the above inequality, we can record the following result:

Proposition 4. *If $\mathbb{P}$ denotes the stationary measure induced by the Max-Exp policy, then*

$$-\limsup_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left[||q^{(n)}(0)||_\infty \geq 1\right] \geq \inf_{s\in[0,S]} \frac{\sum_\alpha \diamond\dot{c}_\alpha(s)\ \Lambda^*_\alpha\left(\phi'_\alpha(s)\right)}{\frac{d}{ds}\Phi(\diamond q(s))}, \tag{14}$$

for any valid Local Fluid Sample Path (LFSP) as specified by (9).

Letting J_* denote the infimum on the RHS of (14) over *all* valid LFSPs, a further lower bound on the buffer overflow exponent of Max-Exp is thus J_*.

6.4.3 Connecting the further Lower Bound to the uniform Upper Bound

The crucial final step in establishing the large-deviations optimality of the Max-Exp algorithm is to show that the lower bound on its decay exponent J_* can actually serve as a uniform upper bound on the decay exponent of any stabilizing scheduling policy, on the lines of Theorem 7.

Theorem 8 (Optimality of Max-Exp). *Let π be any stabilizing scheduling policy for arrival rates $\lambda = (\lambda_1, \ldots, \lambda_n)$, and let $\mathbb{P}^\pi$ be the associated stationary measure. Then,*

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

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.

Proof. Consider an LFSP, specifically the component functions $({}_\diamond q, {}_\diamond c, {}_\diamond g)$, over time $[0, S]$ under the Max-Exp scheduling algorithm. Fix a regular point $s \in [0, S]$. Let

$$\mathcal{O}^* \triangleq \arg\max_{\alpha \in \mathcal{O}} \Phi_\alpha({}_\diamond q(s)) \subseteq \mathcal{O}$$

be the subcollection of "active" observable subsets at time s, i.e., the subsets picked by Max-Exp at s. The regularity of point s and the dynamics of the Max-Exp rule (Lemma 1) implies that the derivatives $\frac{d}{du}\Psi_\alpha(u)\big|_{u=s}$ across all $\alpha \in \mathcal{O}^*$, and $\frac{d}{du}\Psi(u)\big|_{u=s}$, are equal to w', say. For each such α,

$$\begin{aligned} w' &= \sum_{i\in\alpha} e^{{}_\diamond q_i(s)+b_i} (\underbrace{\lambda_i(s)}_{=\lambda_i} - v_i(s)) \\ &= \left\langle e^{{}_\diamond q(s)+b}, \lambda \right\rangle_\alpha - \left\langle e^{{}_\diamond q(s)+b}, v(s) \right\rangle_\alpha \\ &= \left\langle e^{{}_\diamond q(s)+b}, \lambda \right\rangle_\alpha - {}_\diamond\dot{c}_\alpha(s) \left\langle e^{{}_\diamond q(s)+b}, \frac{v(s)}{{}_\diamond\dot{c}_\alpha(s)} \right\rangle_\alpha \\ &= \left\langle e^{{}_\diamond q(s)+b}, \lambda \right\rangle_\alpha - {}_\diamond\dot{c}_\alpha(s) \left[\max_{\eta_\alpha \in V_{\phi'_\alpha(s)}} \left\langle e^{{}_\diamond q(s)+b}, \eta_\alpha \right\rangle_\alpha \right]. \end{aligned} \tag{15}$$

For notational convenience, let us denote, for each α,

$$\rho_\alpha \triangleq \left\langle e^{{}_\diamond q(s)+b}, \lambda \right\rangle_\alpha,$$

$$\xi_\alpha \equiv \xi_\alpha\left(\phi'_\alpha(s)\right) \triangleq \max_{\eta_\alpha \in V_{\phi'_\alpha(s)}} \left\langle e^{{}_\diamond q(s)+b}, \eta_\alpha \right\rangle_\alpha = \sum_{r\in\mathcal{R}_\alpha} \phi'_{\alpha r}(s) \left[\max_{i\in\alpha} \mu^\alpha_{ri} \cdot e^{{}_\diamond q_i(s)+b_i} \right].$$

With this, (15) becomes

$$w' \equiv w'(\phi'_\alpha(s)) = \rho_\alpha - {}_\diamond\dot{c}_\alpha(s) \cdot \xi_\alpha\left(\phi'_\alpha(s)\right).$$

For fixed ${}_\diamond q(s) = q$, the map $\xi_\alpha : \Pi_\alpha \to \mathbb{R}^+$ is linear and hence continuous. Thus, ξ_α induces a good rate function $\tilde{\Lambda}^*_\alpha$ on $\mathbb{R}^+$ [18], given by

$$\tilde{\Lambda}^*_\alpha(\nu'_\alpha) \triangleq \inf\left\{ \Lambda^*_\alpha(\phi'_\alpha) : \phi'_\alpha \in \Pi_\alpha,\ \xi_\alpha(\phi'_\alpha) = \nu'_\alpha \right\}.$$

We have, with $\mathcal{O}^* \subseteq \mathcal{O}^*$ fixed,

$$\begin{aligned} \frac{\sum_{\alpha\in\mathcal{O}^*} {}_\diamond\dot{c}_\alpha(s)\Lambda^*_\alpha(\phi'_\alpha(s))}{\dot{\Psi}(s)} &= \frac{\sum_{\alpha\in\mathcal{O}^*} {}_\diamond\dot{c}_\alpha(s)\Lambda^*_\alpha(\phi'_\alpha(s))}{\rho_\alpha - {}_\diamond\dot{c}_\alpha(s)\cdot\xi_\alpha(\phi'_\alpha(s))} \\ &\geq \inf\left\{ \frac{\sum_{\alpha\in\mathcal{O}^*} c'_\alpha \tilde{\Lambda}^*_\alpha(\nu'_\alpha)}{w'} \,\middle|\, w' > 0, \nu'_\alpha \geq 0, c'_\alpha \geq 0, \sum_{\alpha\in\mathcal{O}^*} c'_\alpha = 1, \rho_\alpha - c'_\alpha\nu'_\alpha = w' \ \ \forall\alpha\in\mathcal{O}^* \right\}. \end{aligned} \tag{16}$$

This exactly corresponds to infimizing the function f^S, given in (37), over the corresponding domain $\mathcal{D}_S$ for the case of singleton observable subsets/individual channels. The correspondence becomes clear when, keeping ${}_\diamond q$ fixed, we *identify each observable subset* α *with a hypothetical queue* having an arrival rate of ρ_α and a "twisted" service rate of ν'_α. Under this correspondence, and due to the fact that $\tilde{\Lambda}^*$ is a (good) rate function, we can employ the same arguments as those in the proof of Lemma 9 to get that

1. There exist $\hat{\nu}'_\alpha \geq 0$, $\alpha \in \mathcal{O}^*$, determining *unique* $\hat{w}' > 0$ and $\hat{c}'_\alpha \geq 0$ feasible for (16), such that the infimum (16) is attained at $(\nu'_\alpha)_{\alpha \in \mathcal{O}^*}$.

2. For every $(d'_\alpha)_{\alpha \in \mathcal{O}^*} \geq 0$ with $\sum_{\alpha \in \mathcal{O}^*} d'_\alpha = 1$, we have

$$\frac{\sum_{\alpha \in \mathcal{O}^*} \hat{c}'_\alpha \tilde{\Lambda}^*_\alpha(\hat{\nu}'_\alpha)}{\hat{w}'} \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \tilde{\Lambda}^*_\alpha(\hat{\nu}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*}(\rho_\alpha - d'_\alpha \hat{\nu}'_\alpha)}. \tag{17}$$

For each of the optimizing $\hat{\nu}'_\alpha$ above, by the lower-semicontinuity of Λ^*_α, we can find $\hat{\phi}'_\alpha \in \Pi_\alpha$ such that $\xi_\alpha(\hat{\phi}'_\alpha) = \hat{\nu}'_\alpha$ and $\tilde{\Lambda}^*_\alpha(\hat{\nu}'_\alpha) = \Lambda^*_\alpha(\hat{\phi}'_\alpha)$. Consider an arbitrary vector $(d'_\alpha)_{\alpha \in \mathcal{O}^*} \geq 0$ with $\sum_{\alpha \in \mathcal{O}^*} d'_\alpha = 1$. Returning to our original LFSP $({}_\diamond q, {}_\diamond c, {}_\diamond g)$, from (16), (17) and the previous remark, we can write

$$\frac{\sum_{\alpha \in \mathcal{O}^*} {}_\diamond\dot{c}_\alpha(s) \Lambda^*_\alpha(\phi'_\alpha(s))}{\dot{\Psi}({}_\diamond q(s))} \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*}(\rho_\alpha - d'_\alpha \cdot \xi_\alpha(\hat{\phi}'_\alpha))}. \tag{18}$$

Considering any $\alpha \in \mathcal{O}^*$, we have

$$\begin{aligned}
\rho_\alpha - d'_\alpha \cdot \xi_\alpha(\hat{\phi}'_\alpha) &= \left\langle e^{{}_\diamond q(s)+b}, \lambda \right\rangle_\alpha - d'_\alpha \cdot \max_{\eta_\alpha \in V_{\hat{\phi}'_\alpha}} \left\langle e^{{}_\diamond q(s)+b}, \eta_\alpha \right\rangle_\alpha \\
&= \left\langle e^{{}_\diamond q(s)+b}, \lambda \right\rangle_\alpha - \max_{\upsilon_\alpha \in d'_\alpha V_{\hat{\phi}'_\alpha}} \left\langle e^{{}_\diamond q(s)+b}, \upsilon_\alpha \right\rangle_\alpha \\
&= \min_{\upsilon_\alpha \in d'_\alpha V_{\hat{\phi}'_\alpha}} \sum_{i \in \alpha} e^{{}_\diamond q_i(s)+b_i} [\lambda_i - \upsilon_{\alpha i}].
\end{aligned} \tag{19}$$

Thanks to the key Lemma 12.2 in [7], we have that there exist

$$l_\alpha > 0, \quad {}_\diamond q^*_{\alpha i} \in [-\infty, \infty), i \in \alpha, \quad \text{and}$$

$$\upsilon^*_\alpha \in \arg \max_{\upsilon_\alpha \in d'_\alpha V_{\hat{\phi}'_\alpha}} \left\langle e^{{}_\diamond q^*_\alpha + b}, \upsilon_\alpha \right\rangle_\alpha$$

such that

$$\begin{aligned}
\forall i \in \alpha \quad \lambda_i - \upsilon^*_{\alpha i} &= l_\alpha, \quad \text{if } e^{{}_\diamond q^*_{\alpha i}} > 0, \\
\lambda_i - \upsilon^*_{\alpha i} &\leq l_\alpha, \quad \text{if } e^{{}_\diamond q^*_{\alpha i}} = 0, \quad \text{and} \\
\min_{\upsilon_\alpha \in d'_\alpha V_{\hat{\phi}'_\alpha}} \sum_{i \in \alpha} e^{{}_\diamond q_i(s)+b_i} [\lambda_i - \upsilon_{\alpha i}] &\leq \Psi_\alpha({}_\diamond q(s))\, l_\alpha \leq \Psi({}_\diamond q(s))\, l_\alpha \\
\Rightarrow \quad \max_{\alpha \in \mathcal{O}^*} \min_{\upsilon_\alpha \in d'_\alpha V_{\hat{\phi}'_\alpha}} \sum_{i \in \alpha} e^{{}_\diamond q_i(s)+b_i} [\lambda_i - \upsilon_{\alpha i}] &\leq \Psi({}_\diamond q(s)) \cdot \max_{\alpha \in \mathcal{O}^*} l_\alpha.
\end{aligned}$$

Using this with (18) and (19) yields

$$
\begin{aligned}
& \frac{\sum_{\alpha \in \mathcal{O}^*} {}_\diamond\dot{c}_\alpha(s) \Lambda^*_\alpha(\hat{\phi}'_\alpha(s))}{\dot{\Psi}({}_\diamond q(s))} \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\Psi({}_\diamond q(s)) \cdot \max_{\alpha \in \mathcal{O}^*} l_\alpha} \\
\Rightarrow \quad & \frac{\sum_{\alpha \in \mathcal{O}^*} {}_\diamond\dot{c}_\alpha(s) \Lambda^*_\alpha(\hat{\phi}'_\alpha(s))}{\left[\frac{\dot{\Psi}({}_\diamond q(s))}{\Psi({}_\diamond q(s))}\right]} \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} l_\alpha} \\
\Rightarrow \quad & \frac{\sum_{\alpha \in \mathcal{O}^*} {}_\diamond\dot{c}_\alpha(s) \Lambda^*_\alpha(\hat{\phi}'_\alpha(s))}{\dot{\Phi}({}_\diamond q(s))} \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} l_\alpha} \\
& \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} \max_{i \in \alpha}(\lambda_i - \upsilon^*_{\alpha i})} \\
& \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} \max_{\upsilon_\alpha \in d'_\alpha V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha}(\lambda_i - \upsilon_{\alpha i})} \\
& \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*, \upsilon_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha}(\lambda_i - d'_\alpha \upsilon_{\alpha i})}. \qquad (20)
\end{aligned}
$$

The above relation holds for any $(d'_\alpha)_{\alpha \in \mathcal{O}^*} \geq 0$ with $\sum_{\alpha \in \mathcal{O}^*} d'_\alpha = 1$. Let $(c'_\alpha)_{\alpha \in \mathcal{O}} \geq 0$ be such that $\sum_{\alpha \in \mathcal{O}} c'_\alpha = 1$. For each $\alpha \in \mathcal{O} \setminus \mathcal{O}^*$, define $\hat{\phi}'_\alpha$ to be the natural probability distribution of sub-states in α, so that $\Lambda^*_\alpha(\hat{\phi}'_\alpha) = 0$ for such α. We can write,

$$
\begin{aligned}
& \frac{\sum_{\alpha \in \mathcal{O}} c'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}, \upsilon_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha}(\lambda_i - c'_\alpha \upsilon_{\alpha i})} = \frac{\sum_{\alpha \in \mathcal{O}^*} c'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}, \upsilon_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha}(\lambda_i - c'_\alpha \upsilon_{\alpha i})} \\
& \leq \frac{\sum_{\alpha \in \mathcal{O}^*} c'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*, \upsilon_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha}(\lambda_i - c'_\alpha \upsilon_{\alpha i})} \\
& \leq \frac{\sum_{\alpha \in \mathcal{O}^*} \tilde{c}'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*, \upsilon_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha}(\lambda_i - \tilde{c}'_\alpha \upsilon_{\alpha i})}, \quad \text{where } \tilde{c}' \triangleq \frac{c'}{\sum_{\alpha \in \mathcal{O}} c'_\alpha}. \qquad (21)
\end{aligned}
$$

Putting (20) and (21) together, we have, for our original LFSP, that

$$
\begin{aligned}
\frac{\sum_\alpha {}_\diamond\dot{c}_\alpha(s) \Lambda^*_\alpha(\hat{\phi}'_\alpha(s))}{\dot{\Phi}({}_\diamond q(s))} & \geq \sup_{\substack{\sum_\alpha c'_\alpha = 1 \\ c'_\alpha \geq 0}} \left[\frac{\sum_\alpha c'_\alpha \Lambda^*_\alpha(\hat{\phi}'_\alpha)}{\max_{\alpha, \upsilon_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha}(\lambda_i - c'_\alpha \upsilon_{\alpha i})} \right] \\
& \geq -\liminf_{n \to \infty} \frac{1}{n} \log \mathbb{P}^\pi \left[||q^{(n)}(0)||_\infty \geq 1 \right], \qquad (22)
\end{aligned}
$$

for the stationary measure $\mathbb{P}^\pi$ of any stabilizing scheduling policy, by Theorem 7. Infimizing (22) over all valid LFSPs and using Proposition 4 yields

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

which finishes the proof. □

7 Conclusions

We developed the Max-Exp and Max-Queue scheduling algorithms for scheduling with only partial wireless Channel State Information (CSI). With tools from the theory of sample-path large deviations, we saw that these algorithms are optimal for the queue overflow exponent for scheduling with partial CSI from subsets of channels.

This work demonstrates that it is possible to design structurally simple scheduling algorithms that intelligently acquire partial CSI and schedule users while guaranteeing high performance. Moreover, to control queue backlogs in such cases, no additional statistical or extraneous information is explicitly required by the scheduling algorithms.

We hope that this work lays the keystone for further investigations of the performance of wireless scheduling under different types of information structures. Future directions for research include studying scheduling with information from general user subsets, temporally varying constraints on available CSI, and performance under delayed CSI with time-correlated channels.

Appendices

A Proof of Proposition 1

Proof. In the nth system, consider the joint channel states for the first nT time slots, i.e., $\big(R^{(n)}(1), R^{(n)}(k), \ldots, R^{(n)}(nT)\big)$, with each $R^{(n)}(k) \in \mathcal{R}^N \subset \mathbb{R}^N$. Since our sampling/scheduling rule is deterministic, the exact time slots in $\{1, \ldots, nT\}$ at which user i is sampled depend entirely on these joint channel states. To avoid heavy notation, we will suppress the superscript (n) as all quantities we deal with refer to the nth queueing system. Let $V = (V_1, \ldots, V_N)$ be the (random) *sampled trace* for the system upto time nT. By this, we mean that each V_i is a vector with elements from $\mathcal{R}$ that represents all the successively observed/sampled rates for user i, i.e. $V_i = (R_i(K_{i_1}), R_i(K_{i_2}), \ldots)$ where user i is chosen precisely at time slots $K_{i_1}, K_{i_2}, \ldots$ In other words, V_i is the ordered row of channel state values sampled by the scheduling policy, so the sum of the lengths of the V_i is exactly nT. In the sequel, we frequently identify each V_i bijectively with its corresponding partial sums process $W_i \equiv W(V_i)$.

We have the following lemma, due to the crucial fact that *for any deterministic sampling rule, the sampled trace uniquely specifies at what times each user was sampled and its sampled channel states at those instants*. By a *valid* sampled trace, we mean a (finite) sampled trace occurring with nonzero probability. For a valid sampled trace w in the n-th system, let $E(w)$ be the set of all extended (i.e. padded upto time nT arbitrarily with values in $\mathcal{R}$) combinations of w, i.e. the set of all $(e_1, \ldots, e_N)$ where each e_i is a size-nT vector with w_i a prefix of it.

Lemma 2. *Let X_{ij}, $i = 1, \ldots, N$, $j = 1, 2, \ldots, nT$ be independent random variables with $X_{ij} \sim R_i(0)$ for all i and j. Let $\hat{\mathbb{P}}^{(nT)}$ be the probability measure induced by $(X_{ij})_{i,j}$. If w is a valid trace in the n-th system, then for any $nq_0 \in (\mathbb{Z}^+)^N$,*

$$\mathbb{P}_{q_0}^{n,T}[W^{(n)} = w] = \hat{\mathbb{P}}^{(nT)}\left[E(w)\right].$$

Proof. Let $w = (w_1, \ldots, w_N)$ with $\sum_{i=1}^N |w_i| = nT$, and let $v = (v_1, \ldots, v_N)$ be the corresponding sampled trace for w, i.e., each w_i is the vector of partial sums for the vector v_i. Associated to w

and v are the time slots $k_{i_1}, k_{i_2}, \ldots$ when user i is sampled, for all i. Furthermore, a key fact is that all the time slots $k_{i_1}, k_{i_2}, \ldots$ when user i is sampled, for all i, are completely specified by v due to the sampling rule being nonrandom.

Recall, from our notation, that the random variable $S(k)$ records which user is sampled at time slot k. We have

$$\begin{aligned}
\mathbb{P}_{q_0}^{n,T}[W^{(n)} = w] &= \mathbb{P}_{q_0}^{n,T}[V^{(n)} = v] \\
&\overset{(a)}{=} \mathbb{P}_{q_0}^{n,T}[V^{(n)} = v, \forall i\; S(k_{i_1}) = i, S(k_{i_2}) = i, \ldots] \\
&= \mathbb{P}_{q_0}^{n,T}\left[\forall i\; R_i(k_{i_1}) = v_{i1}, R_i(k_{i_2}) = v_{i2}, \ldots\right] \\
&\overset{(b)}{=} \prod_{i,j} \mathbb{P}_{q_0}^{n,T}\left[R_i(k_{i_j}) = v_{ij}\right] \\
&\overset{(c)}{=} \hat{\mathbb{P}}^{(nT)}\left[E(w)\right],
\end{aligned}$$

which completes the proof. Here, (a) is by using the key fact in the preceding paragraph; (b) is because channel states are independent across time and the fact that the ${k_{i_j}}_{i,j}$ are all distinct and partition $\{1, 2, \ldots, nT\}$; (c) is due to exchangeability of the (independent across time) channel state process. □

Proceeding with the proof of the proposition, for each $q_0 \in \mathcal{Q}$, we have

$$\mathbb{P}_{q_0}^{n,T}\left[q^{(n)} \in \Gamma\right] \leq \mathbb{P}_{q_0}^{n,T}\left[w^{(n)} \in \Gamma^{(n)}\right],$$

where $\Gamma^{(n)}$ is the set of all valid sampled traces $w^{(n)}$ that, together with some $q_0' \in \mathcal{Q}$, result in queue length paths $q^{(n)} \in \Gamma$ under the scheduling algorithm. For an arbitrary integer $\hat{n}$, we can write

$$\begin{aligned}
\mathbb{P}_{q_0}^{n,T}\left[w^{(n)} \in \Gamma^{(n)}\right] &= \sum_{w \in \Gamma^{(n)}} \mathbb{P}_{q_0}^{n,T}\left[w^{(n)} = w\right] \\
&= \sum_{w \in \Gamma^{(n)}} \hat{\mathbb{P}}^{(nT)}\left[E(w)\right] \quad \text{(by Lemma 2)} \\
&= \hat{\mathbb{P}}^{(nT)}\left[\bigcup_{w \in \Gamma^{(n)}} E(w)\right] \quad \text{(unique prefixes} \Rightarrow \text{disjointness)} \\
&\leq \hat{\mathbb{P}}^{(nT)}\left[\bigcup_{n'=\hat{n}}^{\infty} E\left(\Gamma^{(n')}\right)\right].
\end{aligned}$$

$$\begin{aligned}
\therefore\; -\limsup_{n\to\infty} n^{-1} \log \sup_{q_0 \in \mathcal{Q}} \mathbb{P}_{q_0}^{n,T}\left[w^{(n)} \in \Gamma^{(n)}\right] &\geq -\limsup_{n\to\infty} n^{-1} \log \hat{\mathbb{P}}^{(nT)}\left[\bigcup_{n'=\hat{n}}^{\infty} E\left(\Gamma^{(n')}\right)\right] \\
&\geq \inf\left\{\int_{-T}^{0} \sum_{i=1}^{N} \Lambda_i^*(\dot{w}_i(z))dz : w \in \overline{\bigcup_{n'=\hat{n}}^{\infty} E\left(\Gamma^{(n')}\right)}\right\} \\
&\qquad \text{(by Mogulskii's theorem [18])}
\end{aligned}$$

$$\Rightarrow -\limsup_{n\to\infty} n^{-1} \log \sup_{q_0 \in \mathcal{Q}} \mathbb{P}_{q_0}^{n,T}\left[w^{(n)} \in \Gamma^{(n)}\right] \geq \lim_{\hat{n}\to\infty} \inf\left\{\int_{-T}^{0} \sum_{i=1}^{N} \Lambda_i^*(\dot{w}_i(z))dz : w \in \overline{\bigcup_{n'=\hat{n}}^{\infty} E\left(\Gamma^{(n')}\right)}\right\}. \tag{23}$$

Let the right hand side of (23) be denoted by ζ. For every $\hat{n} = 1, 2, \ldots$, we can choose $w_{\hat{n}}$ such that

$$w_{\hat{n}} \in \overline{\bigcup_{n'=\hat{n}}^{\infty} E\left(\Gamma^{(n')}\right)}, \quad \text{and}$$

$$\lim_{\hat{n}\to\infty} \int_{-T}^{0} \sum_{i=1}^{N} \Lambda_i^*(\dot{w}_{\hat{n},i}(z))dz = \zeta.$$

Since the $w_{\hat{n}}$ are all uniformly Lipschitz continuous and bounded, by the Arzelà-Ascoli theorem, the sequence $(w_{\hat{n}})_{\hat{n}}$ contains a subsequence converging uniformly over the time interval $[-T, 0]$. Without loss of generality, let the subsequence be $\{\hat{n}\}$ itself, and let $\lim_{\hat{n}\to\infty} w_{\hat{n}} = w$. The map $f \mapsto \int_{-T}^{0} \sum_{i=1}^{N} \Lambda_i^*(\dot{f}(z))dz$ is lower-semicontinuous [18], thus

$$\int_{-T}^{0} \sum_{i=1}^{N} \Lambda_i^*(\dot{w}(z))dz \leq \lim_{\hat{n}\to\infty} \int_{-T}^{0} \sum_{i=1}^{N} \Lambda_i^*(\dot{w}_{\hat{n},i}(z))dz = \zeta.$$

We can pick, for each $\hat{n}$, an $\hat{m}_{\hat{n}}$ and a $w_{\hat{m}_{\hat{n}}} \in \bigcup_{n'=\hat{n}}^{\infty} E\left(\Gamma^{(n')}\right)$ such that $||w_{\hat{m}_{\hat{n}}} - w_{\hat{n}}||_\infty < 1/\hat{n}$. Since $w_{\hat{m}_{\hat{n}}} \in \bigcup_{n'=\hat{n}}^{\infty} E\left(\Gamma^{(n')}\right)$, let $w_{\hat{m}_{\hat{n}}} \in E\left(\Gamma^{(\hat{m}'_{\hat{n}})}\right)$ for some $\hat{m}'_{\hat{n}} \geq \hat{n}$. It follows that there exists a corresponding valid queue length path $q_{\hat{m}_{\hat{n}}}$ such that $w_{\hat{m}_{\hat{n}}}$ induces $q_{\hat{m}_{\hat{n}}}$, and moreover, $q_{\hat{m}_{\hat{n}}} \in \Gamma$. We can pick a subsequence of $\{\hat{m}_{\hat{n}}\}_{\hat{n}}$ (let it be $\hat{m}_{\hat{n}}$ without loss of generality) along which $q_{\hat{m}_{\hat{n}}}$ converges to a $q \in \overline{\Gamma} = \Gamma$. We now have $\lim_{\hat{n}\to\infty} w_{\hat{m}_{\hat{n}}} = w$ and $\lim_{\hat{n}\to\infty} q_{\hat{m}_{\hat{n}}} = q$, thus (q, w) is a valid fluid sample path with $q \in \Gamma$ and $q(-T) \in \mathcal{Q}$. This yields

$$-\limsup_{n\to\infty} n^{-1} \log \sup_{q_0 \in \mathcal{Q}} \mathbb{P}_{q_0}^{n,T}\left[w^{(n)} \in \Gamma^{(n)}\right] \geq$$
$$\inf\left\{\int_{-T}^{0} \sum_{i=1}^{N} \Lambda_i^*(\dot{w}_i(z))dz : (w, q) \text{ an FSP}, q \in \Gamma, q(-T) \in \mathcal{Q}\right\}$$

$$\Rightarrow -\limsup_{n\to\infty} \frac{1}{n} \log \sup_{q_0 \in \mathcal{Q}} \mathbb{P}_{q_0}^{n,T}\left[q^{(n)} \in \Gamma\right] \geq$$
$$\inf\left\{\int_{-T}^{0} \sum_{i=1}^{N} \Lambda_i^*(\dot{w}_i(z))dz : (w, q) \text{ an FSP}, q \in \Gamma, q(-T) \in \mathcal{Q}\right\}. \tag{24}$$

Note: By (w, q) being an FSP, in addition to there existing prelimit sequences $w^{(n)} \to w$ and $q^{(n)} \to q$ (uniformly over [-T,0]), we mean that there exist points $z_i \in [-T, 0]$ for all $i = 1, \ldots, N$ such that $z_i^{(n)} \to t_i$ as $n \to \infty$, where $z_i^{(n)}$ is the (scaled by $1/n$) index in the sampled trace $w_i^{(n)}$ beyond which user i is never sampled (i.e. it is the last index at which user i is sampled by the scheduling algorithm if user i's samples are stacked successively and contiguously).

From the observation preceding Lemma 2, each sampled trace $w^{(n)}$ completely specifies the exact instants at which each user was scheduled/sampled and the channel states observed at those instants, i.e. $w^{(n)}$ completely specifies the pair $(m^{(n)}, c^{(n)})$ in $[-T, 0]$. The next lemma relates the large deviations "costs" of the fluid limits of sampled traces to those of the fluid limits of their associated $(m^{(n)}, c^{(n)})$ processes.

Lemma 3. *Let (w, q) be a fluid sample path with $w^{(n)} \to w$ and $q^{(n)} \to q$. Let $(m^{(n)}, c^{(n)})$ be the scaled sampled rate and selection processes specified by $w^{(n)} \to w$. For every subsequential limit (m, c) of $(m^{(n)}, c^{(n)})_n$ (in the $||\cdot||_\infty$ topology on $[-T, 0]$),*

$$\sum_{i=1}^{N} \int_{-T}^{z_i} \Lambda_i^*(\dot{w}_i(z))dz = \int_{-T}^{0} \left[\sum_{i=1}^{N} \dot{c}_i(t) \Lambda_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)} \right) \right] dt.$$

Proof. Assume without loss of generality that $m^{(n)} \to m$ and $c^{(n)} \to c$ uniformly in $[-T, 0]$. Let $-T \le t_1 \le t_2 \le 0$. For all n, by the definition of the sampled traces $w_i^{(n)}$, we have

$$w_i^{(n)} \left(c_i^{(n)}(t_2) \right) - w_i^{(n)} \left(c_i^{(n)}(t_1) \right) = m_i^{(n)}(t_2) - m_i^{(n)}(t_1) + O(1/n). \tag{25}$$

By the (uniform) convergence hypotheses, for $j \in \{1, 2\}$, $c_i^{(n)}(t_j) \to c_i(t_j)$, thus $w_i^{(n)} \left(c_i^{(n)}(t_j) \right) \to w_i(c_i(t_j))$. Letting $n \to \infty$ in (25),

$$w_i\left(c_i(t_2)\right) - w_i\left(c_i(t_1)\right) = m_i(t_2) - m_i(t_1). \tag{26}$$

Since c_i and m_i are nondecreasing Lipschitz-continuous functions, they induces Stieltjes measures dc_i and dm_i respectively on $[-T, 0]$ with $dm_i << dc_i$. In a similar fashion, w_i induces a Stieltjes measure $dw_i << dz$ on $[-T, z_i]$ where dz denotes Lebesgue measure. Let dw_i/dz be the Radon-Nikodym derivative of dw_i with respect to Lebesgue measure, and consider

$$\begin{aligned}
\int_{t_1}^{t_2} \frac{dw_i}{dz} \circ c_i(t) dc_i(t) &= \int_{c_i(t_1)}^{c_i(t_2)} \frac{dw_i}{dz}(z) \left(dc_i \circ c^{-1} \right) \quad \text{(change of variables formula)} \\
&= \int_{c_i(t_1)}^{c_i(t_2)} \frac{dw_i}{dz}(z) dz \quad (dc_i \circ c^{-1} \equiv \text{Lebesgue}[-T, z_i]) \\
&= w_i\left(c_i(t_2)\right) - w_i\left(c_i(t_1)\right) \\
&= m_i(t_2) - m_i(t_1) \quad \text{(thanks to (26))} \\
&= \int_{t_1}^{t_2} dm_i(t) \\
\Rightarrow \quad \frac{dw_i}{dz} \circ c_i(\cdot) &= \frac{dm_i}{dc_i}(\cdot) \quad dc_i\text{-a.e. on } [-T, 0].
\end{aligned}$$

With this, we can finally compute

$$\begin{aligned}
&\int_{-T}^{0} \dot{c}_i(t) \Lambda_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)} \right) = \int_{-T}^{0} \left(\Lambda^* \circ \frac{dm_i}{dc_i} \right)(t) dc_i(t) \\
&= \int_{-T}^{0} \left(\Lambda^* \circ \frac{dw_i}{dz} \circ c_i \right)(t) dc_i(t) = \int_{c_i(-T)}^{c_i(0)} \left(\Lambda^* \circ \frac{dw_i}{dz} \right)(z) \left(dc_i \circ c^{-1} \right) \\
&= \int_{-T}^{z_i} \left(\Lambda^* \circ \frac{dw_i}{dz} \right)(z) dz = \int_{-T}^{z_i} \Lambda^* \left(\dot{w}_i(z) \right) dz.
\end{aligned}$$

This proves the lemma. ☐

Applying the result of Lemma 3 to (24) concludes the proof of Proposition 1. ☐

B Proof of Theorem 3

Proof. Instead of rescaling time to the interval $[-T, 0]$, for the sake of consistency with earlier work, we will shift the time scaling to the interval $[0, T]$, i.e. for any discrete-time process $X^{(n)}$, $n = 0, 1, \ldots$, we will use $x^{(n)}$ to mean the piecewise linearized path of the discrete-time process $\frac{1}{n}X^{(n)}(nt)$, $t = 0, \frac{1}{n}, \frac{2}{n}, \ldots, \frac{T}{n}$. With reference to the proof of a similar theorem (Theorem 8.4) in [7], we can establish the following properties in a completely analogous fashion to complete the proof of the theorem (and hence omit their proofs to avoid repetition):

Lemma 4. *Let $\delta > 0$ and $c > 0$ be given, and let the stopping time $\beta^{(n)} \triangleq \inf\{t \geq 0 : ||q^{(n)}(t)||_\infty \leq \delta\}$. Then, there exists $\Delta > 0$ such that*

$$\limsup_{n\to\infty} \sup_{y:||q(y)||_\infty \leq c} \mathbb{E}_y \beta^{(n)} \leq \Delta c.$$

Lemma 5. *For fixed constants $c > \delta > 0$ and $T > 0$, let*

$$\begin{aligned} K(c, \delta, T) \triangleq \inf_{(m^T, c^T, q^T)} & \int_0^T \left[\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)} \right) \right] dt \\ \text{subject to} \quad & (m^T, c^T, q^T) \text{ an FSP,} \\ & ||q(0)||_\infty \leq c, ||q(t)||_\infty \geq \delta \text{ for all } 0 \leq t \leq T. \end{aligned}$$

Then, uniformly over δ, $K(c, \delta, T) \to \infty$ as $T \to \infty$.

□

C Proof of Theorem 4

Proof. Denote by $\mathcal{C}(\phi'_1, \ldots, \phi'_N)$ the convex hull of the points $(0, \ldots, 0)$, $(\phi'_1, \ldots, 0)$, $(0, \phi'_2, \ldots, 0)$, $\ldots$, $(0, \ldots, \phi'_N)$. The right-hand side of (4) is trivially ∞ if either (a) $\lambda \in \mathcal{C}(\phi'_1, \ldots, \phi'_N)$, or (b) any of the ϕ'_i is not in the effective domain of its corresponding Λ_i^*; we exclude such ϕ'_i and λ in the remainder of the proof.

For each $n = 1, 2, \ldots$, let $t_n \geq 0$ be a nonrandom time, to be specified later (to avoid complications, we assume nt_n is an integer). Consider

$$\begin{aligned} \mathbb{P}^\pi \left[||q^{(n)}(0)||_\infty \geq 1 \right] &= \mathbb{P}^\pi \left[||q^{(n)}(t_n)||_\infty \geq 1 \right] \\ &\geq \mathbb{P}^\pi \left[||q^{(n)}(t_n)||_\infty \geq 1 \mid ||q^{(n)}(0)||_\infty = 0 \right] \mathbb{P}^\pi \left[||q^{(n)}(0)||_\infty = 0 \right] \\ &= \pi((0, 0, \ldots, 0)) \, \mathbb{P}_0^\pi \left[||q^{(n)}(t_n)||_\infty \geq 1 \right]. \end{aligned}$$

Here, $\pi(\cdot)$ is used to denote the stationary distribution that the policy π induces, and $\mathbb{P}_0^\pi$ represents the stationary distribution conditioned on the starting state being the origin (all zeroes).

The non-negativity of queues forces the relation $U^{(n)}(k) \triangleq \lambda n - M^{(n)}(k) \leq Q^{(n)}(k)$, where $U^{(n)}(k) = \sum_{l=1}^k (\lambda - R^{(n)}(l)\delta_{S(l)})$ represents the "unreflected queue lengths" in the n-th queueing system at time k. By suitably rescaling in time and space, we can continue this chain of inequalities as

$$\mathbb{P}^{\pi}\left[||q^{(n)}(0)||_{\infty} \geq 1\right] \geq \pi((0,0,\ldots,0))\,\mathbb{P}_0^{\pi}\left[\max_i u_i^{(n)}(t_n) \geq 1\right]. \tag{27}$$

For each $i = 1,\ldots,N$, since ϕ_i' is in the effective domain of its Cramér rate function Λ_i^*, it follows that there exists $\eta_i' \in \mathbb{R}$ such that $\Lambda_i^*(\phi_i') = \eta_i'\phi_i' - \Lambda_i(\eta_i')$. Define for each i an exponentially tilted measure $\hat{\mathbb{P}}_i$ (with respect to the marginal measure $\mathbb{P}_i$ of the i-th channel state $R_i(0)$) on $\mathbb{R}$ as follows:

$$\hat{\mathbb{P}}_i(dx) \stackrel{\triangle}{=} \exp[\eta_i'x - \Lambda_i(\eta_i')]\,\mathbb{P}_i(dx) = \exp[\eta_i'(x-\phi_i') + \Lambda_i^*(\phi_i')]\,\mathbb{P}_i(dx).$$

A standard computation under the tilted measure yields $\hat{\mathbb{E}}_i[R_i(0)] = \phi_i'$. As with the approach followed in [21], let $\hat{\mathbb{P}}_0^{\pi}$ be the measure defined similarly to $\mathbb{P}_0^{\pi}$ except that the twisted measures $\{\hat{\mathbb{P}}_i\}$ replace $\{\mathbb{P}_i\}$ as the conditional marginal distributions of the sampled channel states/rates, with $\{\hat{\mathbb{E}}_i\}$ being the corresponding expectations.

Let us define

$$t_{\min}^{-1} \stackrel{\triangle}{=} \max_i \lambda_i,$$
$$t_{\max}^{-1} \stackrel{\triangle}{=} \min_{\mu \in \mathcal{C}(\phi_1',\ldots,\phi_N')} \max_i(\lambda_i - \mu_i).$$

Since by hypothesis the arrival rate λ is outside the closed set $\mathcal{C}(\phi_1',\ldots,\phi_N')$, it follows that $0 < t_{\min} \leq t_{\max} < \infty$. The times $t_{\min}$ and $t_{\max}$ represent the earliest and latest time that the maximum queue length can take to overflow to level 1 in a system of queues with "fluid" inputs at rates λ_i that can be drained with instantaneous rates in the convex hull $\mathcal{C}(\phi_1',\ldots,\phi_N')$.

The remainder of the proof is organized into four steps:

1. Showing that for n large enough, under the twisted measure $\hat{\mathbb{P}}$, the service $m_i^{(n)}(t)$ provided to the queue i is approximated with high probability by $\phi_i' c_i(t)$, i.e. we can treat the channel as being deterministic with a service rate of ϕ_i',

2. Under the conditions of the previous step, overflow of the unreflected max-queue $d^{(n)}(\cdot)$ is inevitable by time roughly $t_{\max}$, so with a significant probability the first hitting time of $d^{(n)}(\cdot)$ to level 1 is at most $t_{\max}$. Thus, we can find a time not exceeding $t_{\max}$ at which overflow occurs with a significant probability (i.e. not decaying to 0 exponentially in n)

3. Overflow occurring at the time in the previous step, under the conditions of step 1, forces the scheduling "choice fractions" $c^{(n)}(t)/t$ to be "consistent" with overflow of $d^{(n)}(\cdot)$ occurring at that time

4. Using all the steps to develop the right-hand side of (27) and derive the stated result.

C.0.4 Step 1 of 4

Let us record the following definition. For each $i = 1, \dots, N$, we can write

$$\begin{aligned} m_i^{(n)}(t) \equiv m_i(t) = \frac{1}{n} M_i(nt) = \frac{1}{n} \sum_{l=0}^{nt} R_i(l) X_i(l) = \frac{\overline{M}_i(nt)}{n} + \frac{\phi'_i}{n} \sum_{l=0}^{nt} X_i(l) \\ = \frac{\overline{M}_i(nt)}{n} + \frac{\phi'_i}{n} C_i(nt), \\ = \overline{m}_i(t) + \phi'_i c_i(t), \end{aligned}$$

where $X_i(l)$ is the indicator of the event that user i was scheduled at time slot l, and $\overline{M}_i(k) \stackrel{\triangle}{=} \sum_{l=0}^{k} (R_i(l) - \phi'_i) X_i(l)$ is the (unscaled) "centered" service provided to queue i upto time slot k.

Lemma 6. *Let times t_1 and t_2, such that $0 < t_1 \leq t_2$, and $\delta > 0$ be fixed. Then,*

$$\lim_{n \to \infty} \hat{\mathbb{P}}_0^{\pi} \left[|\overline{m}_i^{(n)}(t)| < \delta t \quad \forall t \in [t_1, t_2] \right] = 1.$$

Proof. Observe that for each i, $\{\overline{M}_i(k)\}_k$ is a martingale (with respect to the measure $\hat{\mathbb{P}}_0^{\pi}$) null at 0 and with differences bounded by $D \stackrel{\triangle}{=} (R_{\max} + \max_i \phi'_i)$, where $R_{\max}$ is the maximum channel rate across all channels in the system. An application of the Azuma-Hoeffding martingale inequality [22] thus gives

$$\hat{\mathbb{P}}_0^{\pi} \left[\left| \frac{\overline{M}_i(k)}{k} \right| \geq \gamma \right] \leq 2 e^{-\frac{k\gamma^2}{2D^2}} \tag{28}$$

for all $k = 1, 2, \dots$. Hence, a union bound gives

$$\begin{aligned} 1 - \hat{\mathbb{P}}_0^{\pi} \left[|\overline{m}_i(t)| < \delta t \quad \forall t \in [t_1, t_2] \right] &= \hat{\mathbb{P}}_0^{\pi} \left[\exists t \in [t_1, t_2] \;\; |\overline{m}_i(t)| \geq \delta t \right] \\ &\leq \sum_{k=nt_1}^{nt_2} \hat{\mathbb{P}}_0^{\pi} \left[\left| \frac{\overline{M}_i(k)}{k} \right| \geq \delta \right] \\ &\leq \sum_{k=nt_1}^{nt_2} 2 e^{-\frac{k\delta^2}{2D^2}} \\ &\leq 2n(t_2 - t_1) e^{-\frac{nt_1\delta^2}{2D^2}} \stackrel{n \to \infty}{\longrightarrow} 0 \quad (\because t_1 > 0), \end{aligned}$$

which is the stated result. ∎

C.0.5 Step 2 of 4

Let us fix $\delta > 0$ small enough, and let $\epsilon > 0$ be such that

$$(t_{\max} + \epsilon)^{-1} = \min_{\mu \in \mathcal{C}(\phi'_1, \dots, \phi'_N)} \max_i (\lambda_i - \delta - \mu_i).$$

Additionally, fix a time $t_0 > 0$ small enough, and let $A \equiv A_n$ denote the event whose (twisted) probability is estimated in Lemma 6, i.e.

$$A_n \equiv A_n(\delta) \equiv A_n(\delta, t_0, t_{\max}) \stackrel{\triangle}{=} \left\{ |\overline{m}_i^{(n)}(t)| < \delta t \quad \forall t \in [t_0, t_{\max}] \right\}.$$

Denote the (unreflected and fluid-scaled) maximum queue length process by $d(\cdot) \equiv d^{(n)}(\cdot) \triangleq \max_i u_i^{(n)}(\cdot)$. It follows that in the event A_n, $d(\cdot)$ must overflow (i.e. hit level 1) at least once by time $(t_{\max} + \epsilon)$. In other words, if we let

$$\tau \equiv \tau_n \triangleq \inf \left\{ t = 0, \frac{1}{n}, \frac{2}{n}, \ldots : d(t) \geq 1 \right\},$$

then $A_n \subseteq \{\tau_n \leq t_{\max} + \epsilon\}$. For each $n = 1, 2, \ldots$, define the (deterministic) time

$$t_n \triangleq \arg \max_{t=0,\frac{1}{n},\ldots,t_{\max}+\epsilon} \hat{\mathbb{P}}_0^\pi \left[\tau_n = t\right]$$

with ties broken in an arbitrary fashion. Observe that t_n does not depend upon δ, t_0 or A_n. Also, note that

$$\begin{aligned} \hat{\mathbb{P}}_0^\pi \left[A_n\right] \leq \hat{\mathbb{P}}_0^\pi \left[\tau_n \leq t_{\max} + \epsilon\right] &\leq \sum_{t=0,\frac{1}{n},\ldots,t_{\max}+\epsilon} \hat{\mathbb{P}}_0^\pi \left[\tau_n = t\right] \\ &\leq n(t_{\max} + \epsilon) \left(\max_{t=0,\frac{1}{n},\ldots,t_{\max}+\epsilon} \hat{\mathbb{P}}_0^\pi \left[\tau_n = t\right] \right) \\ &\leq n(t_{\max} + \epsilon) \hat{\mathbb{P}}_0^\pi \left[\tau_n = t_n\right] \\ \Rightarrow \; \hat{\mathbb{P}}_0^\pi \left[\tau_n = t_n\right] &\geq \frac{\hat{\mathbb{P}}_0^\pi \left[A_n\right]}{n(t_{\max} + \epsilon)} \end{aligned}$$

Since the rate of change of $d^{(n)}(\cdot)$ is bounded by D, we can write

$$\Rightarrow \hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1, 1 + \frac{D}{n}\right] \right] \geq \hat{\mathbb{P}}_0^\pi \left[\tau_n = t_n\right] \geq \frac{\hat{\mathbb{P}}_0^\pi \left[A_n\right]}{n(t_{\max} + \epsilon)}. \tag{29}$$

C.0.6 Step 3 of 4

This step involves showing that when the queues overflow at time t_n then the scheduling choice fractions $c^{(n)}(t_n)/t_n$ at that time are very likely to be the ones that cause "straight-line" overflow at time t_n from the all-empty queue state.

Recall that $\delta > 0$ is a sufficiently small number. We denote by Γ_n the set of *δ-compatible* scheduling fractions for overflow at time t_n as follows:

$$\Gamma_n \equiv \Gamma_n(\delta, t_n) \triangleq \left\{ (f_1', \ldots, f_N') : \sum_i f_i' = 1, f_i' \geq 0, \max_i \left(\lambda_i t_n - \phi_i' f_i' t_n \right) \in \left[1 - \delta t_n, 1 + \delta t_n\right] \right\}.$$

Lemma 7. *For all n large enough and $\delta > 0$,*

$$\hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1 - \frac{D}{n}, 1 + \frac{D}{n}\right], \frac{c(t_n)}{t_n} \notin \Gamma_n \right] \leq 2N e^{-\frac{n t_0 \delta^2}{8D^2}}.$$

Proof. For n sufficiently large,

$$d(t_n) \in \left[1 - \frac{D}{n}, 1 + \frac{D}{n}\right] \Rightarrow d(t_n) \in \left[1 - \frac{\delta}{2} t_n, 1 + \frac{\delta}{2} t_n\right].$$

Also,

$$d(t_n) \in \left[1 - \frac{\delta}{2} t_n, 1 + \frac{\delta}{2} t_n\right], \frac{c(t_n)}{t_n} \notin \Gamma_n \Rightarrow \exists i \; |\overline{m}_i(t_n)| > \frac{\delta}{2} t_n.$$

Thus,

$$\begin{aligned}
&\hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1 - \frac{D}{n}, 1 + \frac{D}{n}\right], \frac{c(t_n)}{t_n} \notin \Gamma_n\right] \leq \hat{\mathbb{P}}_0^\pi \left[\exists i \; |\overline{m}_i(t_n)| > \frac{\delta}{2} t_n\right] \\
&\leq \sum_i \hat{\mathbb{P}}_0^\pi \left[|\overline{m}_i(t_n)| > \frac{\delta}{2} t_n\right] \leq \sum_i 2e^{-\frac{n t_n \delta^2}{8D^2}} \quad \text{(Azuma-Hoeffding (28))} \\
&\leq 2N e^{-\frac{n t_0 \delta^2}{8D^2}}.
\end{aligned}$$

□

Step 4 of 4

We can now finally develop the right-hand side of (27) using the results from the previous steps:

$$\begin{aligned}
&\mathbb{P}_0^\pi \left[\max_i u_i^{(n)}(t_n) \geq 1\right] = \mathbb{P}_0^\pi \left[d^{(n)}(t_n) \geq 1\right] \\
&\geq \mathbb{P}_0^\pi \left[d^{(n)}(t_n) \geq 1, c^{(n)}(t_n)/t_n \in \Gamma_n\right] \\
&= \mathbb{E}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}}\right] \\
&= \hat{\mathbb{E}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} \prod_{l=0}^{n t_n} \exp\left[-\Lambda_{U(l)}^*(\phi'_{U(l)}) - \eta'_{U(l)}\left(R_{U(l)}(l) - \phi'_{U(l)}\right)\right]\right] \\
&= \hat{\mathbb{E}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} \exp\left[-n t_n \left(\sum_i \frac{c_i(t_n)}{t_n} \Lambda_i^*(\phi'_i)\right)\right] \exp\left[-n w(t_n)\right]\right] \\
&\qquad \left(\text{with } w(t_n) \equiv w^{(n)}(t_n) \triangleq \frac{1}{n} W(n t_n) \triangleq \frac{1}{n} \sum_{l=0}^{n t_n} \eta'_{U(l)}\left(R_{U(l)}(l) - \phi'_{U(l)}\right)\right) \\
&= \hat{\mathbb{E}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} \exp\left[-n t_n \left(\sup_{f' \in \Gamma_n} \sum_i f'_i \Lambda_i^*(\phi'_i)\right)\right] \exp\left[-n w(t_n)\right]\right] \\
&= \exp\left[-n t_n \left(\sup_{f' \in \Gamma_n} \sum_i f'_i \Lambda_i^*(\phi'_i)\right)\right] \hat{\mathbb{E}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} e^{-n w(t_n)}\right].
\end{aligned} \tag{30}$$

The second term in the product above can be bounded from below for any $\zeta > 0$ as follows:

$$\begin{aligned}
&\hat{\mathbb{E}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} e^{-n w(t_n)}\right] \geq \hat{\mathbb{E}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n, |w(t_n)| < \zeta\}} e^{-n\zeta}\right] \\
&= e^{-n\zeta} \, \hat{\mathbb{P}}_0^\pi \left[d(t_n) \geq 1, \frac{c(t_n)}{t_n} \in \Gamma_n, |w(t_n)| < \zeta\right] \\
&\geq e^{-n\zeta} \, \hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1, 1 + \frac{D}{n}\right], \frac{c(t_n)}{t_n} \in \Gamma_n, |w(t_n)| < \zeta\right].
\end{aligned} \tag{31}$$

We have

$$
\begin{aligned}
&\hat{\mathbb{P}}_0^\pi\left[d(t_n) \in \left[1, 1+\frac{D}{n}\right], \frac{c(t_n)}{t_n} \in \Gamma_n, |w(t_n)| < \zeta\right] \\
&\geq \hat{\mathbb{P}}_0^\pi\left[d(t_n) \in \left[1, 1+\frac{D}{n}\right], \frac{c(t_n)}{t_n} \in \Gamma_n\right] - \hat{\mathbb{P}}_0^\pi\left[|w(t_n)| \geq \zeta\right] \\
&\geq \hat{\mathbb{P}}_0^\pi\left[d(t_n) \in \left[1, 1+\frac{D}{n}\right]\right] - \hat{\mathbb{P}}_0^\pi\left[d(t_n) \in \left[1, 1+\frac{D}{n}\right], \frac{c(t_n)}{t_n} \notin \Gamma_n\right] - \hat{\mathbb{P}}_0^\pi\left[|w(t_n)| \geq \zeta\right]. \qquad (32)
\end{aligned}
$$

By definition and the properties of the twisted distribution $\hat{\mathbb{P}}$, it can be seen that $\{W(k)\}_{k=0,1,\ldots}$ is again a martingale null at 0 and with bounded increments (bounded by, say, $D_2 \triangleq \max_i \eta_i'(R_{\max} + \phi_i')$). Hence, the Azuma-Hoeffding inequality applied to it yields

$$
\hat{\mathbb{P}}_0^\pi\left[|w(t_n)| \geq \zeta\right] \leq 2e^{-\frac{n\zeta^2}{2t_n D_2^2}} \leq 2e^{-\frac{n\zeta^2}{2t_{\max} D_2^2}}.
$$

Using this and the results of Steps 2 and 3, (32) becomes

$$
\begin{aligned}
&\hat{\mathbb{P}}_0^\pi\left[d(t_n) \in \left[1, 1+\frac{D}{n}\right], \frac{c(t_n)}{t_n} \in \Gamma_n, |w(t_n)| < \zeta\right] \\
&\geq \frac{\hat{\mathbb{P}}_0^\pi\left[A_n\right]}{n(t_{\max}+\epsilon)} - 2Ne^{-\frac{nt_0\delta^2}{8D^2}} - 2e^{-\frac{n\zeta^2}{2t_{\max} D_2^2}} \\
&\geq \frac{1/2}{n(t_{\max}+\epsilon)} - 2Ne^{-\frac{nt_0\delta^2}{8D^2}} - 2e^{-\frac{n\zeta^2}{2t_{\max} D_2^2}} \quad \text{(for } n \text{ large enough, by Step 1).}
\end{aligned}
$$

The first term above decays as n^{-1} while the second and third terms decay exponentially in n, thus

$$
-\liminf_{n\to\infty} \frac{1}{n} \log \hat{\mathbb{P}}_0^\pi\left[d(t_n) \in \left[1, 1+\frac{D}{n}\right], \frac{c(t_n)}{t_n} \in \Gamma_n, |w(t_n)| < \zeta\right] \leq 0. \qquad (33)
$$

What remains is to bound the first term in the product in (30). By definition, for every $f' \in \Gamma_n$, we have

$$
\begin{aligned}
&\max_i(\lambda_i - \phi_i' f_i') \leq \frac{1}{t_n} + \delta \\
\Rightarrow\ & t_n \sup_{f'\in\Gamma_n} \sum_i f_i' \Lambda_i^*(\phi_i') \leq \sup_{f'\in\Gamma_n} \frac{\sum_i f_i' \Lambda_i^*(\phi_i')}{\max_i(\lambda_i - \phi_i' f_i') - \delta} \\
&\leq \sup_{\substack{\sum_i f_i'=1 \\ f_i'\geq 0}} \frac{\sum_i f_i' \Lambda_i^*(\phi_i')}{\max_i(\lambda_i - \phi_i' f_i') - \delta}. \qquad (34)
\end{aligned}
$$

Applying the conclusions of (31), (33) and (34) to (30), we get

$$
-\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}_0^\pi\left[\max_i u_i^{(n)}(t_n) \geq 1\right] \leq \zeta + \sup_{\substack{\sum_i f_i'=1 \\ f_i'\geq 0}} \frac{\sum_i f_i' \Lambda_i^*(\phi_i')}{\max_i(\lambda_i - \phi_i' f_i') - \delta}.
$$

The arbitrary choice of $\zeta > 0$ and $\delta > 0$ implies that

$$-\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}_0^\pi \left[\max_i u_i^{(n)}(t_n) \geq 1 \right] \leq \sup_{\substack{\sum_i f_i'=1 \\ f_i' \geq 0}} \frac{\sum_i f_i' \Lambda_i^*(\phi_i')}{\max_i(\lambda_i - \phi_i' f_i')}.$$

The stationary distribution $\mathbb{P}^\pi$ induced by the (stabilizing) scheduling policy π forces $\pi((0, 0, \ldots, 0)) > 0$, so (27) finally implies

$$-\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}^\pi \left[||q^{(n)}(0)||_\infty \geq 1 \right] \leq \sup_{\substack{\sum_i f_i'=1 \\ f_i' \geq 0}} \frac{\sum_i f_i' \Lambda_i^*(\phi_i')}{\max_i(\lambda_i - \phi_i' f_i')},$$

completing the proof of the theorem. □

D Proof of Theorem 5

Proof. Recall that J_* is the infimum

$$J_* \triangleq \inf_{\substack{T,(m^T,c^T,q^T) \\ 0\leq t\leq T}} \frac{\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)} \right)}{\frac{d}{dt}||q(t)||_\infty} \tag{35}$$

over all feasible Fluid Sample Paths at regular points t. There is nothing to be done if the right hand side above is ∞, so we exclude this case. We have the following characterization of regular points under the Max-Queue scheduling algorithm.

Lemma 8. *Under the Max-Queue policy, let* $s(t) \triangleq \arg\max_{i=1,\ldots,N} q_i(t) \subseteq \{1, \ldots, N\}$. *If t is a regular point, then*

1. $c_i'(t) = 0 \ \forall i \notin s(t)$, *i.e., the non-maximum fluid queues do not receive service,*
2. $\frac{d}{dt}||q(t)||_\infty = \lambda_i - m_i'(t) \ \forall i \in s(t)$, *i.e., all the maximum fluid queues grow at the same rate.*

Thus, by Lemma 8,

$$J_* \geq \inf_{S\subseteq\{1,\ldots,N\}} \frac{\sum_{i\in S} c_i' \Lambda_i^* (\phi_i')}{w'}, \tag{36}$$

for all non-negative $\{c_i'\}_{i\in S}$, $\{\phi_i'\}_{i\in S}$ satisfying $\sum_{i\in S} c_i' = 1$, and $w' = \lambda_i - c_i'\phi_i' \ \forall i \in S$. Note that the denominator w' is strictly positive if and only if $\lambda \notin \mathcal{C}(\phi_1', \ldots, \phi_N')$, and that each ϕ_i' can be restricted to be at most $\mathbb{E}[R_i]$ (since if $\phi_i' > \mathbb{E}[R_i]$, reducing ϕ_i' to $\mathbb{E}[R_i]$ only gives a lesser fraction above).

For a subset $S \subseteq \{1, \ldots, N\}$, let

$$\mathcal{D}_S \triangleq \left\{ (\phi_i')_{i\in S} : R_{\min,i} \leq \phi_i' \leq \mathbb{E}[R_i], \ \exists c_i' \geq 0 \text{ with } \sum_{i\in S} c_i' = 1, \ \forall i,j \in S \ \lambda_i - c_i'\phi_i' = \lambda_j - c_j'\phi_j' = w' > 0 \right\}.$$

It can be shown that for each such tuple $\phi' \in \mathcal{D}_S$, there is a *unique* corresponding tuple c' and hence a *unique* w' (this follows because the point $(c'_i \phi'_i)_{i \in S}$ must be the point where the diagonal from ϕ' pierces the convex hull $\mathcal{C}((\phi'_i)_{i \in S})$). Thus, if we define a map $f^S : \mathcal{D}_S \to \mathbb{R}^+$ by

$$f^S(\phi') \triangleq \frac{\sum_{i \in S} c'_i \Lambda^*_i (\phi'_i)}{w'}, \tag{37}$$

then (36) is just

$$J_* \geq \min_S \inf_{\phi'_S \in \mathcal{D}_S} f^S(\phi'_S). \tag{38}$$

The next lemma contains the key result needed to prove Theorem 5:

Lemma 9. *Let $S \subseteq \{1, \ldots, N\}$ be such that $\mathcal{D}_S \neq \emptyset$. Then,*

1. *f^S attains its infimum over $\mathcal{D}_S$ at a point $\hat{\phi}'_S \in \mathcal{D}_S$.*
2. *Let $\hat{c}'$ be the (unique) tuple corresponding to $\hat{\phi}'_S$ such that $0 < \lambda_i - \hat{c}'_i \hat{\phi}'_i = \lambda_j - \hat{c}'_j \hat{\phi}'_j \ \forall i, j \in S$. Then, for every $\{c'_i\}_{i \in S}$ with $c'_i \geq 0$, $\sum_{i \in S} c'_i = 1$, we have*

$$f^S(\hat{\phi}'_S) \geq \frac{\sum_{i \in S} c'_i \Lambda^*_i(\hat{\phi}'_i)}{\max_{i \in S}(\lambda_i - c'_i \hat{\phi}'_i)}.$$

Proof. Without loss of generality, we will assume $S = \{1, \ldots, N\}$. Denote $\mu_i \triangleq \mathbb{E}[R_i]$. λ is a stabilizable vector of arrival rates, so $\lambda \in \mathcal{C}(\mu_1, \ldots, \mu_N)$ (here μ_i is overloaded to denote the N-tuple with the i-th coordinate being μ_i and the remaining coordinates being 0). Hence, there exists $\delta > 0$ such that $\sum_{i=1}^N \frac{\lambda_i}{\mu_i} = 1 - \delta$.

For any $\phi' \in \mathcal{D}_S$, we have $\sum_{i=1}^N \frac{\lambda_i}{\phi'_i} > 1$ by definition. Thus, $\sum_{i=1}^N \frac{\lambda_i}{\phi'_i} > \left(\frac{1}{1-\delta}\right) \sum_{i=1}^N \frac{\lambda_i}{\mu_i}$, so $\phi'_j < (1 - \delta)\mu_j$ for at least one j. It follows from the properties of the Cramér rate function for finite alphabets [18] that for each i, $\Lambda^*_i(\cdot)$ is strictly decreasing on $[R_{\min,i}, \mu_i]$, with $\Lambda^*_i(\mu_i) = 0$. Denote by γ the positive number $\min_i \Lambda^*_i((1 - \delta)\mu_i)$. Fix $\epsilon > 0$ small enough. If additionally (for $\phi' \in \mathcal{D}_S$) $w' < \epsilon$, then

$$f^S(\phi') = \frac{\sum_i c'_i \Lambda^*_i(\phi'_i)}{w'} \geq \frac{c'_j \Lambda^*_j(\phi'_j)}{w'} = \left(\frac{\lambda_j - w'}{\phi'_j w'}\right) \Lambda^*_j(\phi'_j) > \left(\frac{\lambda_j - \epsilon}{\mu_j \epsilon}\right) \gamma.$$

This means that for every $B > 0$, there exists $\epsilon_B > 0$ such that $\{\phi' \in \mathcal{D}_S : w' < \epsilon_B\} \subseteq \{\phi' \in \mathcal{D}_S : f^S(\phi') > B\}$. Thus,

$$\inf_{\phi' \in \mathcal{D}_S} f^S(\phi') = \inf_{\substack{\phi' \in \mathcal{D}_S: \\ w' \geq \epsilon_B}} f^S(\phi').$$

Observe that $\{\phi' \in \mathcal{D}_S : w' \geq \epsilon_B\}$ is a compact set, and that the lower-semicontinuity of $\Lambda^*_i(\cdot)$ [18] forces f^S to be lower-semicontinous on this compact set. It follows that f^S achieves its infimum on this set and thus on $\mathcal{D}_S$. This proves the first part of the lemma.

Turning to the second part, let $\hat{\phi}'_S \in \mathcal{D}_S$ infimize $f^S(\cdot)$ over $\mathcal{D}_S$, with $R_{\min,i} \leq \hat{\phi}'_i \leq \mu_i \ \forall i \in S$. Fix any $i \in S$. Since $\hat{\phi}'_S$ is a minimizer, increasing $\phi'_i = \hat{\phi}'_i$ by a small amount (keeping the other coordinates unchanged and ϕ'_S within $\mathcal{D}_S$) cannot decrease $f^S(\phi'_S)$, i.e., $\frac{\partial}{\partial \phi'_i} f^S(\phi'_S)\Big|_{\hat{\phi}'_S} \geq 0$. From the definition of f^S (37), we can write $\frac{\partial}{\partial \phi'_i} f^S(\phi'_S) = \frac{\partial}{\partial \phi'_i} \frac{N}{D}$, where $N \equiv N(\phi'_S) = \sum_{i \in S} c'_i \Lambda^*_i(\phi'_i)$, and $D \equiv D(\phi'_S) = w' \equiv w'(\phi'_S)$. Thus,

$$0 \leq \frac{\partial}{\partial \phi'_i} f^S(\phi'_S)\Big|_{\hat{\phi}'_S} = \frac{1}{D^2(\hat{\phi}'_S)} \left(D(\hat{\phi}'_S) \frac{\partial}{\partial \phi'_i} N(\phi'_S) - N(\hat{\phi}'_S) \frac{\partial}{\partial \phi'_i} D(\phi'_S) \right)\Big|_{\hat{\phi}'_S}. \tag{39}$$

Define, for each i, $\eta'_i \triangleq \frac{\Lambda^*_i(\phi'_i)}{\phi'_i}$ (and $\hat{\eta}'_i \triangleq \frac{\Lambda^*_i(\hat{\phi}'_i)}{\hat{\phi}'_i}$). Noticing that $\frac{\partial}{\partial \phi'_i} D(\phi'_S) = \frac{\partial}{\partial \phi'_i} (\lambda_j - c'_j \phi'_j) = -\frac{\partial}{\partial \phi'_i} (c'_j \phi'_j)$ for all $j \in S$, we can write

$$\frac{\partial}{\partial \phi'_i} N(\phi'_S) = \frac{\partial}{\partial \phi'_i} \sum_{j \in S} c'_j \phi'_j \eta'_j = -\frac{\partial D(\phi'_S)}{\partial \phi'_i} \cdot \sum_{j \in S} \eta'_j + c'_i \phi'_i \cdot \frac{\partial \eta'_i}{\partial \phi'_i}.$$

Along with (39), this implies (evaluated at $\phi_S = \hat{\phi}'_S$)

$$\begin{aligned}
0 &\leq -D(\phi'_S) \cdot \frac{\partial D(\phi'_S)}{\partial \phi'_i} \cdot \sum_{j \in S} \eta'_j + D(\phi'_S) \cdot c'_i \phi'_i \cdot \frac{\partial \eta'_i}{\partial \phi'_i} - N(\phi'_S) \frac{\partial D(\phi'_S)}{\partial \phi'_i} \\
&= -\frac{\partial D(\phi'_S)}{\partial \phi'_i} \left[D(\phi'_S) \cdot \sum_{j \in S} \eta'_j + N(\phi'_S) \right] + D(\phi'_S) \cdot c'_i \phi'_i \cdot \frac{\partial \eta'_i}{\partial \phi'_i} \\
&= -\frac{\partial D(\phi'_S)}{\partial \phi'_i} \left[D(\phi'_S) \cdot \sum_{j \in S} \eta'_j + N(\phi'_S) \right] + D(\phi'_S) \cdot c'_i \phi'_i \cdot \frac{\phi'_i \frac{\partial \Lambda^*_i(\phi'_i)}{\partial \phi'_i} - \Lambda^*_i(\phi'_i)}{\phi'^2_i} \\
&= -\frac{\partial D(\phi'_S)}{\partial \phi'_i} \left[D(\phi'_S) \cdot \sum_{j \in S} \eta'_j + N(\phi'_S) \right] - D(\phi'_S) c'_i \eta'_i + D(\phi'_S) \cdot c'_i \cdot \underbrace{\frac{\partial \Lambda^*_i(\phi'_i)}{\partial \phi'_i}}_{\leq 0} \\
&\leq -\underbrace{\frac{\partial D(\phi'_S)}{\partial \phi'_i}}_{\leq 0} \left[D(\phi'_S) \cdot \sum_{j \in S} \eta'_j + N(\phi'_S) \right] - D(\phi'_S) c'_i \eta'_i \\
\Rightarrow \frac{N}{D} &\geq -\frac{c'_i \eta'_i}{\left(\frac{\partial D}{\partial \phi'_i} \right)} - \sum_{j \in S} \eta'_j.
\end{aligned} \tag{40}$$

Since $D = \lambda_j - c'_j \phi'_j$ and $\sum_{j \in S} c'_j = 1$, we have

$$\sum_{j \in S} \frac{\lambda_j - D}{\phi'_j} = 1 \;\Rightarrow\; D = \frac{\sum_{j \in S} \frac{\lambda_j}{\phi'_j} - 1}{\sum_{j \in S} \frac{1}{\phi'_j}}.$$

Using this, some calculus yields

$$
\begin{aligned}
-\frac{c_i'}{\left(\frac{\partial D}{\partial \phi_i'}\right)} &= \phi_i' \cdot \sum_{j\in S} \frac{1}{\phi_j'} \\
\Rightarrow\ \frac{N}{D} &\geq \eta_i'\phi_i' \cdot \sum_{j\in S} \frac{1}{\phi_j'} - \sum_{j\in S} \eta_j' \quad \text{(by (40))} \\
\Rightarrow\ f^S(\hat{\phi}_S') = \frac{N}{D} &\geq \left(\max_{i\in S} \hat{\eta}_i'\hat{\phi}_i'\right) \cdot \sum_{j\in S} \frac{1}{\hat{\phi}_j'} - \sum_{j\in S} \hat{\eta}_j'.
\end{aligned} \tag{41}
$$

Now consider any tuple $\{d_i'\}_{i\in S}$ with $d_i' \geq 0$ and $\sum_{i\in S} d_i' = 1$. Let $\delta_i' \triangleq d_i' - \hat{c}_i'$ for all $i \in S$, so that $\sum_{i\in S} \delta_i' = 0$, and for $t \in [0,1]$, define

$$
g(t) \triangleq \frac{\sum_{i\in S}(\hat{c}_i' + t\delta_i')\Lambda_i^*(\hat{\phi}_i')}{\max_{i\in S}(\lambda_i - (\hat{c}_i' + t\delta_i')\hat{\phi}_i')},
$$

so that $g(0) = f^S(\hat{\phi}_S')$. To prove the second part of the lemma, we proceed to show that $g(0) \geq g(1)$. First, note that since (for $t = 0$) $\lambda_i - \hat{c}_i'\hat{\phi}_i'$ is equal for all $i \in S$, we can assume without loss of generality that $1 \in S$ and that the denominator in the definition of $g(t)$ above is equal to $\lambda_1 - \hat{c}_1'\hat{\phi}_1' - t\delta_1'\hat{\phi}_1' = D(\hat{\phi}_S') - t\delta_1'\hat{\phi}_1'$, with $\delta_1'\hat{\phi}_1' \leq \delta_i'\hat{\phi}_i'$ for each $i \in S$. This makes $g(\cdot)$ a quotient of affine functions on $[0,1]$, and thus monotone. It just remains to show that $g'(t) \leq 0$ for all t.

Consider

$$
\begin{aligned}
\frac{d}{dt}g(t) = \frac{d}{dt}\left(\frac{N + t\sum_{i\in S}\delta_i'\Lambda^*(\hat{\phi}_i')}{D - t\delta_1'\hat{\phi}_1'}\right) &\leq 0 \\
\Leftrightarrow \quad D \cdot \sum_{i\in S} \delta_i'\Lambda^*(\hat{\phi}_i') + N \cdot \delta_1'\hat{\phi}_1' &\leq 0 \\
\Leftrightarrow \quad \frac{\sum_{i\in S}\delta_i'\Lambda^*(\hat{\phi}_i')}{\underbrace{-\delta_1'\hat{\phi}_1'}_{>0}} &\leq \frac{N}{D} \\
\Leftrightarrow \quad \sum_{i\in S}\left(\frac{\delta_i'\hat{\phi}_i'}{-\delta_1'\hat{\phi}_1'}\right)\hat{\eta}_i' &\leq \frac{N}{D}.
\end{aligned}
$$

By (41), we will be done if we can show that

$$
\left(\max_{j\in S} \hat{\eta}_j'\hat{\phi}_j'\right) \cdot \sum_{j\in S} \frac{1}{\hat{\phi}_j'} - \sum_{j\in S} \hat{\eta}_j' \geq \sum_{j\in S}\left(\frac{\delta_j'\hat{\phi}_j'}{-\delta_1'\hat{\phi}_1'}\right)\hat{\eta}_j'.
$$

But notice that

$$\sum_{j\in S}\left(\frac{\delta'_j\hat{\phi}'_j}{-\delta'_1\hat{\phi}'_1}\right)\hat{\eta}'_j+\sum_{j\in S}\hat{\eta}'_j=\sum_{j\in S}\hat{\eta}'_j\left[1+\frac{\delta'_j\hat{\phi}'_j}{-\delta'_1\hat{\phi}'_1}\right]$$
$$\le\left(\max_{j\in S}\hat{\eta}'_j\hat{\phi}'_j\right)\sum_{j\in S}\frac{1}{\hat{\phi}'_j}\left[1+\frac{\delta'_j\hat{\phi}'_j}{-\delta'_1\hat{\phi}'_1}\right]$$
$$=\left(\max_{j\in S}\hat{\eta}'_j\hat{\phi}'_j\right)\sum_{j\in S}\frac{1}{\hat{\phi}'_j}+\left(\max_{j\in S}\hat{\eta}'_j\hat{\phi}'_j\right)\underbrace{\sum_{j\in S}\frac{\delta'_j}{-\delta'_1\hat{\phi}'_1}}_{=0}$$
$$=\left(\max_{j\in S}\hat{\eta}'_j\hat{\phi}'_j\right)\sum_{j\in S}\frac{1}{\hat{\phi}'_j}.$$

This completes the proof of the lemma. □

Using this lemma, we can finish the proof of the theorem. Let S be a subset of channels that achieves the minimum in (38); according to the lemma there exists $\hat{\phi}'_S$ that infimizes f^S over $\mathcal{D}_S$. Extend $\hat{\phi}'_S \in \mathbb{R}^{|S|}$ to an N-tuple $\hat{\phi}' \in \mathbb{R}^N$ by setting coordinates $i \notin S$ to their respective mean channel rates $\mathbb{E}[R_i]$. This means that $\Lambda^*_i(\hat{\phi}'_i) = 0$ for $i \notin S$, so for any N-tuple e' on the simplex, because $\sum_{i\in S} e'_i \le 1$, the lemma gives

$$\frac{\sum_{i=1}^N e'_i\Lambda^*_i(\hat{\phi}'_i)}{\max_{1\le i\le N}(\lambda_i-e'_i\hat{\phi}'_i)}=\frac{\sum_{i\in S} e'_i\Lambda^*_i(\hat{\phi}'_i)}{\max_{1\le i\le N}(\lambda_i-e'_i\hat{\phi}'_i)}\le\frac{\sum_{i\in S} e'_i\Lambda^*_i(\hat{\phi}'_i)}{\max_{i\in S}(\lambda_i-e'_i\hat{\phi}'_i)}\le f^S(\hat{\phi}'_S)\le J_*,$$

completing the proof. □

References

[1] Motorola, "Long term evolution (LTE): Overview of LTE air-interface," *White paper*, 2007.

[2] A. Gopalan, C. Caramanis, and S. Shakkottai, "On wireless scheduling with partial channel state information," in *Proc. Allerton Conference on Communication, Control and Computing*, 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 Informational Sciences*, vol. 14, pp. 191–217, 2004.

[5] 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.

[6] 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.

[7] A. Stolyar, "Large deviations of queues sharing a randomly time-varying server," *Queueing Systems*, pp. 1–35, 2008.

[8] 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.

[9] 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.

[10] 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.

[11] A. Stolyar and K. Ramanan, "Largest weighted delay first scheduling: Large deviations and optimality," *Annals of Applied Probability*, pp. 1–48, 2001.

[12] 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.

[13] S. Shakkottai, "Effective capacity and QoS for wireless scheduling," *IEEE Transactions on Automatic Control*, February 2008.

[14] 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/paper/it09-ldp-updated2.pdf.

[15] 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.

[16] 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.

[17] 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.

[18] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications.* Jones and Bartlett, 1993.

[19] A. Gopalan, C. Caramanis, and S. Shakkottai, "On wireless scheduling with partial channel state information," *To appear in the IEEE Trans. on Information Theory*, 2011. Previous version in the Proc. of the Allerton Conference on Communication, Control and Computing, 2007.

[20] 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.

[21] N. Shimkin, "Extremal large deviations in controlled i.i.d. processes with applications to hypothesis testing," *Advances in Applied Probability*, vol. 25, no. 4, pp. pp. 875–894, 1993.

[22] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis.* New York, NY, USA: Cambridge University Press, 2005.