

Wireless Scheduling with Partial Channel State Information: Large Deviations and Optimality

Aditya Gopalan · Constantine
Caramanis · Sanjay Shakkottai

Received: date / Accepted: date

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 sub-

This work was partially supported by NSF grants CNS-0721380, CNS-1017549, CNS-1161868, CNS-0721532, and EFRI-0735905. A preliminary version of this work appears in the proceedings of the 31st IEEE International Conference on Computer Communications (IEEE INFOCOM), Orlando, March 2012 [9].

Aditya Gopalan
Department of Electrical Communication Engineering
Indian Institute of Science
Bangalore 560012, India.
E-mail: aditya@ece.iisc.ernet.in

Constantine Caramanis
Department of Electrical and Computer Engineering
The University of Texas at Austin
Austin, Texas 78712, USA.
E-mail: constantine@utexas.edu

Sanjay Shakkottai
Department of Electrical and Computer Engineering
The University of Texas at Austin
Austin, Texas 78712, USA.
E-mail: shakkott@austin.utexas.edu

sequently 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 that base their decision on the current (or any finite past history of) system state, of the tail of the longest queue. To accomplish this, we employ novel analytical techniques for studying the performance of scheduling algorithms using partial state, which may be of independent interest. These include new sample-path large deviations results for processes obtained by non-random, predictable sampling of sequences of independent and identically distributed random variables. A consequence of these results is that scheduling with partial state information yields a rate function significantly different from scheduling with full channel information. In the special case when the observable subsets are singleton flows, i.e., when there is effectively no *a priori* channel-state information, Max-Exp reduces to simply serving the flow with the longest queue; thus, our results show that to always serve the longest queue in the absence of any channel-state information is large-deviations optimal.

Keywords Wireless scheduling · Large deviations

1 Introduction

Next-generation wireless cellular systems such as LTE-Advanced [15] and WiMAX [1] promise high-speed packet-switched data services for a variety of applications, including file transfer, peer-to-peer sharing and real-time audio/video streaming. This demands effective scheduling in typical wireless environments with time-varying channels and limited resources, to guarantee high data rates to the users. Together with maximizing data rates or throughput, the scheduling algorithm at the cellular base station must keep packet delays in the system low, in order to support highly delay-sensitive applications like real-time video streaming.

There has been much recent work to develop wireless scheduling algorithms with optimal throughput and/or delay performance [28, 2, 25, 17, 30]. Such opportunistic scheduling algorithms utilize instantaneous wireless Channel State Information (CSI) from all users to make good scheduling decisions. However, in a practical situation with a large number of users in the network, channel state feedback resources could potentially be limited, i.e., it might be infeasible to acquire complete instantaneous CSI from all channels due to bandwidth and latency limitations. Instead, it might be possible to request CSI feedback from only a *subset of users* each time. Thus, it is important to develop algorithms that can schedule using only partial CSI rather than complete CSI, and at the same time afford the best possible delay performance.

Using partial CSI – from subsets of channels – entails a new dimension of opportunism in wireless scheduling. The scheduling algorithm needs to make a careful choice of which subsets to sample, together with how to use the sampled CSI for scheduling. Recently, natural extensions of complete-CSI scheduling algorithms to the partial-CSI setting have shown to have throughput-optimal

properties [10], yet it is not clear how they perform in the sense of packet delays. The general structure of low-delay, partial-CSI scheduling algorithms remains unknown, i.e., how an algorithm should choose “good” subsets of channels, whether any additional backlog or statistical information is needed for picking subsets, and if so, how much, how users should be scheduled in the observed subset etc.

In this work, we develop algorithms for wireless scheduling that use only partial CSI, i.e., from subsets of channels, and that also enjoy high performance guarantees. We consider a wireless downlink where a base station schedules users using partial CSI from subsets of channels. Viewing the system queue lengths as a surrogate for packet delays, we seek scheduling strategies that can keep the longest queue in the system as short as possible, i.e., minimize the likelihood of overflow of the longest queue. We design a new scheduling algorithm, that we term *Max-Exp*, that obtains partial CSI relying on just current queue lengths and no other auxiliary information. Employing sample-path large deviations techniques, we show that when the observable channel subsets are disjoint, Max-Exp yields the best decay rate for the longest-queue overflow probability, across all scheduling strategies which use subset-based CSI to schedule users. To the best of our knowledge, this is the first work that analyzes queue-overflow performance for scheduling with the information structure of partial CSI, and that provides a simple scheduling algorithm needing no extra statistical information which is actually rate-function optimal for buffer overflow.

From a technical standpoint, sample-path large deviations techniques have successfully been used to analyze wireless scheduling algorithms [3, 25, 17, 30]; yet, significant new analytical challenges emerge when studying the large deviations behavior of scheduling strategies that cannot access the full state of the system. A chief difference in this regard arises from the fact that when scheduling is carried out by observing the *complete* state/randomness of the system, large deviations occur depending on how the scheduler responds to atypical channel state behavior. In other words, a natural cause-effect relationship between the channel state process and scheduling actions is the basis for the analysis of large deviations performance. On the other hand, when *partial* channel state is *acquired* selectively by a scheduling algorithm, this cause-effect sequence is reversed – it is the algorithm that first decides what part of the channel state to sample; subsequently, this dynamic portion of the channel state can respond by behaving atypically. Viewed differently, the scheduling information structure no longer falls into an “experts” setting (all channel rates known in advance) but rather into a “bandit” setting (only chosen channel rates known) [12], implying a fundamental change in the large deviations dynamics. Indeed, we are able to show that this difference results in a significantly different rate function than that encountered in the former complete-CSI case.

Also, the standard approach of analyzing queue overflow probability exponents using continuity of queue-length/delays as functions of the arrivals and channel processes [18, 31] becomes cumbersome due to the complex two-stage

sampling and scheduling structure of scheduling with partial CSI. Thus, we are led to develop new sample-path large deviations results for processes with dynamically (and predictably) sampled randomness, which help to bound the resulting rate functions via connections to appropriate variational problems. We believe that these techniques and results are of independent interest as tools to analyze the behavior of scheduling policies that can only sample parts of the system state.

1.1 Related Work

For scheduling with complete CSI, there is a rich body of work on throughput-optimal scheduling algorithms, starting from the pioneering approach of Tassiulas et al. [28] to develop the Backpressure algorithm. A host of scheduling algorithms such as Max-Weight/Backpressure [28, 2], the Exponential rule [21, 20, 25] and the Log rule [17] have been developed for scheduling using full CSI. Many optimality results are now known for the delay/queue-length performance of the above full-CSI algorithms. These include expected queue length/delay bounds via Lyapunov function techniques [14, 8], tail probability decay rates for queue lengths [26, 29, 17, 18, 30, 25, 31], heavy-traffic optimality [24] etc.

Throughput-maximizing scheduling has been studied with different forms of partial CSI, including infrequent channel state measurements [11], group/random-access based quantized channel state feedback [16, 27], optimal channel state probing with costs [5, 4], delayed CSI [32] and subset-based CSI [10]. However, to date, neither the structure nor performance results for queue overflow tails under scheduling with partial CSI are known.

1.2 Contributions

We describe a new scheduling algorithm – Max-Exp – for scheduling over a wireless downlink when Channel State Information (CSI) is restricted to a collection of *observable channel subsets*. Max-Exp picks a subset of channels to observe their states, depending on an appropriate exponentiated sum of the subset queue lengths. Having done that, it uses the well-known Exponential rule [21] to schedule a user from the subset using the obtained CSI. Thus, Max-Exp *does not* need any additional information (e.g. traffic/channel statistics) other than queue lengths to dynamically pick subsets, and only the instantaneous subset channel states to schedule users.

Our main contributions can be summarized as follows:

1. We derive a lower bound on the rate function for overflow of the longest queue under the Max-Exp scheduling algorithm, using sample-path large deviations tools and their connection to variational optimal-control problems. A key technical contribution here is developing large deviations properties for processes obtained by predictably sampling independent and

identically distributed (iid) sequences. These results help to show that the sample-path large deviations rate function, for algorithms that sample portions of the channel state, not only depends on the standard Cramér empirical rate functions of the sampled portions, but also relies crucially on the *sampling frequencies* of the portions.

Conversely, we also show *universal* (i.e., over *all* scheduling algorithms that use partial, subset-based CSI) upper bounds on the rate function of queue overflow. Here again, a technical challenge arises due to the fact that for an arbitrary¹ scheduling algorithm, the large-deviations “cost” of buffer overflow depends crucially on its subset *sampling behavior* – 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.

2. In the case where the collection of observable subsets available to the scheduler is disjoint, we prove that the lower bound on the large deviations buffer overflow rate function for Max-Exp matches the uniform upper bound on the rate function over all algorithms. This not only characterizes the exact buffer overflow exponent of the Max-Exp algorithm, but also shows rather surprisingly that *the simple Max-Exp strategy yields the optimal overflow exponent across all scheduling rules using partial CSI*². As a side consequence, this shows that for scheduling with singleton subsets of users, merely scheduling the user with the longest queue at each time slot – a greedy strategy when no CSI is available beforehand – is large-deviations rate function-optimal.

Technically, showing that the lower and upper bounds for the queue overflow rate function match involves solving a complex and non-convex variational problem arising from the rate function for predictably sampled random processes, and is another contribution of this work.

2 System Model

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

¹ In the context of this work, an arbitrary scheduling algorithm is to be understood as any map that is based on the current (or any finite past history) of system state.

² By optimal, we mean optimal among all scheduling algorithms that base their decision on the current (or any finite past history of) system state.

1. **Arrivals:** An integer number of data packets $A_i(k)$ arrives to user i , $i = 1, \dots, N$. Packets get queued at their respective users if they are not immediately transmitted.
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*. At time slot k , we denote the instantaneous service rates by $(R_1(k), \dots, R_N(k))$.
3. **Scheduling:** One user $U(k) \in \{1, \dots, N\}$ is picked for service, and a number of packets not exceeding its instantaneous service rate is removed from its queue. Let $D_i(k)$ denote whether user i is scheduled in time slot k ($D_i(k) = 1$), or not ($D_i(k) = 0$). Then, user i 's queue length (denoted by $Q_i(\cdot)$) evolves as $Q_i(k+1) = [Q_i(k) + A_i(k) - D_i(k)R_i(k)]^+$, where $x^+ \equiv \max(x, 0)$.

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

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

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

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

- *Step 1:* Pick a *subset* $S(k)$ of the N channels, from a given collection \mathcal{O} of *observable subsets*³. This choice can depend on all random variables in time slots up to and including k except the channel state $R(k)$.
- *Step 2:* Once the subset $S(k)$ of channels is chosen, the instantaneous service rates $(R_i(k))_{i \in S(k)}$ are revealed/available to the scheduling algorithm, and it chooses a user $U(k) \in S(k)$ for service, possibly depending on these service rates.

Note that at each time, the channel state information available to the scheduling algorithm is restricted to the chosen subset $S(k)$ of channels, as opposed to the full CSI case where all the service rates $(R_i(k))_{i=1}^N$ are available. For further detailed discussion about this scheduling model and how it abstracts limited channel state information at the wireless physical layer, etc., we refer the reader to [10].

³ The collection of observable subsets models the collection of subsets of channels for which the wireless scheduler can obtain instantaneous channel state information, as described in the Introduction (1).

3 Objective, Algorithms and Main Results

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

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

(when the limit exists), for scheduling algorithms that *observe only partial channel state* while scheduling. Note that for large n , $\mathbb{P}[\|Q(k)\|_\infty \geq n] \approx e^{-nI}$, so maximizing the exponent I gives smaller overflow probabilities. Also, it is well-known that packet delays are closely related to queue lengths [30], which justifies using I as our performance objective.

With this objective in mind, we introduce a new scheduling algorithm *Max-Exp* (Algorithm 1). The algorithm may be interpreted as locally (in Step 2) using the Exponential scheduling rule [19, 22, 20, 21], and globally (in Step 1) using the Exponential rule metric without the (observed) instantaneous rate to pick a subset of channels.

It is well-known that in case the entire set of channels is observable (i.e., the full-information setting), the Exponential rule maximizes the exponent of the queue overflow probability [25], hence it is a natural candidate for the in-subset scheduling rule in Step 2 of the Max-Exp algorithm. The rule used to choose subsets in Step 1 is chosen so as to match the in-subset Exponential rule, and guarantees properties that are required in the fluid limit scaled description of the dynamics in order to show our main result.

Algorithm 1 Max-Exp

At each time slot k , breaking ties arbitrarily,

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

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

is maximized (here $\bar{Q}(k) \triangleq \frac{1}{N} \sum_{i=1}^N Q_i(k)$ is the length of the average queue at time slot k).

2. Schedule a user $i \in S(k)$ such that $R_i(k) \exp \left(\frac{Q_i(k)}{1 + \sqrt{\bar{Q}(k)}} \right)$ is maximized (the Exponential rule [21]).
-

By our probabilistic assumptions on the channel state process, Max-Exp makes the vector process of queue lengths at each time a discrete-time Markov chain. Following standard convention [2, 14, 8], we term the set of arrival rates

$\lambda \equiv (\lambda_i)_{i=1}^N$ for which this Markov chain is positive-recurrent as the *throughput region* of Max-Exp. To not deviate from the main focus of this work, we state that when the observable subsets are disjoint, the throughput region of Max-Exp contains that of any other scheduling algorithm, i.e., Max-Exp is *throughput-optimal*⁴. The proof of throughput-optimality is analogous to that of the Max-Sum-Queue scheduling algorithm [10]. Our main result states that Max-Exp yields the best (exponential) rate of decay of the tail of the longest queue over all strategies that use partial CSI from disjoint subsets:

Theorem 1 (Large Deviations Optimality of Max-Exp) *Let the system's arrival rates λ lie in the interior of the throughput region of the Max-Exp scheduling algorithm. There exists $J_* > 0$ such that the following holds.*

1. Let \mathbb{P} denote the stationary probability distribution that the Max-Exp algorithm induces on the vector of queue lengths. Then,

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

2. Let π be an arbitrary scheduling rule⁵ that induces a stationary distribution \mathbb{P}^π on the vector of queue lengths. If the system of observable subsets \mathcal{O} is disjoint, then

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

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

Theorem 1 highlights the striking property that Max-Exp, using only current queue length information to sample channel subsets and the Exponential rule to schedule a sampled channel, yields the *fastest decay* of the buffer overflow probability across the whole spectrum of partial-CSI scheduling algorithms – including those that potentially use additional statistical information, traffic characteristics etc. The crucial scheduling step in Max-Exp is Step 1, which essentially samples the “right” channel subset depending on queue lengths. The result shows that queue length feedback is sufficient to guarantee good delay performance, provided suitable subsets of channel states are sampled as with the Max-Exp scheduling algorithm. We remark that the optimality of Max-Exp continues to hold even when all queue lengths are delayed by any bounded amount. We also remark that the restriction to disjoint observable subsets is necessary since otherwise, even throughput-optimality of Max-Exp-style scheduling rules does not hold [10]. This is briefly because the geometry of the throughput region is fundamentally different when subsets are

⁴ We mean throughput-optimal among all scheduling algorithms that base their decision on the current (or any finite past history of) system state.

⁵ An arbitrary scheduling rule is any map that is based on the current (or any finite past history) of system state.

disjoint, and its properties play a key role in the optimality proof for Max-Exp here.

En route to proving Theorem 1, we develop lower bounds for the large deviations exponents of partially and deterministically sampled iid processes, that are of independent interest. This results in a new rate function formulation in terms of variational optimization, that differs significantly from existing rate functions [29, 17, 18, 30, 25, 31] by explicitly incorporating partial channel state sampling behavior. Standard optimal control approaches for the full-CSI case cannot be applied to analyze partial-CSI scheduling algorithms – since only a portion of the channel state is revealed to the scheduler, the channel state process can cause large deviations by behaving atypically just in the revealed portion, and not jointly as a whole.

A related challenge arises in the process of finding universal upper bounds on the decay rate for arbitrary partial-CSI scheduling policies⁶. Recent large-deviations work in full-CSI scheduling [30, 25] accomplishes this by calculating the “cost” of universal channel-state sample paths that cause buffer overflow under any scheduling algorithm; however, this procedure fails for algorithms actively sampling the channel state, since the cost of such sample paths intimately depends on the subset sampling behavior. To overcome this, we use a martingale-based argument in a novel way with the standard exponential tilting method to prove universal upper bounds on the exponent.

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

Algorithm 2 Max-Queue

At each time slot k , breaking ties arbitrarily,

1. Schedule a user i such that $Q_i(k)$ is maximized.
-

Thus, an immediate corollary of Theorem 1 is the following optimality result for Max-Queue when the observable user subsets are restricted to singletons, i.e., when there is effectively no CSI to use in scheduling:

Corollary 1 (Large Deviations Optimality of Max-Queue for singleton observable subsets) *If the system’s arrival rates λ lie in the interior of the throughput region of the Max-Queue scheduling algorithm, then Max-Queue has the optimal⁷ large-deviations exponent of the queue overflow probability over all stabilizing scheduling policies that can sample only individual channel states.*

Road map to prove Theorem 1: Though Theorem 1 for Max-Exp is our chief result, we prove it by first establishing the optimality result for Max-Queue

⁶ In the context of this work, an arbitrary scheduling policy is to be understood as any map that is based on the current (or any finite past history) of system state.

⁷ By optimal, we mean optimal among all scheduling algorithms that base their decision on the current (or any finite past history of) system state.

(Corollary 1), and then extending the argument to the setting of general disjoint subsets. This is mainly because the essence of the optimality lies in the key subset selection step, and restricting attention to the case of singleton observable subsets allows us to concentrate on how subset selection influences the large deviations rate function of buffer overflow. Technically, another reason for this order of working is that Max-Queue can naturally be analyzed with the standard $O(n)$ fluid scaling, whereas showing the optimality property for Max-Exp requires using a more delicate fluid limit framework at the $O(\sqrt{n})$ “local” fluid time-scale [21, 25].

4 Preliminaries and Sample Path Large Deviations Framework

This section lays down preliminaries for the sample-path large deviations techniques we use to study overflow probabilities of wireless scheduling algorithms. Much of this framework is standard in large deviations analyses of wireless systems [30, 25, 17], but we include it for completeness.

Throughout this work, we denote by $(\Omega, \mathcal{F}, \mathbb{P})$ a common probability space that supports all defined random variables and processes. Fix an integer $T > 0$, and consider a sequence of (independent) queueing systems indexed by $n = 1, 2, \dots$, each with its own arrival and channel state processes, and evolving as described in Section 2. Henceforth, we explicitly reference by the superscript (n) any quantity associated with the n th system. For any (possibly vector-valued) random process $X^{(n)}(k)$, $k = 0, 1, 2, \dots$ in the n th system, let us define its scaled (by $1/n$), shifted and piecewise linear version $x^{(n)}(\cdot)$ on the interval $[0, T]$ as follows:

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

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

For the n th queueing system, with k a nonnegative integer, we define the following random processes central to our study of the evolution of the system:

- $F_i^{(n)}(k)$: The total number of packets to queue i that arrived by time slot k ,
- $\hat{F}_i^{(n)}(k)$: The total number of packets that were served from queue i by time slot k ,
- $C_\alpha^{(n)}(k)$: The total number of time slots before k when the observable subset α was chosen by the scheduling algorithm,
- **(Sub-state)** $R_\alpha^{(n)}(k)$: The vector of instantaneous service rates $R^{(n)}(k)$ restricted to the coordinates of α , i.e., $R_\alpha^{(n)}(k) = (R_i^{(n)}(k))_{i \in \alpha}$,

- $G_r^{\alpha, (n)}(k)$: The total number of time slots before time slot k when the subset α was picked and its sub-state was r ,
- $\hat{G}_{ri}^{\alpha, (n)}(k)$: 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,
- $Q_i^{(n)}(k)$: The length of queue i at time slot k , whose evolution is specified in Section 2,
- $M^{(n)}(k)$: The (vector-valued) partial sums process corresponding to the sampled rates $R^{(n)}(k)\delta_{S(k)}$, i.e., $M^{(n)}(k) \triangleq \sum_{j=0}^k R^{(n)}(j)\delta_{S(j)}$. (Here, δ_S denotes the indicator vector of the subset S .)

For right-continuous, non-decreasing functions $u : \mathbb{R} \rightarrow \mathbb{R}$ and $v : \mathbb{R} \rightarrow \mathbb{R}$, we overload notation and denote by u and v their respective induced Stieltjes measures on \mathbb{R} , whenever the context is understood. Furthermore, when $v \ll u$ (i.e., when dv is absolutely continuous wrt du), we denote by $\frac{dv}{du}$ the Radon-Nikodym derivative⁸ of v wrt u .

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

Note 1. Wherever a scheduling algorithm is being explicitly considered, we will use the term *valid FSP* to denote an FSP that occurs with positive probability under the scheduling algorithm.

Note 2. We use \dot{f} and f' interchangeably, in the paper, to denote the derivative of a (differentiable) function f .

5 Analysis: Singleton Subsets and Max-Queue

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

⁸ The Radon-Nikodym derivative $\frac{dv}{du}$ is uniquely defined du -a.e.

5.1 Lower-bound for Max-Queue's Decay Rate

Consider the queueing system operating under an arbitrary *nonrandom* scheduling algorithm, i.e., the algorithm's choice of a singleton user in the current time slot is a deterministic function of the entire history of observed users' indices and channel states, and does not depend on the unobserved channel states in the past⁹. Max-Queue with deterministic tie-breaking (e.g., pick the lowest-indexed queue when there are two or more longest queues) is an example of a nonrandom scheduling algorithm, since the current user chosen depends on accumulated queue lengths, which in turn depend directly on the channel rates obtained as a result of past scheduling choices.

The sequence of observed users and their channel states under a nonrandom scheduling algorithm is an outcome of sampling an iid vector-valued process (i.e., the full channel state) in a nonrandom and *predictable* (i.e., with sampling indices depending only on past observed history) manner. Our first key result (Proposition 1) essentially furnishes an upper bound for the deviation probability of the queue-length process (equivalently the cumulative process of observed channel states) in time slots $0, \dots, nT$, in terms of a novel sample-path large deviations rate function of the user selection and channel state paths.

Let us fix $T > 0$. For $q_0 \in \mathbb{R}^N$, let $\mathbb{P}_{q_0}^{n,T}$ be the probability measure of the n -th queueing system conditioned on starting the system at $Q^{(n)}(0) = nq_0$ (i.e. $q^{(n)}(0) = q_0$). If we denote by $\mathcal{C}_{\mathcal{L}}^+([0, T])$ the space of nonnegative \mathbb{R}^N -valued Lipschitz functions on $[0, T]$ equipped with the supremum norm, then we have:

Proposition 1 (Large Deviation Bound for a Finite Horizon) *Let Γ be a closed set of trajectories in $\mathcal{C}_{\mathcal{L}}^+([0, T])$. Then, under any nonrandom scheduling policy,*

$$\begin{aligned} & -\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_0^{n,T} \left[q^{(n)} \in \Gamma \right] \\ & \geq \inf_{(m^T, c^T, q^T)} \int_0^T \left[\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{dm_i}{dc_i}(t) \right) \right] dt \quad (1) \\ & \text{subject to } (m^T, c^T, q^T) \text{ a valid FSP,} \\ & \quad q^T(0) = 0, q^T \in \Gamma, \end{aligned}$$

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

Proposition 1 states that the “correct” sample-path large deviations rate function, for algorithms that can sample only singleton subsets of channels, is a combination of the standard rate functions Λ_i^* for the empirical means of

⁹ In formal terms, an arbitrary scheduling algorithm is to be understood as any map that is based on the current (or any finite past history) of system state.

individual channel rates *weighted by the corresponding channel selection frequencies* \dot{c}_i . Note the crucial dependence of the rate function on the subset selection process, captured by weighting Λ_i^* by \dot{c}_i in (1) – a significant departure from the rate function studied for the standard case of full channel state information where there is no pre-weighting by the algorithm-dependent factor \dot{c} [30, 25].

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

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), we now proceed to extend this result to the queue overflow rate function for the *stationary distribution* under Max-Queue. Recall that a unique stationary distribution exists since Max-Queue makes the irreducible and aperiodic system state Markov chain positive recurrent [10]. Intuitively, we expect that the finite horizon probability distribution $\mathbb{P}_{q_0}^{n,T}$ somehow “tends” to the stationary distribution \mathbb{P} . Thus, we show that minimizing the right hand side of (1) 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 earlier, using techniques from Friedlin-Wentzell large-deviations theory [30, 25]. 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.

Proposition 2 (Large Deviation Bound for the Stationary Distribution) *Let \mathbb{P} denote the stationary probability distribution of the system state under the Max-Queue scheduling algorithm. Then,*

$$\begin{aligned}
& -\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left[\|q^{(n)}(0)\|_\infty \geq 1 \right] \\
& \geq \inf_{T, (m^T, c^T, q^T)} \int_0^T \left[\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{dm_i}{dc_i}(t) \right) \right] dt \\
& \text{subject to } (m^T, c^T, q^T) \text{ a valid FSP in } [0, T], \\
& \quad q^T(0) = 0, \|q^T(T)\|_\infty \geq 1, \\
& \quad T \geq 0.
\end{aligned} \tag{2}$$

The reader is referred to Appendix B for the proof details.

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

$$\int_0^T \left[\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{dm_i}{dc_i}(t) \right) \right] dt \geq \inf_{t \in \mathcal{B}} \frac{\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{dm_i}{dc_i}(t) \right)}{\frac{d}{dt} \|q(t)\|_{\infty}},$$

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

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

with the infimum over all FSPs (m^T, c^T, q^T) feasible for (2), all *regular points* t , and all *finite horizons* T . This results in the following (weaker) lower bound on the rate function of Max-Queue’s stationary queue overflow probability:

Proposition 3 (Lower bound for Max-Queue’s Queue Overflow Rate Function)

$$-\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left[\|q^{(n)}(0)\|_{\infty} \geq 1 \right] \geq J_{**}. \quad (3)$$

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

5.2 Universal Large Deviations Upper Bound

We next derive a *uniform* upper bound for the stationary buffer overflow probability decay rate, over all singleton-CSI scheduling algorithms. A popular approach followed in recent work [25, 30, 17] to do this is by estimating the cost of “straight-line” joint channel state sample paths that universally cause buffer overflow. However, when only a *dynamically selected portion of the channel state* is visible to the scheduling algorithm, the cost (1) of such straight-line paths depends explicitly on the algorithm’s sampling behavior, so the standard approach fails.

For every i , let $\phi_i \geq 0$ denote a “twisted” mean rate for channel i , and consider the quantity $\frac{\sum_i c'_i \Lambda_i^*(\phi_i)}{[\max_i (\lambda_i - c'_i \phi_i)]^+}$. Here, we assume that $\sum_i c'_i = 1$, and that the fraction is ∞ whenever the denominator is 0. Suppose a scheduling policy samples each channel i with frequency c'_i . Then, (a) the numerator of the above expression corresponds to the “instantaneous large deviations cost” of witnessing each channel i ’s mean rate be ϕ_i (by (1)), while (b) the denominator can be interpreted as the average rate with which the longest queue grows when

each channel i is sampled with a frequency c'_i . Maximizing the expression over all possible user sampling/scheduling frequencies $\{c'_i : \sum_i c'_i = 1, c'_i \geq 0\}$ induced by scheduling algorithms should thus give the highest possible large deviations cost for buffer overflow. This intuition is formalized in the following key result:

Proposition 4 (Universal Upper Bound on Decay Rate for any Algorithm) *Let π be a stabilizing scheduling policy¹⁰ for the arrival rate $\lambda = (\lambda_1, \dots, \lambda_N)$, and let \mathbb{P}^π be its associated stationary measure. For any $\phi_i \in \mathbb{R}^+$, $i = 1, \dots, N$,*

$$-\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}^\pi \left[\|q^{(n)}(0)\|_\infty \geq 1 \right] \leq \sup_{\substack{\sum_i c'_i = 1 \\ c'_i \geq 0}} \frac{\sum_i c'_i \Lambda_i^*(\phi_i)}{[\max_i (\lambda_i - c'_i \phi_i)]^+}. \quad (4)$$

Note: Each choice of the twisted means $(\phi_i)_i$ above yields such an upper bound on the decay rate. Thus, the best possible upper bound is obtained by minimizing (4) over all choices $(\phi_i)_i$.

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

We refer the reader to Appendix C for the proof of Proposition 4. At the heart of the proof of Proposition 4 is a twisted measure construction where each channel’s marginal rate is ϕ_i . Observing that the cumulative fluid service process $m(\cdot)$ is a submartingale under the twisted measure for any scheduling algorithm, the Doob-Meyer decomposition [7] allows us to express $m(\cdot)$ as the predictable algorithm-dependent component $\phi_i c_i(\cdot)$ plus a martingale noise component $\bar{m}(\cdot)$. This shows that with high probability, the service provided to each queue i is approximated by $\phi_i c_i(\cdot)$, i.e., we can effectively treat each channel i as having a deterministic fluid service rate of ϕ_i . Analyzing this deterministic fluid system for overflow and translating the results back to the original probabilistic system gives us the result.

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

The final step in the proof of optimality of Max-Queue (Corollary 1) is carried out by showing that the lower bound for Max-Queue (3) in fact dominates the uniform upper bound (4) over all scheduling policies:

¹⁰ By a stabilizing scheduling policy π , we mean a scheduling rule that operates under the scheduling model described in Section 2, and which makes the discrete time Markov chain of queue lengths aperiodic, irreducible and positive recurrent.

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

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

The proof of this result involves solving the non-convex problem for the rate function lower bound given in Proposition 3, and relating the solution to a suitable uniform upper bound of the type prescribed by Proposition 4. It utilizes the convexity and lower-semicontinuity of the rate functions Λ_i^* , and is accomplished by considering the properties of the $(\phi_i)_i$ which minimize the upper bound (4). The full proof appears in Appendix D.

6 Analysis: General Subsets and Max-Exp

In this section, we extend the queue overflow optimality result for Max-Queue to the general setting of arbitrary disjoint subsets of observable channels and the Max-Exp scheduling algorithm. For this, we follow the same key steps in obtaining the Max-Queue result – (a) prove lower bounds on the buffer overflow exponent for Max-Exp, (b) derive universal upper bounds on the buffer overflow exponent across all scheduling algorithms using subset channel state information, and (c) demonstrate that the upper and lower bounds match.

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

Hence, our analysis for Max-Exp proceeds by looking at sample paths of the system’s processes over intervals of $O(\sqrt{n})$ time slots. For Step (a) above, analogous to Proposition 1, we establish a “refined” Mogulskii-type theorem for sample-path large deviations of predictably sampled processes over a sub- $O(n)$ timescale (a corresponding result for the full-CSI case was first proved in [25]). Next, we use the framework of *Local Fluid Sample Paths* (LFSPs, introduced in [21]) to obtain a lower bound on the decay exponent of Max-Exp’s overflow probability. LFSPs allow us to “magnify” the standard $O(n)$ fluid limit processes to examine events on the $O(\sqrt{n})$ “local fluid” timescale,

and this helps us match the lower and upper bounds for the decay exponent to establish the optimality of Max-Exp.

6.1 Lower Bounding Max-Exp's Decay Rate: Refined-timescale Large Deviations for Sampled Processes and Local FSPs

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

The general approach for studying scheduling behavior on finer-than- $O(n)$ timescales is to introduce a positive integer function $u(n)$, such that $u(n) \rightarrow \infty$ and $u(n)/n \rightarrow 0$ as $n \rightarrow \infty$ (see Stolyar [25]). We take $u(n) = \lceil \sqrt{n} \rceil$, which is the relevant timescale for the dynamics of the Max-Exp scheduling rule (1).

For our analysis of the queue overflow rate function, we will need to use this idea, along with the following *variable time discretization* for each observable subset. For any non-decreasing, right-continuous-with-left-limits (RCLL) scalar function h on $[0, \infty)$, and any non-decreasing continuous function $\chi : [0, \infty) \rightarrow [0, \infty)$, let $U_\chi^n h$ denote the continuous and piecewise-linearized (according to χ) version of h constructed as follows: we divide $[0, \infty)$ into the contiguous subintervals $[0, \chi(u(n)/n)]$, $[\chi(u(n)/n), \chi(2u(n)/n)]$, $[\chi(2u(n)/n), \chi(3u(n)/n)]$, \dots , and linearize h between its endpoints in each subinterval. For $t \geq 0$, let $\theta^{(n)}(t)$ be the largest right-endpoint of a sub-interval that does not exceed t . When the functions h and χ are vector-valued of (the same) finite dimension, we employ the same notation $U_\chi^n h$ to mean the above linearization performed for each of the individual scalar component functions in h and its counterpart function in χ . In this case, the definition of $\theta^{(n)}(t)$ is similarly extended in a component-wise fashion.

For each observable subset α , let A_α^* be the Sanov rate function [6] for the empirical marginal distribution of the state of its channels $(R_i(1))_{i \in \alpha}$. The domain of A_α^* is the $|\mathcal{R}_\alpha|$ -dimensional simplex where \mathcal{R}_α is the set of all possible sub-states for subset α .

Sampled Trace of the Queueing System: We define here a random object crucial to the analysis of Max-Exp. Consider the evolution of the n -th queueing system in the time slots $1, 2, \dots, nT$, and suppose that subset α is picked by the scheduling algorithm precisely at time slots $K_\alpha(1), K_\alpha(2), \dots, K_\alpha(C_\alpha(nT)) \in \{1, 2, \dots, nT\}$. Recall that \mathcal{R}_α is the set of all possible sub-states wrt subset α . For each such sub-state $r \in \mathcal{R}_\alpha$, we will find it convenient to associate it with the unit vector \mathbf{e}_r which is simply the $|\mathcal{R}_\alpha|$ -dimensional vector with 1 in the r -th position (according to a fixed ordering) and zeros everywhere else.

For each subset α , set

$$V^{\alpha, (n)} \equiv V^\alpha \triangleq (\mathbf{e}_{R_\alpha(K_\alpha(1))}, \mathbf{e}_{R_\alpha(K_\alpha(2))}, \mathbf{e}_{R_\alpha(K_\alpha(3))}, \dots, \mathbf{e}_{R_\alpha(K_\alpha(C_\alpha(nT)))})$$

i.e., the j -th element of V^α simply records what sub-state was sampled when α was picked for the j -th time $K_\alpha(j)$.

We call $V^{(n)} \equiv V \triangleq (V^\alpha)_{\alpha \in \mathcal{O}}$ the *sampled trace* of the queueing system. The sampled trace represents, in words, the sequence of sub-state observations seen by the scheduling algorithm, organized according to the subsets sampled during the operation of the scheduling algorithm. Note also that for any *deterministic* scheduling algorithm, the sampled trace completely specifies the entire sample path of the queue lengths (in conjunction with the arrival sequence which is assumed to be deterministic).

Corresponding to each possible sampled trace V , we define its partial sums process

$$W^{\alpha, (n)}(k) \equiv W^\alpha = \sum_{j=1}^k V^\alpha(j), \quad 1 \leq k \leq C_\alpha(nT)$$

for each observable subset $\alpha \in \mathcal{O}$. We then define $W^{(n)} \equiv W \triangleq (W^\alpha)_{\alpha \in \mathcal{O}}$. Note that each sampled trace V corresponds bijectively to its partial-sums process W . Also, as per convention, we use $w \equiv w^{(n)}$ and $v \equiv v^{(n)}$ to denote the rescaled (by n) versions of W and V respectively.

Let us define the candidate sample-path large deviations rate function for our queueing system as follows:

$$\begin{aligned} \hat{J}_t(z, s) &\triangleq \int_0^{z_\alpha(t)} \sum_{\alpha \in \mathcal{O}} A_\alpha^*(s'_\alpha(u)) du, \\ s_\alpha &\in \mathcal{A}([0, T] \rightarrow \mathbb{R}^{|\mathcal{R}_\alpha|}), \\ z_\alpha &\in \mathcal{A}([0, T] \rightarrow \mathcal{R}), \quad t \in [0, T]. \end{aligned}$$

Here, \mathcal{A} is used to denote the set of absolutely continuous functions.

In order to track large deviations costs over the refined $u(n)$ timescale, let us introduce the notion of a Generalized Fluid Sample Path (GFSP) [25], built upon the framework of standard FSPs.

Definition 1 (Generalized Fluid Sample Path (GFSP)) Suppose that there exists an increasing subsequence $\{n\}$ of the sequence of positive integers such that

1. For each n , there is a valid realization $(f^{(n)}, \hat{f}^{(n)}, c^{(n)}, g^{(n)}, \hat{g}^{(n)}, q^{(n)}, m^{(n)}, w^{(n)})$.
2. As $n \rightarrow \infty$, we have the u.o.c. convergence

$$(f^{(n)}, \hat{f}^{(n)}, c^{(n)}, g^{(n)}, \hat{g}^{(n)}, q^{(n)}, m^{(n)}, w^{(n)}) \rightarrow (f, \hat{f}, c, g, \hat{g}, q, m, w)$$

for a set of limiting, Lipschitz continuous functions $(f, \hat{f}, c, g, \hat{g}, q, m, w)$, and the u.o.c. convergence

$$\bar{J}^{(n)} \equiv (\bar{J}_t^{(n)}, t \in [0, T]) \triangleq \left(\hat{J}_t \left(c^{(n)}, U_{c^{(n)}}^n w^{(n)} \right), t \in [0, T] \right) \rightarrow \bar{J} = (\bar{J}_t, t \in [0, T])$$

for a non-negative non-decreasing Lipschitz-continuous function \bar{J} .

Then, the entire construction

$$\{[n]; (f^{(n)}, \hat{f}^{(n)}, c^{(n)}, g^{(n)}, \hat{g}^{(n)}, q^{(n)}, m^{(n)}, w^{(n)}), \bar{J}^{(n)}; (f, \hat{f}, c, g, \hat{g}, q, m, w), \bar{J}\}$$

is called a generalized fluid sample path (GFSP). The non-decreasing function \bar{J} will be called the *refined cost function* of the GFSP.

We note that for any $0 \leq t_1 < t_2 < \infty$,

$$\bar{J}_{t_2} - \bar{J}_{t_1} \geq \hat{J}_{t_2}(c, w) - \hat{J}_{t_1}(c, w), \quad (5)$$

as a result of convexity of the Λ_α^* , $\alpha \in \mathcal{O}$, and Jensen's inequality.

The following finite-horizon result strengthens Proposition 1. It states that for any *nonrandom* scheduling algorithm, the sample path large deviations rate function for the queue length process is lower-bounded by the minimum *refined cost* over valid GFSPs.

Proposition 6 (Refined-time-scale Lower Bound on Large Deviation Rate Function) *Let Γ be a closed set of trajectories in $\mathcal{C}_{\mathcal{L}}^+([0, T])$. Then, under a nonrandom scheduling policy,*

$$-\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_0^{n, T} [q^{(n)} \in \Gamma] \geq \inf \{ \bar{J}_0 : \exists \text{GFSP } \psi \text{ on } [0, T], \bar{J} \in \psi, q \in \psi, q \in \Gamma \}. \quad (6)$$

The proof appears below. The proof uses ideas from the large deviations of sampling (in the manner of Proposition 1), the crucial concept of *sampled traces*, and a *variable discretization*-version of a refined Mogulskii theorem first shown by Stolyar [25], in order to establish the rate function bound (6).

Proof For an observable subset α , let $\tilde{\mathbb{P}}_\alpha$ be the probability measure on $\{\mathbf{e}_r : r \in \mathcal{R}_\alpha\}$ such that $\tilde{\mathbb{P}}_\alpha[\mathbf{e}_r] = \mathbb{P}[R_\alpha(1) = r] \forall r \in \mathcal{R}_\alpha$. Form the ‘‘marginal’’ product distribution for subset α as $\hat{\mathbb{P}}_\alpha \triangleq \tilde{\mathbb{P}}_\alpha \times \tilde{\mathbb{P}}_\alpha \times \dots$ (i.e., extend $\tilde{\mathbb{P}}_\alpha$ to countably infinite sequences in an iid fashion), and finally take the product of these marginal measures, across observable subsets, to get $\hat{\mathbb{P}} \triangleq \prod_{\alpha \in \mathcal{O}} \hat{\mathbb{P}}_\alpha$. For any candidate sampled trace $v = (v^\alpha)_{\alpha \in \mathcal{O}}$, we understand $\hat{\mathbb{P}}[v]$ as $\prod_{\alpha \in \mathcal{O}} \hat{\mathbb{P}}_\alpha[v^\alpha] = \prod_{\alpha \in \mathcal{O}} \prod_j \tilde{\mathbb{P}}_\alpha[v^\alpha(j)]$.

The lemma below states that for any *deterministic* scheduling algorithm, the probability distribution of the sampled trace of the queueing system is *identical* under both the original measure $\mathbb{P}_{q_0}^{n, T}$ and the product-of-marginals measure $\hat{\mathbb{P}}$ defined above. This will subsequently allow us to apply sample-path large-deviations results on the iid measure $\hat{\mathbb{P}}$ instead of the more complex, correlated measure $\mathbb{P}_{q_0}^{n, T}$.

Lemma 1 *For every sampled trace $v = (v^\alpha)_{\alpha \in \mathcal{O}}$, $\mathbb{P}_{q_0}^{n, T}[V = v] = \hat{\mathbb{P}}[v]$.*

Continuing with the proof of the proposition, let $\Gamma^{(n)}$ (resp. $\mathcal{C}^{(n)}$) be the set of all valid rescaled sampled traces (resp. all valid rescaled subset-sampling trajectories $c^{(n)}$) in the n -th queueing system¹¹ that, starting with initial queue lengths q_0 , result in queue length sample paths belonging to Γ . We also let $\Gamma\mathcal{C}^{(n)} \triangleq \{(w^{(n)}, c^{(n)}) : w^{(n)} \in \Gamma^{(n)}\}$ denote the set of valid (sampled trace, sampling frequency) pairs corresponding to sampled trace trajectories belonging to $\Gamma^{(n)}$. We have, due to Lemma 1,

$$\begin{aligned} \mathbb{P}_{q_0}^{nT} [w^{(n)} \in \Gamma^{(n)}] &= \sum_{w \in \Gamma^{(n)}} \mathbb{P}_{q_0}^{nT} [W^{(n)} = w] \\ &= \sum_{w \in \Gamma^{(n)}} \hat{\mathbb{P}}[w] \\ &= \hat{\mathbb{P}} \left[\left\{ w : w \in \Gamma^{(n)} \right\} \right], \end{aligned}$$

where the final step is due to the fact that the sampled trace uniquely specifies the queueing system's complete trajectory, and so sampled traces corresponding to different sample paths of the system must necessarily be different. Passing in this fashion to the iid measure $\hat{\mathbb{P}}$ allows us to use a refinement of Mogulskii's theorem [6] first established by Stolyar [25, Theorem 7.1], to estimate the large deviations rate function. As a consequence, we can write¹²

$$-\limsup_{n \rightarrow \infty} \frac{1}{n} \log \hat{\mathbb{P}} \left[\left\{ w^{(n)} : w^{(n)} \in \Gamma^{(n)} \right\} \right] \geq \liminf_{n \rightarrow \infty} \inf \left\{ \hat{J}_T(c, U_c^n w) : (w, c) \in \Gamma\mathcal{C}^{(n)} \right\}. \quad (7)$$

Let the limit inferior on the right-hand side of (7) above be denoted by ζ . It follows that we can find for each n a $w_n \in \Gamma^{(n)}$ and $c_n = (c_{n,\alpha})_{\alpha \in \mathcal{O}} \in \mathbb{R}^{|\mathcal{O}|}$, such that $(w_n, c_n) \in \Gamma\mathcal{C}^{(n)}$ and $\hat{J}_0(c_n, U_{c_n}^n w_n) \rightarrow \zeta$. Using uniform Lipschitz continuity of the $\{w_n\}$ and $\{c_n\}$, we can extract a subsequence of trajectories (w_n, c_n) which converges and forms a GFSP with refined cost ζ , and which satisfies, by construction, the conditions on the right-hand side of (6). Thus, we get

$$\zeta \geq \inf \left\{ \bar{J}_0 : \exists \text{ GFSP } \psi \text{ on } [0, T], \bar{J} \in \psi, q \in \psi, q \in \Gamma \right\},$$

completing the proof.

¹¹ Since sampled traces and their partial sums processes are in one-to-one correspondence, we take the liberty of referring to them interchangeably.

¹² [25, Theorem 7.1] derives the rate function bound under a *fixed* discretization of the time axis where the discretization rate is always unity, i.e., $\chi(x) = x \forall x \geq 0$. When the discretization rate is variable and depends on the subset selection frequency $c^{(n)}$, it is not hard to see that the result of [25, Theorem 7.1] extends with $U^n h$ replaced by $U_{c^{(n)}}^n h$ – the key property that affords us this extension is that the *size* of each subinterval $(u(n)/n$ in [25] and $\chi((k+1)u(n)/n) - \chi(ku(n)/n)$ here) does not matter; it is the *number* of discretized intervals $(Tu(n)/n$ both in [25] and here) that is the crucial ingredient in the bound in [25, Theorem 7.1].

Similar to extending the result of Proposition 1 to the stationary queue length distribution, minimizing the RHS of (6) across FSPs over all finite time horizons $T > 0$ yields a lower bound for the large deviations rate of the *stationary* queue overflow probability. This uses standard tools (see, for instance, [25,30]), and we omit the proof for brevity.

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 Proposition 2), we can use standard Friedlin-Wentzell-type techniques to extend Proposition 6 to a large-deviations lower bound [30,25] for the stationary measure under the Max-Exp scheduling policy. Note that this requires showing that Max-Exp is throughput-optimal¹³ – a fact whose proof we omit for brevity, but which results from a fairly straightforward modification of the proof of throughput-optimality of the Max-Sum Queue algorithm (see [10] for details).

Theorem 2 *Let \mathbb{P} denote the stationary measure induced by the Max-Exp policy. Then,*

$$\begin{aligned} & -\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left[\|q^{(n)}(0)\|_{\infty} \geq 1 \right] \\ & \geq \inf_{T \geq 0} \inf \left\{ \bar{J}_t : \exists \text{ GFSP } \psi \text{ on } [0, T], \bar{J} \in \psi, q \in \psi, t \in [0, T], q(0) = 0, \|q(t)\|_{\infty} \geq 1 \right\}. \end{aligned} \tag{8}$$

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 Proposition 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}_{α} denotes the (finite) set of all possible (joint) sub-states that can be observed channels in α . We use Π_{α} to denote the $|\mathcal{R}_{\alpha}|$ -valued simplex, i.e., the set of all probability measures on the sub-states of α . Any distribution $\phi_{\alpha} \in \Pi_{\alpha}$ induces a *subset throughput region* $V_{\phi_{\alpha}}$, which represents all the long-term average service rates that can be sustained to users in α when the sub-states are distributed as ϕ_{α} (see also [2,10]). The uniform large-deviations upper-bound can now be stated for any stabilizing scheduling policy π :

¹³ Again, the throughput-optimality holds among all scheduling algorithms that base their decision on the current (or any finite past history of) system state.

Theorem 3 *Let π be a stabilizing scheduling policy for arrival rates $\lambda = (\lambda_1, \dots, \lambda_n)$, and let \mathbb{P}^π be the associated stationary measure. Let distributions $\phi_\alpha \in \Pi_\alpha$ be fixed, for every α , such that $\lambda \notin \mathcal{CH}((V_{\phi_\alpha})_\alpha)$. Then,*

$$-\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}^\pi \left[\|q^{(n)}(0)\|_\infty \geq 1 \right] \leq \sup_{\substack{\sum_\alpha c'_\alpha = 1 \\ c'_\alpha \geq 0}} \left[\frac{\sum_\alpha c'_\alpha \Lambda_\alpha^*(\phi'_\alpha)}{\max_{\alpha, v_\alpha \in V_{\phi_\alpha}} \max_{i \in \alpha} (\lambda_i - c'_\alpha v_{\alpha, i})} \right]. \quad (9)$$

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 rules¹⁴. For this, we leverage the large-deviations lower bound for the Max-Exp scheduling algorithm (Theorem 2) and show that it is actually a uniform upper bound over all scheduling rules as prescribed by Theorem 3.

Our approach at the high level is comprised of the following steps:

1. Consider a feasible FSP (q, \bar{J}) on $[0, T]$ for Theorem 2, i.e., $q(0) = 0$, $q(t) = 1$ for some $t \in [0, T]$. We show, by “magnifying” the FSP about some $\tau \in [0, T]$ and taking “local” fluid limits, that the “unit large-deviations cost” of raising the longest queue in the associated Local Fluid Sample Path (LFSP) [21, 25] at τ is close to the total FSP cost \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 longest queue” over all feasible local fluid sample paths – call it J_* .
3. In the context of Theorem 3, we exhibit suitable twisted subset distributions $\phi_\alpha \in \Pi_\alpha \forall \alpha$ such that the RHS of (9) 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 (8) 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.

¹⁴ By optimal, we mean optimal among all stabilizing scheduling algorithms that base their decision on the current (or any finite past history of) system state.

The formal LFSP construction is along the lines of that used in [21,25], and is as follows. Consider a standard fluid sample path on $[0, T]$ (along with its prelimit functions) and call it ψ . Let us introduce the “recentered” queue lengths

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

where $b_i, i = 1, \dots, N$ are such that for each 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

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

so we have the uniform convergence

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

and

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

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

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

It follows that we can choose a subsequence of n along which the following uniform convergence to Lipschitz functions holds on $[0, S]$ [25]:

$$\begin{aligned} & \left(\diamond q^{(n)}, \diamond q_*^{(n)}, \diamond f^{(n)}, \diamond \hat{f}^{(n)}, (\diamond c_\alpha^{(n)})_\alpha, (\diamond g_r^{\alpha, (n)})_{\alpha r}, (\diamond \hat{g}_{ri}^{\alpha, (n)})_{\alpha r, i}, \diamond m^{(n)} \right) \rightarrow \\ & \left(\diamond q, \diamond q_*, \diamond f, \diamond \hat{f}, (\diamond g_r^\alpha)_{\alpha r}, (\diamond \hat{g}_{ri}^\alpha)_{\alpha r, i}, \diamond m \right). \end{aligned} \quad (10)$$

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 call the tuple on the right-hand side of (10) above a *Local Fluid Sample Path* at (scaled) time τ .

We also have the following consequence of the (marginal) convexity of \hat{J} , in close analogy with (5):

$$\liminf_{n \rightarrow \infty} \frac{1}{\sigma_n} [\bar{J}_{t_2}^{(n)} - \bar{J}_{t_1}^{(n)}] \geq \hat{J}_S \left(\diamond c, \frac{d_\diamond g}{d_\diamond c} \right). \quad (11)$$

The following key lemma, along the lines of Lemma 9.1 in [25], is crucial to understand the local timescale dynamics of the Max-Exp scheduling algorithm:

Lemma 2 *For any LFSP over an interval $[0, S]$,*

1. *The following derivatives exist Lebesgue-a.e.¹⁵ and are finite:*

$$\diamond \dot{q}, \quad \diamond \dot{q}_*, \quad \diamond \dot{c}_\alpha, \quad \diamond \dot{f}, \quad v \triangleq \diamond \dot{\hat{f}}, \quad \diamond \dot{g}_r^\alpha, \quad \diamond \dot{\hat{g}}_{ri}^\alpha, \quad \diamond \dot{m}.$$

2. *For every $\alpha \in \mathcal{O}$ and $r \in \mathcal{R}_\alpha$, $\diamond g_r^\alpha \ll \diamond c_\alpha$ (wrt the corresponding Lebesgue-Stieltjes induced measures). Thus, there exists (c_α -a.e.) a version of the Radon-Nikodym derivative $\phi_{\alpha r} \triangleq \frac{d_\diamond g_r^\alpha}{d_\diamond c_\alpha}$.*

¹⁵ As convention, we take $\frac{d}{dt}(\infty) = 0$.

3. The following relations hold, whenever the relevant derivatives exist, for $s \in [0, S]$:

$$\diamond \dot{f}(s) = \lambda, \quad (12)$$

$$\diamond \dot{q}(s) = \diamond \dot{f}(s) - v(s), \quad (13)$$

$$\diamond q_*(s) = \max_i \diamond q_i(s), \quad (14)$$

$$\diamond \dot{q}_*(s) = \diamond \dot{q}_i(s) \quad \text{for each } i \text{ such that } \diamond q_i(s) = \diamond q_*(s), \quad (15)$$

$$v_i(s) = \sum_{r \in \mathcal{R}_\alpha} \diamond \dot{g}_{ri}^\alpha(s) \mu_{ri}^\alpha \quad \text{for each } i \in \alpha, \alpha \in \mathcal{O}, \quad (16)$$

$$\sum_{i \in \alpha} \diamond \dot{g}_{ri}^\alpha(s) = \diamond \dot{g}_r^\alpha(s), \quad (17)$$

$$\sum_{r \in \mathcal{R}_\alpha} \diamond \dot{g}_r^\alpha(s) = \diamond \dot{c}_\alpha(s), \quad (18)$$

$$\sum_{r \in \mathcal{R}_\alpha} \diamond \phi_{\alpha r}(s) = 1, \quad (19)$$

$$\sum_{\alpha \in \mathcal{O}} \diamond \dot{c}_\alpha(s) = 1, \quad (20)$$

$$\forall \alpha \in \mathcal{O} \quad v_\alpha(s) = \diamond \dot{c}_\alpha(s) \times \arg \max_{\eta \in V_{\phi_\alpha}(s)} \langle e^{\diamond q(s)+b}, \eta \rangle_\alpha, \quad (21)$$

$$\sum_{i \in \beta} e^{\diamond q_i(s)+b_i} < \max_{\alpha \in \mathcal{O}} \sum_{i \in \alpha} e^{\diamond q_i(s)+b_i} \Rightarrow \diamond \dot{c}_\beta(s) = 0 \quad \text{for each } \beta \in \mathcal{O}, \quad (22)$$

$$\left. \frac{d}{du} \sum_{i \in \beta} e^{\diamond q_i(u)+b_i} \right|_{u=s} = \left. \frac{d}{du} \sum_{i \in \gamma} e^{\diamond q_i(u)+b_i} \right|_{u=s}$$

whenever $\beta, \gamma \in \arg \max_{\alpha \in \mathcal{O}} \sum_{i \in \alpha} e^{\diamond q_i(s)+b_i}$, and

$$s \text{ is a regular point of } \sum_{i \in \beta} e^{\diamond q_i(s)+b_i}, \sum_{i \in \gamma} e^{\diamond q_i(s)+b_i} \text{ and } \max_{\alpha \in \{\beta, \gamma\}} \sum_{i \in \alpha} e^{\diamond q_i(s)+b_i}. \quad (23)$$

Proof The first assertion of the lemma follows due to the absolute continuity of the LFSP functions being considered, which, in turn, is a consequence of the corresponding Lipschitz-continuous prelimit functions.

The second assertion of the lemma is due to the fact that $\diamond g_r^{\alpha, (n)} \ll \diamond c_\alpha^{(n)}$ for the prelimit functions – a queue belonging to a subset cannot be scheduled without first choosing the subset.

As regards the third assertion, properties (12)-(20) follow due to the corresponding properties of their prelimit LFSP functions, together with the (Lebesgue-a.e.) derivatives.

Property (21) is a key property of the Exponential scheduling rule, and has been established previously in the work of Shakkottai and Stolyar [21]. It is a consequence of the ratios of the $\exp(\cdot)$ terms (for different queues i) in the definition of the intra-subset Exponential rule (i.e., Step (2) in Algorithm 1) converging to the ratios $e^{\circ q_i(s)+b_i}$ on the LFSP time scale.

Property (22) follows from Max-Exp's subset selection criterion (i.e., Step (1) in Algorithm 1) applied to the prelimit LFSP functions, along with the convergence of the $\exp(\cdot)$ terms to the ratios $e^{\circ q_i(s)+b_i}$ as noted above.

To show (23), we argue by contradiction. If s is a regular point of the subset functions $L_\beta(\cdot) \triangleq \sum_{i \in \beta} e^{\circ q_i(\cdot)+b_i}$ and $L_\gamma(\cdot) \triangleq \sum_{i \in \gamma} e^{\circ q_i(\cdot)+b_i}$, and $L_\beta(s) = L_\gamma(s)$ holds, but

$$\left. \frac{d}{du} L_\beta(s) \right|_{u=s} \neq \left. \frac{d}{du} L_\gamma(s) \right|_{u=s},$$

then a simple argument shows that the function $\max(L_\beta(\cdot), L_\gamma(\cdot))$ cannot be differentiable at s , yielding a contradiction.

With this framework of LFSPs set up, we can resume the main development from Theorem 2. Consider a feasible GFSP ψ on $[0, T]$ for the right-hand side of (8) (i.e., for which $q(0) = 0$ and $\|q(t)\|_\infty = 1$ for some $t \in [0, T]$), and whose refined cost is \bar{J}_t . Fix an arbitrary $\epsilon > 0$. Then, there must exist a time point $\tau \in (0, t)$ such that $q_*(\tau) > 0$, $q'_*(\tau) > 0$, $\bar{J}'_\tau > 0$, and

$$\frac{\bar{J}'_\tau}{q'_*(\tau)} < \bar{J}_t + \epsilon.$$

Continuing using a technique similar to that in [25, Section 11], we can show that for an arbitrary $S > 0$ and sufficiently large n , we can find intervals $[t_1^{(n)}, t_2^{(n)}]$ and $[t_1, t_2]$ such that $\tau \in [t_1^{(n)}, t_2^{(n)}] \subset [t_1, t_2]$, and with

$$t_2^{(n)} - t_1^{(n)} = \frac{S}{\sqrt{n}} \sqrt{\bar{q}^{(n)}(t_1^{(n)})},$$

$$\tilde{q}_*^{(n)}(t_2^{(n)}) - \tilde{q}_*^{(n)}(t_1^{(n)}) > 0, \quad (24)$$

$$\frac{\bar{J}_{t_2^{(n)}}^{(n)} - \bar{J}_{t_1^{(n)}}^{(n)}}{\tilde{q}_*^{(n)}(t_2^{(n)}) - \tilde{q}_*^{(n)}(t_1^{(n)})} < \bar{J}_t + 3\epsilon. \quad (25)$$

We can choose a subsequence of $\{n\}$ above so that, for some $\tau_1 \in [t_1, t_2]$, we have the left endpoints $t_1^{(n)} \rightarrow \tau_1$ (so that $t_2^{(n)} \rightarrow \tau_1$ as well). Then, let us choose a further subsequence such that

$$\left(\diamond q^{(n)}, \diamond q_*^{(n)}, \diamond f^{(n)}, \diamond \hat{f}^{(n)}, (\diamond c_\alpha^{(n)})_\alpha, (\diamond g_r^{\alpha, (n)})_{\alpha r}, (\diamond \hat{g}_{ri}^{\alpha, (n)})_{\alpha r, i}, \diamond m^{(n)} \right)$$

converges to an LFSP $\left(\diamond q, \diamond q_*, \diamond f, \diamond \hat{f}, (\diamond g_r^\alpha)_{\alpha r}, (\diamond \hat{g}_{ri}^\alpha)_{\alpha r, i}, \diamond m \right)$ on the *local* time interval $[0, S]$.

We claim that there must exist $\epsilon_1 > 0$ such that

$$\hat{J}_S(\diamond c, \phi) - \hat{J}_0(\diamond c, \phi) \geq \epsilon_1 S$$

(recall that $\phi = \frac{d_{\diamond} q}{d_{\diamond} c}$ as defined in Lemma 2). If not, then $\hat{J}_S(\diamond c, \phi) = \hat{J}_0(\diamond c, \phi)$, which means that all observed channel state distributions over subsets are *exactly* typical. Since, by hypothesis, the arrival rate vector λ lies in the interior of the throughput region, this contradicts the fact that the longest queue in the system does not decrease (24).

The inequality above, together with the lower bound (11) and the relation (25), gives us that there exists $\epsilon_2 > 0$ such that

$$\diamond q_*(S) - \diamond q_*(0) \geq \epsilon_2 S. \quad (26)$$

We thus arrive at the inequality

$$\frac{\hat{J}_S(\diamond c, \phi) - \hat{J}_0(\diamond c, \phi)}{\diamond q_*(S) - \diamond q_*(0)} \leq \bar{J}_t + 3\epsilon. \quad (27)$$

In other words, we are able to approximate the cost of FSPs arbitrarily well with the “unit cost” of raising $\diamond q_*$ in suitably constructed LFSPs.

6.4.2 A Relaxed 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:

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

together with its logarithm

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

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 infimum¹⁶ (8) in Theorem 2. By combining the above fact with the conclusions of the previous section (i.e. properties (26) and (27)), we have that for an arbitrarily small $\epsilon > 0$, an LFSP can be constructed on $[0, S]$, with $S > 0$ suitably large, so that the following properties hold with $\epsilon_2 > 0$:

$$\Phi(\diamond q(S)) - \Phi(\diamond q(0)) \geq (\epsilon_2/2)S, \quad (28)$$

$$\frac{\hat{J}_S(\diamond c, \phi) - \hat{J}_0(\diamond c, \phi)}{\Phi(\diamond q(S)) - \Phi(\diamond q(0))} \leq \bar{J}_t + 2\epsilon. \quad (29)$$

¹⁶ If the infimum is not attainable, it suffices to consider an FSP ϵ' -close to the infimum, with $\epsilon' > 0$.

In the sequel, we will concentrate on the LHS of (29) – modulo an arbitrarily small $\epsilon > 0$, it is a lower bound on the original FSP cost \bar{J}_t . We have

$$\begin{aligned} \frac{\hat{J}_S(\diamond c, \phi) - \hat{J}_0(\diamond c, \phi)}{\Phi(\diamond q(S)) - \Phi(\diamond q(0))} &= \frac{\int_0^S \frac{d}{ds} \hat{J}_s(\diamond c, \phi) ds}{\int_0^S \frac{d}{ds} \Phi(\diamond q(s)) ds} \\ &\geq \inf_{s \in [0, S]} \frac{\frac{d}{ds} \hat{J}_s(\diamond c, \phi)}{\frac{d}{ds} \Phi(\diamond q(s))} \\ &= \inf_{s \in [0, S]} \frac{\sum_{\alpha} \diamond \dot{c}_{\alpha}(s) \Lambda_{\alpha}^*(\phi_{\alpha}(s))}{\frac{d}{ds} \Phi(\diamond q(s))} \quad (\text{Lemma 2}). \end{aligned}$$

As a consequence of the above inequality¹⁷, we can record the following result:

Proposition 7 *If \mathbb{P} denotes the stationary measure induced by the Max-Exp policy, then*

$$-\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left[\|q^{(n)}(0)\|_{\infty} \geq 1 \right] \geq \inf_{s \in [0, S]} \frac{\sum_{\alpha} \diamond \dot{c}_{\alpha}(s) \Lambda_{\alpha}^*(\phi_{\alpha}(s))}{\frac{d}{ds} \Phi(\diamond q(s))}, \quad (30)$$

for any valid Local Fluid Sample Path (LFSP) as specified by (10).

Letting J_* denote the infimum on the RHS of (30) over *all* valid LFSPs, a further lower bound on the buffer overflow exponent of Max-Exp is thus J_* .

6.4.3 Connecting the relaxed 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_* is, in fact, a uniform upper bound on the decay exponent of any stabilizing scheduling policy, on the lines of Theorem 3. The proof may be found in Appendix E, and uses the disjointness of the collection of observable subsets \mathcal{O} in a key way.

Theorem 4 (Optimality of Max-Exp) *Let π be any stabilizing scheduling policy (i.e., a stabilizing policy that bases its decision on the current (or any finite past history of) system state) for arrival rates $\lambda = (\lambda_1, \dots, \lambda_n)$, and let \mathbb{P}^{π} be the associated stationary measure. Then,*

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

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

¹⁷ We have abused notation to indicate that the infimum above is, in fact, over the (Lebesgue-a.e.) regular points $s \in [0, S]$.

7 Conclusion

For scheduling with only partial wireless Channel State Information (CSI), we developed the Max-Exp and Max-Queue scheduling algorithms yielding optimal queue overflow tails. This work shows that structurally simple scheduling algorithms which use partial CSI can guarantee 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 partial 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.

References

1. Andrews, J.G., Ghosh, A., Muhamed, R.: *Fundamentals of WiMAX: Understanding Broadband Wireless Networking*. Prentice Hall (2007)
2. Andrews, M., Kumaran, K., Ramanan, K., Stolyar, A.L., Vijayakumar, R., Whiting, P.: Scheduling in a queueing system with asynchronously varying service rates. *Probability in Engineering and Informational Sciences* **14**, 191–217 (2004)
3. Bertsimas, D., Paschalidis, I., Tsitsiklis, J.: Asymptotic buffer overflow probabilities in multiclass multiplexers: An optimal control approach. *IEEE Trans. Automat. Contr.* **43**(3), 315–335 (1998)
4. Chang, N., Liu, M.: Optimal channel probing and transmission scheduling for opportunistic spectrum access. In: *ACM Int. Conf. on Mobile Computing and Networking (MobiCom)*. Montreal, Canada (2007)
5. Chaporkar, P., Proutiere, A., Asnani, H., Karandikar, A.: Scheduling with limited information in wireless systems. In: *ACM Mobihoc*, pp. 75–84 (2009)
6. Dembo, A., Zeitouni, O.: *Large Deviations Techniques and Applications*. Jones and Bartlett (1993)
7. Durrett, R.: *Probability: Theory and Examples*. Brooks/Cole - Thomson Learning, Belmont, CA (2005)
8. Eryilmaz, A., Srikant, R., Perkins, J.R.: Stable scheduling policies for fading wireless channels. *IEEE/ACM Trans. Networking* **13**, 411–424 (2005)
9. Gopalan, A., Caramanis, C., Shakkottai, S.: Low-delay wireless scheduling with partial channel-state information. In: *Proc. IEEE INFOCOM*, pp. 1071–1079 (2012)
10. Gopalan, A., Caramanis, C., Shakkottai, S.: On wireless scheduling with partial channel state information. *IEEE Trans. Inform. Theory* **58**(1), 403–420 (2012)
11. Kar, K., Luo, X., Sarkar, S.: Throughput-optimal scheduling in multichannel access point networks under infrequent channel measurements. *IEEE Trans. Wireless Commun.* **7**(7), 2619–2629 (2008)
12. Mannor, S., Shamir, O.: From bandits to experts: On the value of side-observations. In: J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, K. Weinberger (eds.) *Advances in Neural Information Processing Systems 24*, pp. 684–692. Curran Associates, Inc. (2011)
13. Mitzenmacher, M., Upfal, E.: *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, New York, NY, USA (2005)
14. Neely, M., Modiano, E., Rohrs, C.: Dynamic power allocation and routing for time-varying wireless networks. *IEEE J. Sel. Areas Commun.* **23**(1), 89–103 (2005)
15. Parkvall, S., Astely, D.: The evolution of LTE towards IMT-advanced. *Journal of Communications* **4**(3) (2009)

16. Patil, S., de Veciana, G.: Reducing feedback for opportunistic scheduling in wireless systems. *IEEE Trans. Wireless Commun.* **12**(6), 4227–38 (2007)
17. Sadiq, B., de Veciana, G.: Optimality and large deviations for multi-user wireless systems under pseudo-log opportunistic scheduling. In: *Proc. Ann. Allerton Conf. Communication, Control and Computing*, pp. 1–8 (2008)
18. Shakkottai, S.: Effective capacity and QoS for wireless scheduling. *IEEE Trans. Automat. Contr.* (2008)
19. Shakkottai, S., Srikant, R., Stolyar, A.L.: Pathwise optimality and state space collapse for the exponential rule. In: *Proc. IEEE Int. Symp. Information Theory (ISIT)*. Lausanne, Switzerland (2002)
20. Shakkottai, S., Srikant, R., Stolyar, A.L.: Pathwise optimality of the exponential scheduling rule for wireless channels. *Adv. Appl. Prob.* **36**(4), 1021–1045 (2004)
21. Shakkottai, S., Stolyar, A.: Scheduling for multiple flows sharing a time-varying channel: The exponential rule. *Analytic Methods in Applied Probability*. American Mathematical Society Translations, Series 2, A volume in memory of F. Karpelevich, Yu. M. Suhov, Editor **207** (2002)
22. Shakkottai, S., Stolyar, A.: 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* **207** (2002)
23. Shimkin, N.: Extremal large deviations in controlled IID processes with applications to hypothesis testing. *Adv. Appl. Prob.* **25**(4), pp. 875–894 (1993)
24. Stolyar, A.: Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *Annals of Appl. Prob.* **14**, No.1, 1–53 (2004)
25. Stolyar, A.: Large deviations of queues sharing a randomly time-varying server. *Queueing Systems - Theory and Applications* pp. 1–35 (2008)
26. Stolyar, A., Ramanan, K.: Largest weighted delay first scheduling: Large deviations and optimality. *Annals of Appl. Prob.* pp. 1–48 (2001)
27. Tang, T., Heath, R.W.: Opportunistic feedback for downlink multiuser diversity. *IEEE Commun. Letters* **9**, 948–950 (2005)
28. Tassiulas, L., Ephremides, A.: Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automat. Contr.* **4**, 1936–1948 (1992)
29. Venkataramanan, V.J., Lin, X.: Structural properties of LDP for queue-length based wireless scheduling algorithms. In: *Proc. Ann. Allerton Conf. Communication, Control and Computing* (2007)
30. Venkataramanan, V.J., Lin, X.: On the queue-overflow probability of wireless systems: A new approach combining large deviations with Lyapunov functions. *IEEE Trans. Inform. Theory* **59**, 6367–6392 (2013)
31. Ying, L., Dullerud, G., Srikant, R.: A large deviations analysis of scheduling in wireless networks. *IEEE Trans. Inform. Theory* **52** (2006)
32. Ying, L., Shakkottai, S.: On throughput optimality with delayed network-state information. *IEEE Trans. Inform. Theory* **57**(8), 5116–5132 (2011)

A Proof of Proposition 1

In the n th system, consider the joint channel states for the first nT time slots, i.e., $(R^{(n)}(1), R^{(n)}(k), \dots, R^{(n)}(nT))$, 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 $\{1, \dots, nT\}$ 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 n th queueing system. Let $V = (V_1, \dots, V_N)$ 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. $V_i = (R_i(K_{i_1}), R_i(K_{i_2}), \dots)$ where user i is chosen precisely at time slots K_{i_1}, K_{i_2}, \dots . 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 combinations of w , i.e. the set of all (e_1, \dots, e_N) where each e_i is a vector in \mathcal{R}^{nT} such that w_i is a prefix of e_i .

Lemma 3 *Let Z_{ij} , $i = 1, \dots, N$, $j = 1, 2, \dots, nT$ be independent random variables with $Z_{ij} \sim R_i(0)$ for all i and j . Let $\hat{\mathbb{P}}^{(nT)}$ be the probability measure induced by $(Z_{ij})_{i,j}$. If w is a valid trace in the n -th system, then for any $nq_0 \in (\mathbb{Z}^+)^N$,*

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

Proof Let $w = (w_1, \dots, w_N)$ with $\sum_{i=1}^N |w_i| = nT$, and let $v = (v_1, \dots, v_N)$ 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}, \dots when user i is sampled, for all i . Furthermore, a key fact is that all the time slots k_{i_1}, k_{i_2}, \dots 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

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

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_{ij} are all distinct and partition $\{1, 2, \dots, nT\}$; (c) is due to exchangeability of the (independent across time) channel state process.

Proceeding with the proof of the proposition, we have

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

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

$$\begin{aligned} \mathbb{P}_{q_0}^{n,T}[w^{(n)} \in \Gamma^{(n)}] &= \sum_{w \in \Gamma^{(n)}} \mathbb{P}_{q_0}^{n,T}[w^{(n)} = w] \\ &= \sum_{w \in \Gamma^{(n)}} \hat{\mathbb{P}}^{(nT)}[E(w)] \quad (\text{by Lemma 3}) \\ &= \hat{\mathbb{P}}^{(nT)} \left[\bigcup_{w \in \Gamma^{(n)}} E(w) \right] \quad (\text{unique prefixes} \Rightarrow \text{disjointness}) \\ &\leq \hat{\mathbb{P}}^{(nT)} \left[\bigcup_{n'=\hat{n}}^{\infty} E(\Gamma^{(n')}) \right]. \end{aligned}$$

$$\begin{aligned}
\therefore -\limsup_{n \rightarrow \infty} n^{-1} \log \mathbb{P}_{q_0}^{n,T} [w^{(n)} \in \Gamma^{(n)}] &\geq -\limsup_{n \rightarrow \infty} n^{-1} \log \hat{\mathbb{P}}^{(n,T)} \left[\bigcup_{n'=\hat{n}}^{\infty} E(\Gamma^{(n')}) \right] \\
&\geq \inf \left\{ \int_0^T \sum_{i=1}^N A_i^*(\dot{w}_i(z)) dz : w \in \overline{\bigcup_{n'=\hat{n}}^{\infty} E(\Gamma^{(n')})} \right\} \\
&\quad \text{(by Mogulskii's theorem [6])} \\
\Rightarrow -\limsup_{n \rightarrow \infty} n^{-1} \log \mathbb{P}_{q_0}^{n,T} [w^{(n)} \in \Gamma^{(n)}] &\geq \liminf_{\hat{n} \rightarrow \infty} \inf \left\{ \int_0^T \sum_{i=1}^N A_i^*(\dot{w}_i(z)) dz : w \in \overline{\bigcup_{n'=\hat{n}}^{\infty} E(\Gamma^{(n')})} \right\}.
\end{aligned} \tag{31}$$

Let the right hand side of (31) be denoted by ζ . For every $\hat{n} = 1, 2, \dots$, we can choose $w_{\hat{n}}$ such that

$$\begin{aligned}
w_{\hat{n}} &\in \overline{\bigcup_{n'=\hat{n}}^{\infty} E(\Gamma^{(n')})}, \quad \text{and} \\
\lim_{\hat{n} \rightarrow \infty} \int_0^T \sum_{i=1}^N A_i^*(\dot{w}_{\hat{n},i}(z)) dz &= \zeta.
\end{aligned}$$

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 $[0, T]$. Without loss of generality, let the subsequence be $\{\hat{n}\}$ itself, and let $\lim_{\hat{n} \rightarrow \infty} w_{\hat{n}} = w$. The map $f \mapsto \int_0^T \sum_{i=1}^N A_i^*(\dot{f}(z)) dz$ is lower-semicontinuous [6], thus

$$\int_0^T \sum_{i=1}^N A_i^*(\dot{w}(z)) dz \leq \lim_{\hat{n} \rightarrow \infty} \int_0^T \sum_{i=1}^N A_i^*(\dot{w}_{\hat{n},i}(z)) dz = \zeta.$$

We can pick, for each \hat{n} , an $\hat{m}_{\hat{n}}$ and a $w_{\hat{m}_{\hat{n}}} \in \bigcup_{n'=\hat{n}}^{\infty} E(\Gamma^{(n')})$ such that $\|w_{\hat{m}_{\hat{n}}} - w_{\hat{n}}\|_{\infty} < 1/\hat{n}$. Since $w_{\hat{m}_{\hat{n}}} \in \bigcup_{n'=\hat{n}}^{\infty} E(\Gamma^{(n')})$, let $w_{\hat{m}_{\hat{n}}} \in E(\Gamma^{(\hat{m}'_{\hat{n}})})$ 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 the sequence $q_{\hat{m}_{\hat{n}}}$ converges to a $q \in \bar{\Gamma} = \Gamma$. We now have $\lim_{\hat{n} \rightarrow \infty} w_{\hat{m}_{\hat{n}}} = w$ and $\lim_{\hat{n} \rightarrow \infty} q_{\hat{m}_{\hat{n}}} = q$, thus (q, w) is a valid fluid sample path with $q \in \Gamma$. This yields

$$\begin{aligned}
-\limsup_{n \rightarrow \infty} n^{-1} \log \mathbb{P}_{q_0}^{n,T} [w^{(n)} \in \Gamma^{(n)}] &\geq \\
&\inf \left\{ \int_0^T \sum_{i=1}^N A_i^*(\dot{w}_i(z)) dz : (w, q) \text{ an FSP}, q \in \Gamma \right\} \\
\Rightarrow -\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{q_0}^{n,T} [q^{(n)} \in \Gamma] &\geq \\
&\inf \left\{ \int_0^T \sum_{i=1}^N A_i^*(\dot{w}_i(z)) dz : (w, q) \text{ an FSP}, q \in \Gamma \right\}.
\end{aligned} \tag{32}$$

Note: By (w, q) being an FSP, in addition to there existing prelimit sequences $w^{(n)} \rightarrow w$ and $q^{(n)} \rightarrow q$ (uniformly over $[0, T]$), we mean that there exist points $z_i \in [0, T]$ for all $i = 1, \dots, N$ such that $z_i^{(n)} \rightarrow z_i$ as $n \rightarrow \infty$, 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 3, 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 $[0, T]$. 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 4 *Let (w, q) be a fluid sample path with $w^{(n)} \rightarrow w$ and $q^{(n)} \rightarrow q$. For each integer $n \geq 1$, let $(m^{(n)}, c^{(n)})$ be the scaled sampled rate and selection processes, which are completely specified by $w^{(n)}$. Then, for every subsequential limit (m, c) of $(m^{(n)}, c^{(n)})_n$ (in the $\|\cdot\|_\infty$ topology on $[0, T]$),*

$$\sum_{i=1}^N \int_0^{z_i} \Lambda_i^*(\dot{w}_i(z)) dz = \int_0^T \left[\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)} \right) \right] dt.$$

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

$$w_i^{(n)}(c_i^{(n)}(t_2)) - w_i^{(n)}(c_i^{(n)}(t_1)) = m_i^{(n)}(t_2) - m_i^{(n)}(t_1) + O(1/n). \quad (33)$$

By the (uniform) convergence hypotheses, for $j \in \{1, 2\}$, $c_i^{(n)}(t_j) \rightarrow c_i(t_j)$, thus $w_i^{(n)}(c_i^{(n)}(t_j)) \rightarrow w_i(c_i(t_j))$. Letting $n \rightarrow \infty$ in (33),

$$w_i(c_i(t_2)) - w_i(c_i(t_1)) = m_i(t_2) - m_i(t_1). \quad (34)$$

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

$$\begin{aligned} \int_{t_1}^{t_2} \frac{dw_i}{dz} \circ c_i(t) dc_i(t) &= \int_{c_i(t_1)}^{c_i(t_2)} \frac{dw_i}{dz}(z) (dc_i \circ c^{-1}) \quad (\text{change of variables formula}) \\ &= \int_{c_i(t_1)}^{c_i(t_2)} \frac{dw_i}{dz}(z) dz \quad (dc_i \circ c^{-1} \equiv \text{Lebesgue}[0, z_i]) \\ &= w_i(c_i(t_2)) - w_i(c_i(t_1)) \\ &= m_i(t_2) - m_i(t_1) \quad (\text{thanks to (34)}) \\ &= \int_{t_1}^{t_2} dm_i(t) \\ \Rightarrow \frac{dw_i}{dz} \circ c_i(\cdot) &= \frac{dm_i}{dc_i}(\cdot) \quad dc_i\text{-a.e. on } [0, T]. \end{aligned}$$

With this, we can finally compute

$$\begin{aligned} \int_0^T \dot{c}_i(t) \Lambda_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)} \right) dt &= \int_0^T \left(\Lambda_i^* \circ \frac{dm_i}{dc_i} \right) (t) dc_i(t) \\ &= \int_0^T \left(\Lambda_i^* \circ \frac{dw_i}{dz} \circ c_i \right) (t) dc_i(t) = \int_{c_i(0)}^{c_i(T)} \left(\Lambda_i^* \circ \frac{dw_i}{dz} \right) (z) (dc_i \circ c^{-1}) \\ &= \int_0^{z_i} \left(\Lambda_i^* \circ \frac{dw_i}{dz} \right) (z) dz = \int_0^{z_i} \Lambda_i^*(\dot{w}_i(z)) dz. \end{aligned}$$

This proves the lemma.

Applying the result of Lemma 4 to (32) concludes the proof of Proposition 1.

B Proof of Proposition 2

With reference to the proof of a similar result [25, Theorem 8.4], we can establish the following properties in a completely analogous fashion to complete the proof of the proposition (the proofs are omitted to avoid repetition):

Lemma 5 *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*

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

Lemma 6 *For fixed constants $c > \delta > 0$ and $T > 0$, let*

$$K(c, \delta, T) \triangleq \inf_{(m^T, c^T, q^T)} \int_0^T \left[\sum_{i=1}^N \dot{c}_i(t) A_i^* \left(\frac{\dot{m}_i(t)}{\dot{c}_i(t)} \right) \right] dt$$

subject to (m^T, c^T, q^T) an FSP,
 $\|q(0)\|_\infty \leq c, \|q(t)\|_\infty \geq \delta$ for all $0 \leq t \leq T$.

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

C Proof of Proposition 4

Denote by $\mathcal{C}(\phi_1, \dots, \phi_N)$ the convex hull of the points $(0, \dots, 0)$, $(\phi_1, \dots, 0)$, $(0, \phi_2, \dots, 0)$, \dots , $(0, \dots, \phi_N)$. The right-hand side of (4) is trivially ∞ if either (a) $\lambda \in \mathcal{C}(\phi_1, \dots, \phi_N)$, or (b) any of the ϕ_i is not in the effective domain of its corresponding A_i^* ; we exclude such ϕ_i and λ in the remainder of the proof.

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

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

Here, $\pi(\cdot)$ is used to denote the stationary distribution that the policy π induces, and \mathbb{P}_0^π 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

$$\mathbb{P}^\pi \left[\|q^{(n)}(0)\|_\infty \geq 1 \right] \geq \pi((0, 0, \dots, 0)) \mathbb{P}_0^\pi \left[\max_i u_i^{(n)}(t_n) \geq 1 \right]. \quad (35)$$

For each $i = 1, \dots, N$, since ϕ_i is in the effective domain of its Cramér rate function A_i^* , it follows that there exists $\eta'_i \in \mathbb{R}$ such that $A_i^*(\phi_i) = \eta'_i \phi_i - A_i(\eta'_i)$. 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:

$$\hat{\mathbb{P}}_i(dx) \triangleq \exp[\eta'_i x - A_i(\eta'_i)] \mathbb{P}_i(dx) = \exp[\eta'_i(x - \phi_i) + A_i^*(\phi_i)] \mathbb{P}_i(dx).$$

A standard computation under the tilted measure yields $\hat{\mathbb{E}}_i[R_i(0)] = \phi_i$. As with the approach followed in [23], let $\hat{\mathbb{P}}_0^\pi$ be the measure defined similarly to \mathbb{P}_0^π except that the twisted measures $\{\mathbb{P}_i\}$ 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

$$t_{\min}^{-1} \triangleq \max_i \lambda_i,$$

$$t_{\max}^{-1} \triangleq \min_{\mu \in \mathcal{C}(\phi_1, \dots, \phi_N)} \max_i (\lambda_i - \mu_i).$$

Since by hypothesis the arrival rate λ is outside the closed set $\mathcal{C}(\phi_1, \dots, \phi_N)$, 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 $\mathcal{C}(\phi_1, \dots, \phi_N)$.

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 (35) 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 $\overline{M}_i(k) \triangleq \sum_{l=0}^k (R_i(l) - \phi_i) X_i(l)$ is the (unscaled) “centered” service provided to queue i upto time slot k .

Lemma 7 *Let times t_1 and t_2 , such that $0 < t_1 \leq t_2$, and $\delta > 0$ be fixed. Then,*

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

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 $D \triangleq (R_{\max} + \max_i \phi_i)$, where R_{\max} is the maximum channel rate across all channels in the system. An application of the Azuma-Hoeffding martingale inequality [13] thus gives

$$\hat{\mathbb{P}}_0^\pi \left[\left| \frac{\overline{M}_i(k)}{k} \right| \geq \gamma \right] \leq 2e^{-\frac{k\gamma^2}{2D^2}} \quad (36)$$

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

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

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 7, i.e.

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

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

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

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

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

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

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

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

$$\Rightarrow \hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1, 1 + \frac{D}{n} \right] \right] \geq \hat{\mathbb{P}}_0^\pi [\tau_n = t_n] \geq \frac{\hat{\mathbb{P}}_0^\pi [A_n]}{n(t_{\max} + \epsilon)}. \quad (37)$$

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:

$$\Gamma_n \equiv \Gamma_n(\delta, t_n) \triangleq \left\{ (f'_1, \dots, f'_N) : \sum_i f'_i = 1, f'_i \geq 0, \max_i (\lambda_i t_n - \phi_i f'_i t_n) \in [1 - \delta t_n, 1 + \delta t_n] \right\}.$$

Lemma 8 For all n large enough and $\delta > 0$,

$$\hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1 - \frac{D}{n}, 1 + \frac{D}{n} \right], \frac{c(t_n)}{t_n} \notin \Gamma_n \right] \leq 2Ne^{-\frac{nt_0\delta^2}{8D^2}}.$$

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,

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

Thus,

$$\begin{aligned} & \hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1 - \frac{D}{n}, 1 + \frac{D}{n} \right], \frac{c(t_n)}{t_n} \notin \Gamma_n \right] \leq \hat{\mathbb{P}}_0^\pi \left[\exists i \mid \bar{m}_i(t_n) > \frac{\delta}{2} t_n \right] \\ & \leq \sum_i \hat{\mathbb{P}}_0^\pi \left[\bar{m}_i(t_n) > \frac{\delta}{2} t_n \right] \leq \sum_i 2e^{-\frac{nt_n\delta^2}{8D^2}} \quad (\text{by the Azuma-Hoeffding inequality (36)}) \\ & \leq 2Ne^{-\frac{nt_0\delta^2}{8D^2}}. \end{aligned}$$

Step 4 of 4

We can now finally develop the right-hand side of (35) using the results from the previous steps:

$$\begin{aligned} & \mathbb{P}_0^\pi \left[\max_i u_i^{(n)}(t_n) \geq 1 \right] = \mathbb{P}_0^\pi \left[d^{(n)}(t_n) \geq 1 \right] \\ & \geq \mathbb{P}_0^\pi \left[d^{(n)}(t_n) \geq 1, c^{(n)}(t_n)/t_n \in \Gamma_n \right] \\ & = \mathbb{E}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} \right] \\ & = \hat{\mathbb{E}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} \prod_{l=0}^{nt_n} \exp \left[-\Lambda_{U(l)}^*(\phi'_{U(l)}) - \eta'_{U(l)} \left(R_{U(l)}(l) - \phi'_{U(l)} \right) \right] \right] \\ & = \hat{\mathbb{E}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} \exp \left[-nt_n \left(\sum_i \frac{c_i(t_n)}{t_n} \Lambda_i^*(\phi_i) \right) \right] \exp[-nw(t_n)] \right] \\ & \quad \left(\text{with } w(t_n) \equiv w^{(n)}(t_n) \triangleq \frac{1}{n} W(nt_n) \triangleq \frac{1}{n} \sum_{l=0}^{nt_n} \eta'_{U(l)} \left(R_{U(l)}(l) - \phi'_{U(l)} \right) \right) \\ & = \hat{\mathbb{E}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} \exp \left[-nt_n \left(\sup_{f' \in \Gamma_n} \sum_i f'_i \Lambda_i^*(\phi_i) \right) \right] \exp[-nw(t_n)] \right] \\ & = \exp \left[-nt_n \left(\sup_{f' \in \Gamma_n} \sum_i f'_i \Lambda_i^*(\phi_i) \right) \right] \hat{\mathbb{E}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} e^{-nw(t_n)} \right]. \quad (38) \end{aligned}$$

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

$$\begin{aligned} & \hat{\mathbb{P}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n\}} e^{-nw(t_n)} \right] \geq \hat{\mathbb{P}}_0^\pi \left[\mathbb{1}_{\{d(t_n) \geq 1, c(t_n)/t_n \in \Gamma_n, |w(t_n)| < \zeta\}} e^{-n\zeta} \right] \\ & = e^{-n\zeta} \hat{\mathbb{P}}_0^\pi \left[d(t_n) \geq 1, \frac{c(t_n)}{t_n} \in \Gamma_n, |w(t_n)| < \zeta \right] \\ & \geq e^{-n\zeta} \hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1, 1 + \frac{D}{n} \right], \frac{c(t_n)}{t_n} \in \Gamma_n, |w(t_n)| < \zeta \right]. \end{aligned} \quad (39)$$

We have

$$\begin{aligned} & \hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1, 1 + \frac{D}{n} \right], \frac{c(t_n)}{t_n} \in \Gamma_n, |w(t_n)| < \zeta \right] \\ & \geq \hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1, 1 + \frac{D}{n} \right], \frac{c(t_n)}{t_n} \in \Gamma_n \right] - \hat{\mathbb{P}}_0^\pi [|w(t_n)| \geq \zeta] \\ & \geq \hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1, 1 + \frac{D}{n} \right] \right] - \hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1, 1 + \frac{D}{n} \right], \frac{c(t_n)}{t_n} \notin \Gamma_n \right] - \hat{\mathbb{P}}_0^\pi [|w(t_n)| \geq \zeta]. \end{aligned} \quad (40)$$

By definition and the properties of the twisted distribution $\hat{\mathbb{P}}$, it can be seen that $\{W(k)\}_{k=0,1,\dots}$ 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

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

Using this and the results of Steps 2 and 3, (40) becomes

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

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

$$-\liminf_{n \rightarrow \infty} \frac{1}{n} \log \hat{\mathbb{P}}_0^\pi \left[d(t_n) \in \left[1, 1 + \frac{D}{n} \right], \frac{c(t_n)}{t_n} \in \Gamma_n, |w(t_n)| < \zeta \right] \leq 0. \quad (41)$$

What remains is to bound the first term in the product in (38). By definition, for every $f' \in \Gamma_n$, we have

$$\begin{aligned} & \max_i (\lambda_i - \phi_i f'_i) \leq \frac{1}{t_n} + \delta \\ \Rightarrow & t_n \sup_{f' \in \Gamma_n} \sum_i f'_i \Lambda_i^*(\phi_i) \leq \sup_{f' \in \Gamma_n} \frac{\sum_i f'_i \Lambda_i^*(\phi_i)}{\max_i (\lambda_i - \phi_i f'_i) - \delta} \\ & \leq \sup_{\substack{\sum_i f'_i = 1 \\ f'_i \geq 0}} \frac{\sum_i f'_i \Lambda_i^*(\phi_i)}{\max_i (\lambda_i - \phi_i f'_i) - \delta}. \end{aligned} \quad (42)$$

Applying the conclusions of (39), (41) and (42) to (38), we get

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

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

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

The stationary distribution \mathbb{P}^π induced by the (stabilizing) scheduling policy π forces $\pi((0, 0, \dots, 0)) > 0$, so (35) finally implies

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

which completes the proof of the proposition.

D Proof of Proposition 5

Recall that J_* is the infimum

$$J_* \triangleq \inf_{\substack{T, (m^T, c^T, q^T) \\ 0 \leq t \leq T}} \frac{\sum_{i=1}^N \dot{c}_i(t) \Lambda_i^* \left(\frac{m_i(t)}{c_i(t)} \right)}{\frac{d}{dt} \|q(t)\|_\infty} \quad (43)$$

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 9 *Under the Max-Queue policy, let $s(t) \triangleq \arg \max_{i=1, \dots, N} q_i(t) \subseteq \{1, \dots, N\}$. 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 9,

$$J_* \geq \inf_{S \subseteq \{1, \dots, N\}} \frac{\sum_{i \in S} c'_i \Lambda_i^*(\phi_i)}{w'}, \quad (44)$$

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 $\lambda \notin \mathcal{C}(\phi_1, \dots, \phi_N)$, 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 $S \subseteq \{1, \dots, N\}$, 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 follows that for each such tuple $\phi \in \mathcal{D}_S$, there is a *unique* corresponding tuple c' and hence a *unique* w' . Thus, if we define a map $f^S : \mathcal{D}_S \rightarrow \mathbb{R}^+$ by

$$f^S(\phi) \triangleq \frac{\sum_{i \in S} c'_i \Lambda_i^*(\phi_i)}{w'}, \quad (45)$$

then (44) is just

$$J_* \geq \min_S \inf_{\phi_S \in \mathcal{D}_S} f^S(\phi_S). \quad (46)$$

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

Lemma 10 *Let $S \subseteq \{1, \dots, N\}$ 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. 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 $S = \{1, \dots, N\}$. Denote $\mu_i \triangleq \mathbb{E}[R_i]$. λ is a stabilizable vector of arrival rates, so $\lambda \in \mathcal{C}(\mu_1, \dots, \mu_N)$ (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 [6] 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 $\Lambda_i^*(\cdot)$ [6] forces f^S to be lower-semicontinuous 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 $\hat{\phi}_S \in \mathcal{D}_S$ minimize $f^S(\cdot)$ 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., $\left.\frac{\partial}{\partial \phi_i} f^S(\phi_S)\right|_{\hat{\phi}_S} \geq 0$. From the definition of f^S (45), we can write $\left.\frac{\partial}{\partial \phi_i} f^S(\phi_S)\right|_{\hat{\phi}_S} = \frac{\partial}{\partial \phi_i} \frac{N}{D}$, 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,

$$0 \leq \left.\frac{\partial}{\partial \phi_i} f^S(\phi_S)\right|_{\hat{\phi}_S} = \frac{1}{D^2(\hat{\phi}_S)} \left(D(\hat{\phi}_S) \frac{\partial}{\partial \phi_i} N(\phi_S) - N(\hat{\phi}_S) \frac{\partial}{\partial \phi_i} D(\phi_S) \right) \Big|_{\hat{\phi}_S}. \quad (47)$$

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 (47), this implies (evaluated at $\phi_S = \hat{\phi}_S$)

$$\begin{aligned}
0 &\leq -D(\phi_S) \cdot \frac{\partial D(\phi_S)}{\partial \phi_i} \cdot \sum_{j \in S} \eta'_j + D(\phi_S) \cdot c'_i \phi_i \cdot \frac{\partial \eta'_i}{\partial \phi_i} - N(\phi_S) \frac{\partial D(\phi_S)}{\partial \phi_i} \\
&= -\frac{\partial D(\phi_S)}{\partial \phi_i} \left[D(\phi_S) \cdot \sum_{j \in S} \eta'_j + N(\phi_S) \right] + D(\phi_S) \cdot c'_i \phi_i \cdot \frac{\partial \eta'_i}{\partial \phi_i} \\
&= -\frac{\partial D(\phi_S)}{\partial \phi_i} \left[D(\phi_S) \cdot \sum_{j \in S} \eta'_j + N(\phi_S) \right] + D(\phi_S) \cdot c'_i \phi_i \cdot \frac{\phi_i \frac{\partial \Lambda_i^*(\phi_i)}{\partial \phi_i} - \Lambda_i^*(\phi_i)}{\phi_i^2} \\
&= -\frac{\partial D(\phi_S)}{\partial \phi_i} \left[D(\phi_S) \cdot \sum_{j \in S} \eta'_j + N(\phi_S) \right] - D(\phi_S) c'_i \eta'_i + D(\phi_S) \cdot c'_i \cdot \underbrace{\frac{\partial \Lambda_i^*(\phi_i)}{\partial \phi_i}}_{\leq 0} \\
&\leq -\underbrace{\frac{\partial D(\phi_S)}{\partial \phi_i}}_{\leq 0} \left[D(\phi_S) \cdot \sum_{j \in S} \eta'_j + N(\phi_S) \right] - D(\phi_S) c'_i \eta'_i \\
\Rightarrow \frac{N}{D} &\geq -\frac{c'_i \eta'_i}{\left(\frac{\partial D}{\partial \phi_i}\right)} - \sum_{j \in S} \eta'_j. \tag{48}
\end{aligned}$$

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

$$\begin{aligned}
-\frac{c'_i}{\left(\frac{\partial D}{\partial \phi_i}\right)} &= \phi_i \cdot \sum_{j \in S} \frac{1}{\phi_j} \\
\Rightarrow \frac{N}{D} &\geq \eta'_i \phi_i \cdot \sum_{j \in S} \frac{1}{\phi_j} - \sum_{j \in S} \eta'_j \quad (\text{by (48)}) \\
\Rightarrow f^S(\hat{\phi}_S) &= \frac{N}{D} \geq \left(\max_{i \in S} \hat{\eta}'_i \hat{\phi}_i \right) \cdot \sum_{j \in S} \frac{1}{\hat{\phi}_j} - \sum_{j \in S} \hat{\eta}'_j. \tag{49}
\end{aligned}$$

Now consider any tuple $\{d'_i\}_{i \in S}$ with $d'_i \geq 0$ and $\sum_{i \in S} d'_i = 1$. Let \hat{c}' be the (unique) tuple corresponding to $\hat{\phi}_S$ such that $0 < \lambda_i - \hat{c}'_i \hat{\phi}_i = \lambda_j - \hat{c}'_j \hat{\phi}_j \forall i, j \in S$. 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

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

By (49), 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

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

This completes the proof of the lemma.

Using this lemma, we can finish the proof of the proposition. Let S be a subset of channels that achieves the minimum in (46); 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 $\sum_{i \in S} e'_i \leq 1$, the lemma gives

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

completing the proof.

E Proof of Theorem 4

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 2) implies that the derivatives $\left. \frac{d}{du} \Psi_\alpha(u) \right|_{u=s}$ across all $\alpha \in \mathcal{O}^*$, and $\left. \frac{d}{du} \Psi(u) \right|_{u=s}$, are equal to w' , say. For each such α ,

$$\begin{aligned}
w' &= \sum_{i \in \alpha} e^{\circ q_i(s) + b_i} \underbrace{(\lambda_i(s) - v_i(s))}_{=\lambda_i} \\
&= \langle e^{\circ q(s) + b}, \lambda \rangle_\alpha - \langle e^{\circ q(s) + b}, v(s) \rangle_\alpha \\
&= \langle e^{\circ q(s) + b}, \lambda \rangle_\alpha - \circ \dot{c}_\alpha(s) \langle e^{\circ q(s) + b}, \frac{v(s)}{\circ \dot{c}_\alpha(s)} \rangle_\alpha \\
&= \langle e^{\circ q(s) + b}, \lambda \rangle_\alpha - \circ \dot{c}_\alpha(s) \left[\max_{\eta_\alpha \in V_{\phi_\alpha(s)}} \langle e^{\circ q(s) + b}, \eta_\alpha \rangle_\alpha \right]. \tag{50}
\end{aligned}$$

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

$$\rho_\alpha \triangleq \langle e^{\circ q(s) + b}, \lambda \rangle_\alpha,$$

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

With this, (50) becomes

$$w' \equiv w'(\phi_\alpha(s)) = \rho_\alpha - \circ \dot{c}_\alpha(s) \cdot \xi_\alpha(\phi_\alpha(s)).$$

For fixed $\circ q(s) = q$, the map $\xi_\alpha : \Pi_\alpha \rightarrow \mathbb{R}^+$ is linear and hence continuous. Thus, ξ_α induces a good rate function \tilde{A}_α^* on \mathbb{R}^+ [6], given by

$$\tilde{A}_\alpha^*(\nu'_\alpha) \triangleq \inf \{ A_\alpha^*(\phi_\alpha) : \phi_\alpha \in \Pi_\alpha, \xi_\alpha(\phi_\alpha) = \nu'_\alpha \}.$$

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

$$\begin{aligned}
\frac{\sum_{\alpha \in \mathcal{O}^*} \circ \dot{c}_\alpha(s) A_\alpha^*(\phi_\alpha(s))}{\dot{\Psi}(s)} &= \frac{\sum_{\alpha \in \mathcal{O}^*} \circ \dot{c}_\alpha(s) A_\alpha^*(\phi_\alpha(s))}{\rho_\alpha - \circ \dot{c}_\alpha(s) \cdot \xi_\alpha(\phi_\alpha(s))} \\
&\geq \inf \left\{ \frac{\sum_{\alpha \in \mathcal{O}^*} c'_\alpha \tilde{A}_\alpha^*(\nu'_\alpha)}{w'} \mid w' > 0, \nu'_\alpha \geq 0, c'_\alpha \geq 0, \sum_{\alpha \in \mathcal{O}^*} c'_\alpha = 1, \rho_\alpha - c'_\alpha \nu'_\alpha = w' \quad \forall \alpha \in \mathcal{O}^* \right\}. \tag{51}
\end{aligned}$$

This exactly corresponds to infimizing the function f^S , given in (45), over the corresponding domain \mathcal{D}_S for the case of singleton observable subsets/individual channels. The correspondence becomes clear when, keeping $\circ q$ 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{A}^* is a (good) rate function, we can employ the same arguments as those in the proof of Lemma 10 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 (51), such that the infimum (51) 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

$$\frac{\sum_{\alpha \in \mathcal{O}^*} \hat{c}'_\alpha \tilde{A}_\alpha^*(\hat{\nu}'_\alpha)}{\hat{w}'} \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \tilde{A}_\alpha^*(\hat{\nu}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} (\rho_\alpha - d'_\alpha \hat{\nu}'_\alpha)}. \tag{52}$$

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 $\Lambda_\alpha^*(\hat{\nu}'_\alpha) = \Lambda_\alpha^*(\hat{\phi}'_\alpha)$. 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 (51), (52) and the previous remark, we can write

$$\frac{\sum_{\alpha \in \mathcal{O}^*} \diamond \dot{c}_\alpha(s) \Lambda_\alpha^*(\hat{\phi}'_\alpha(s))}{\hat{\Psi}(\diamond q(s))} \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} (\rho_\alpha - d'_\alpha \cdot \xi_\alpha(\hat{\phi}'_\alpha))}. \quad (53)$$

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

$$\begin{aligned} \rho_\alpha - d'_\alpha \cdot \xi_\alpha(\hat{\phi}'_\alpha) &= \langle e^{\diamond q(s)+b}, \lambda \rangle_\alpha - d'_\alpha \cdot \max_{\eta_\alpha \in V_{\hat{\phi}'_\alpha}} \langle e^{\diamond q(s)+b}, \eta_\alpha \rangle_\alpha \\ &= \langle e^{\diamond q(s)+b}, \lambda \rangle_\alpha - \max_{v_\alpha \in d'_\alpha V_{\hat{\phi}'_\alpha}} \langle e^{\diamond q(s)+b}, v_\alpha \rangle_\alpha \\ &= \min_{v_\alpha \in d'_\alpha V_{\hat{\phi}'_\alpha}} \sum_{i \in \alpha} e^{\diamond q_i(s)+b_i} [\lambda_i - v_{\alpha i}]. \end{aligned} \quad (54)$$

Thanks to the key Lemma 12.2 in [25], we have that there exist

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

$$v_\alpha^* \in \arg \max_{v_\alpha \in d'_\alpha V_{\hat{\phi}'_\alpha}} \langle e^{\diamond q_\alpha^*+b}, v_\alpha \rangle_\alpha$$

such that

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

Using this with (53) and (54) yields

$$\begin{aligned} & \frac{\sum_{\alpha \in \mathcal{O}^*} \diamond \dot{c}_\alpha(s) \Lambda_\alpha^*(\hat{\phi}'_\alpha(s))}{\hat{\Psi}(\diamond q(s))} \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda_\alpha^*(\hat{\phi}'_\alpha)}{\Psi(\diamond q(s)) \cdot \max_{\alpha \in \mathcal{O}^*} l_\alpha} \\ \Rightarrow & \frac{\sum_{\alpha \in \mathcal{O}^*} \diamond \dot{c}_\alpha(s) \Lambda_\alpha^*(\hat{\phi}'_\alpha(s))}{\left[\frac{\hat{\Psi}(\diamond q(s))}{\Psi(\diamond q(s))} \right]} \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} l_\alpha} \\ \Rightarrow & \frac{\sum_{\alpha \in \mathcal{O}^*} \diamond \dot{c}_\alpha(s) \Lambda_\alpha^*(\hat{\phi}'_\alpha(s))}{\hat{\Phi}(\diamond q(s))} \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} l_\alpha} \\ & \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} \max_{i \in \alpha} (\lambda_i - v_{\alpha i}^*)} \\ & \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} \max_{v_\alpha \in d'_\alpha V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha} (\lambda_i - v_{\alpha i})} \\ & \geq \frac{\sum_{\alpha \in \mathcal{O}^*} d'_\alpha \Lambda_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*} \max_{v_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha} (\lambda_i - d'_\alpha v_{\alpha i})}. \end{aligned} \quad (55)$$

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 $A_\alpha^*(\hat{\phi}'_\alpha) = 0$ for such α . We can write,

$$\begin{aligned} & \frac{\sum_{\alpha \in \mathcal{O}} c'_\alpha A_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}, v_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha} (\lambda_i - c'_\alpha v_{\alpha i})} = \frac{\sum_{\alpha \in \mathcal{O}^*} c'_\alpha A_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}, v_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha} (\lambda_i - c'_\alpha v_{\alpha i})} \\ & \leq \frac{\sum_{\alpha \in \mathcal{O}^*} c'_\alpha A_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*, v_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha} (\lambda_i - c'_\alpha v_{\alpha i})} \\ & \leq \frac{\sum_{\alpha \in \mathcal{O}^*} \tilde{c}'_\alpha A_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha \in \mathcal{O}^*, v_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha} (\lambda_i - \tilde{c}'_\alpha v_{\alpha i})}, \quad \text{where } \tilde{c}'_\alpha \triangleq \frac{c'_\alpha}{\sum_{\alpha \in \mathcal{O}} c'_\alpha}. \end{aligned} \quad (56)$$

Putting (55) and (56) together, we have, for our original LFSP, that

$$\begin{aligned} \frac{\sum_{\alpha} \circ \dot{c}_\alpha(s) A_\alpha^*(\hat{\phi}'_\alpha(s))}{\dot{\Phi}(\circ q(s))} & \geq \sup_{\substack{\sum_{\alpha} c'_\alpha = 1 \\ c'_\alpha \geq 0}} \left[\frac{\sum_{\alpha} c'_\alpha A_\alpha^*(\hat{\phi}'_\alpha)}{\max_{\alpha, v_\alpha \in V_{\hat{\phi}'_\alpha}} \max_{i \in \alpha} (\lambda_i - c'_\alpha v_{\alpha i})} \right] \\ & \geq - \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}^\pi \left[\|q^{(n)}(0)\|_\infty \geq 1 \right], \end{aligned} \quad (57)$$

for the stationary measure \mathbb{P}^π of any stabilizing scheduling policy, by Theorem 3. Infimizing (57) over all valid LFSPs and using Proposition 7 yields

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

which finishes the proof.