Large Deviations of Queues Sharing a Randomly Time-varying Server

Alexander L. Stolyar
Bell Labs, Alcatel-Lucent
600 Mountain Ave., 2C-322
Murray Hill, NJ 07974
stolyar@research.bell-labs.com

Submitted July 3, 2006.
Revised: June 26, 2007; December 20, 2007;
April 22, 2008; May 16, 2008

Abstract

We consider a discrete-time model where multiple queues, each with its own exogenous arrival process, are served by a server whose capacity varies randomly and asynchronously with respect to different queues. This model is primarily motivated by the problem of efficient scheduling of transmissions of multiple data flows sharing a wireless channel.

We address the following problem of controlling large deviations of the queues: find a scheduling rule, which is optimal in the sense of maximizing

$$\min_i \left[\lim_{n\to\infty} \frac{-1}{n} \log P\left(a_i Q_i > n\right) \right] \,, \tag{0.1}$$

where Q_i is the length of the i-th queue in a stationary regime, and $a_i > 0$ are parameters. Thus, we seek to maximize the minimum of the exponential decay rates of the tails of distributions of weighted queue lengths $a_i Q_i$. We give a characterization of the upper bound on (0.1) under any scheduling rule, and of the lower bound on (0.1) under the *exponential* (EXP) rule. We prove that the two bounds match, thus proving optimality of the EXP rule. The EXP rule is very parsimonious in that it does not require any "pre-computation" of its parameters, and uses only current state of the queues and of the server.

The EXP rule is not invariant with respect to scaling of the queues, which complicates its analysis in the large deviations regime. To overcome this, we introduce and prove a refined sample path large deviations principle, or *refined Mogulskii theorem*, which is of independent interest.

Key words and phrases: Queueing networks, dynamic scheduling, sample path large deviations principle, refined Mogulskii theorem, local fluid limit, time-varying server, quality of service, exponential (EXP) rule

Abbreviated Title: Queues Sharing a Time-varying Server

AMS Subject Classification: 60K25, 60F10, 90B18

1 Introduction

The model we consider in this paper is motivated primarily by the problem of scheduling transmissions of multiple data users (flows) sharing the same wireless channel (server). As an example, one can think of the following scenario: a wireless access point, or base station, receives data traffic flows destined to several different mobile users, and needs to schedule data transmissions to the users over a shared wireless channel, so that the channel is used efficiently. (Cf. [3, 1, 18] for a more detailed discussion of this scenario.) The distinctive feature of this model, which separates it from more "conventional" queueing models, is the fact that the capacity (service rate) of the channel varies with time randomly and *asynchronously* with respect to different users.

A little more precisely (but still informally), the model is as follows. There are N exogenous input (traffic) flows, which are queued in separate (infinite capacity) buffers, before they can be served by a channel. Time is divided into slots. The channel can serve only one of the flows in one slot. The "aggregate state" of the channel varies randomly from slot to slot. If the channel state in a given slot is m and flow i is chosen for service in this slot, the service rate is $\mu_i^m \geq 0$, i.e., μ_i^m customers (bits of data) of flow i are served (transmitted) and leave the system. This and related models received a significant amount of attention in recent years (cf. [14] for an overview). It is well known that efficient scheduling rules cannot be "channel state oblivious." However, it is also known that large classes of rather "parsimonious" algorithms, making scheduling decisions based only on the current channel state and current queue lengths (and/or current head-of-the-line queueing delays) information can in fact achieve certain notions of efficiency. For example, MaxWeight-type algorithms (cf. [2] and references therein) and the Exponential (EXP) algorithm [10] are *throughput optimal* in the sense that they ensure stochastic stability of the queues as long as such is feasible at all, under any rule. Also, both MaxWeight and EXP rules exhibit optimal behavior under heavy traffic conditions (see [14, 11]).

In this paper we would like to address the following issue. Suppose we want to find a scheduling algorithm (rule), or queueing discipline, under which the following Quality-of-Service condition is satisfied:

$$P\{Q_i > B_i\} \leq \delta_i \ , \ i = 1, \dots, N \ , \tag{1.1}$$

where Q_i is the steady state queue length for flow i, $B_i > 0$ is a predefined threshold, and δ_i is the maximum acceptable probability of queue length exceeding the threshold. (This problem appears in a variety of applications, cf. [4, 12, 13] for a further discussion and reviews.)

If the thresholds B_i are "large," then conditions (1.1) can be "approximately" replaced by the following asymptotic - "tail" - conditions

$$\beta(Q_i) \geq a_i \ , \ i = 1, \ldots, N \ , \tag{1.2}$$

where we use the notation

$$\beta(X) \doteq \lim_{n\to\infty} -\frac{1}{n} \log P(X > n) \tag{1.3}$$

for the exponential decay rate of the tail of the distribution of random variable X (assuming the above limit exists), and

$$a_i = -\log(\delta_i)/B_i.$$

This is precisely what we will do in this paper. *We consider the problem of finding a scheduling rule such that the tail conditions (1.2) are satisfied for some fixed set of positive parameters* a_i.

This problem in turn is equivalent to solving the following *optimization* problem

$$\text{maximize} \quad \min_{i=1,\ldots,N} a_i^{-1} \beta(Q_i) \ , \tag{1.4}$$

where the maximization is over all scheduling disciplines. Indeed, a discipline satisfying (1.2) exists if and only if the maximum in (1.4) is 1 or greater (and the maximum is attained, to be precise). Finally, if we denote by

$$Q_* \doteq \max_i a_i Q_i$$

the *maximal weighted queue length*, and observe that $\min_i a_i^{-1}\beta(Q_i) = \beta(Q_*)$, we see that the problem (1.4) is equivalent to

$$\text{maximize} \quad \beta(Q_*) \ . \tag{1.5}$$

To summarize, *we want to find a scheduling rule solving problem (1.5), i.e. a rule maximizing the exponential decay rate of the tail of the distribution of the maximal weighted queue length* Q_*, *with some fixed "weights"* $a_i > 0$.

In the case when the channel is *not* time-varying, i.e., there is only one channel state and therefore the (potential) service rates μ_i are constant, our model essentially fits into the framework of [12], where, in particular, it is proved that an extremely simple rule always choosing for service the queue maximizing $a_i Q_i$ is an optimal solution to problem (1.4). (This result was extended in [13] to a queueing network setting.) However, for our model, where the channel *is* time-varying, the above simple rule *cannot possibly be optimal for problem (1.4)*, because it ignores the current state of the channel; moreover, except for degenerate cases, this rule is not even throughput-optimal - it can make queues unstable in cases when stability (under a different rule) is feasible. The main goal of this paper is to establish optimality of the EXP rule for the problem (1.4). The EXP rule is defined as follows:

$$\text{When the channel is in state } m, \text{ serve flow maximizing } \mu_i^m \exp\left(\frac{a_i Q_i}{1+\overline{Q}^\eta}\right), \tag{1.6}$$

where $\overline{Q} \doteq (1/N)\sum_i a_i Q_i$, and $\eta \in (0,1)$ is a fixed parameter.

Problems like (1.4) are naturally approached using Large Deviations (LD) theory techniques. It is well known in LD theory that, roughly speaking, the value of $\beta(Q_*)$ under a given scheduling rule is determined by a "most likely path" for the process $Q_*(t)$ to reach level n, starting from 0. (See the definition of $\beta(\cdot)$ in (1.3).) Or, equivalently, this is a most likely path for a "fluid-scaled" process $(1/n)Q_*(nt)$ to reach level 1. In turn, the likelihoods of such rescaled paths are determined by a sample path large deviations principle (Mogulskii theorem) for the sequence of fluid-scaled "driving processes" - namely, input flow and channel state processes, as $n \to \infty$. (If the value of the corresponding LD rate function of a path - or path "cost" - is c, then the "probability" of the path is "approximately" e^{-cn}, when n is large.)

One of the difficulties in the LD analysis of the EXP rule is that the "standard" sample path large deviations principle (SP-LDP) is not sufficient for "keeping track" of the path costs. The basic reason for this is that *EXP rule is not asymptotically invariant with respect to scaling of queue lengths.* Informally, an "asymptotically scaling-invariant" rule is such that, when queue lengths are large, a scaling of all queue lengths by the same factor at any given time, (roughly speaking) does not change the scheduling choice. (An example of asymptotically scaling-invariant rule is a MaxWeight-type algorithm, choosing for service a flow i maximizing $c_i Q_i^\kappa \mu_i^m$, where κ and all c_i are arbitrary positive parameters. A slightly more general rule, maximizing $[c_i Q_i^\kappa + d_i]\mu_i^m$, where d_i's are additional parameters, is also asymptotically scaling-invariant.)

Fluid scaling is the "relevant" one to study the dynamics of the queue lengths under an asymptotically scaling-invariant rule in an (unscaled) time interval of the order of $O(n)$ (because rescaling of queue lengths by $1/n$, for any n, "preserves the information" on which scheduling choices are made), and a standard SP-LDP gives the likelihood of trajectories under this scaling. In contrast, the EXP rule is not asymptotically scaling-invariant, as seen from the expression in (1.6). Even if ultimately we are interested in the dynamics of the queue lengths under EXP rule over an interval of the order $O(n)$, the "relevant" time and space scale which determines such dynamics is of the order $O(n^\eta)$. (The value of $\overline{Q}$ is "typically" $O(n)$. Therefore, the differences of the order $O(n^\eta)$ between weighted queue lengths $a_i Q_i$ result in the order $O(1)$ ratios of the exponent terms in (1.6) for different flows i. But, these ratios are what determine the scheduling choices.) Consequently, we need the likelihoods of (unscaled) trajectories over order $O(n^\eta)$ time intervals; fluid scaling, however, does not "preserve" this information. To resolve this difficulty, we introduce and prove what can be called a "refined" SP-LDP, or a *refined Mogulskii theorem* (RMT). Using RMT we introduce the notions of a *generalized fluid sample path* (GFSP) and its *refined cost*. (Roughly speaking, the refined cost of a GFSP "takes into account" the behavior of (unscaled) process trajectories on time scales that are "finer" than $O(n)$.) We show that the likelihood of building large value of Q_* under EXP rule can be given in terms of GFSP refined costs.

Our RMT result (Theorem 7.1) and the notions of GFSP and its refined cost are generic and are of independent interest. In particular, as the above discussion demonstrates, they are instrumental in LD analysis of scheduling rules that are not scaling-invariant.

The **main results** of the paper are as follows. We prove the upper bound $\beta(Q_*) \le J_*$, which holds under any scheduling rule, where J_* is defined in terms of lowest cost "simple" (linear) paths to raise Q_*. The proof of this upper bound involves only a standard Mogulskii theorem for the sequence of fluid-scaled input flow and channel state processes. We introduce and prove a refined Mogulskii theorem, and introduce the related notion of GFSP. We then give the lower bound $\beta(Q_*) \ge J_{**}$, which holds for the EXP rule, where J_{**} is defined in terms of the lowest refined cost of a GFSP to raise Q_*. Finally, for the EXP rule, we prove that the lower and upper bounds on $\beta(Q_*)$ match, that is $\beta(Q_*) = J_{**} = J_*$, thus proving that the EXP rule is indeed an optimal solution to problem (1.4).

Previous work on the large deviations regime for queues served by a time-varying server includes [19], which contains results for a MaxWeight-type rule (maximizing $Q_i\mu_i^m$) in a symmetric model. ("Symmetric" means: all input flows have equal rate and are non-random; channel state $m = (m_1, \ldots, m_N)$ is a direct product of N independent and identically distributed channel states m_i of the individual flows.) The optimality problem (1.4) is not addressed in [19], and the analysis relies in an essential way on the symmetry assumptions.

The rest of the paper is organized as follows. In Section 2 we introduce basic notations, definitions, conventions used in the paper. The system model, formal definition of the EXP rule, and our main results (Theorem 3.2) characterizing the bounds on $\beta(Q_*)$ under an arbitrary rule and EXP rule, and proving EXP optimality, are given in Section 3. The necessary definitions of a sequence of scaled processes and a standard SP-LDP (Mogulskii theorem) are presented in Sections 4 and 5, respectively. In Section 6 we prove the bound $\beta(Q_*) \le J_*$ (Theorem 3.2(i)) for any scheduling discipline. A refined Mogulskii theorem is formulated and proved in Section 7. Section 8 contains the definition of a GFSP and proof of the bound $\beta(Q_*) \ge J_{**}$ (Theorem 3.2(ii)) under EXP rule. A part of Theorem 3.2(ii) proof is postponed until Section 10, because it requires the notion of *local fluid sample path* (LFSP), developed in Section 9; LFSPs describe the dynamics of the queue lengths in time intervals of the order $O(n^\eta)$ and serve as a key tool for deriving the bound J_{**} under EXP rule. Finally, in Sections 11 and 12, using again the LFSP construction, we prove that, in fact, $J_{**} = J_*$ and therefore EXP rule is optimal (Theorem 3.2(iii)).

2 Basic notation and definitions

We denote by $\mathbb{R}$, $\mathbb{R}_+$ and $\mathbb{R}_{++}$ the sets of real, real non-negative and real positive numbers, respectively. The corresponding k-times product spaces are $\mathbb{R}^k$, $\mathbb{R}^k_+$ and $\mathbb{R}^k_{++}$. For vectors

$a, b \in I\!\!R^k$, the scalar product is $a \cdot b \doteq \sum_i a_i b_i$, the norm of b is $\|b\| \doteq (b \cdot b)^{1/2}$; we also use non-standard notations $e^b \doteq (e^{b_1}, \ldots, e^{b_k})$ and $a \times b \doteq (a_1 b_1, \ldots, a_k b_k)$.

We denote minimum and maximum of two real numbers ξ_1 and ξ_2 by $\xi_1 \wedge \xi_2$ and $\xi_1 \vee \xi_2$, respectively; and by $\lfloor \xi \rfloor$ and $\lceil \xi \rceil$ the integer part and the ceiling of a real number ξ, respectively. The infimum of a function over an empty set is interpreted as $+\infty$.

Let $\mathcal{D}$ be the space of RCLL functions (i.e. right continuous functions with left limits) defined on $[0, \infty)$ and taking values in $I\!\!R$. Unless otherwise specified, we assume $\mathcal{D}$ is endowed with the topology of uniform convergence on compact sets (u.o.c.). As a measurable space, we always assume that $\mathcal{D}$ is endowed with the σ-algebra generated by the cylinder sets. By $\mathcal{A}$ we denote the subset of absolutely continuous functions in $\mathcal{D}$, and by $\mathcal{A}_0 \subset \mathcal{A}$ the subset of functions $h(\cdot)$ with $h(0) = 0$. For any function space S, and any $0 \leq c < d$, $\zeta_c^d S$ denotes the space of functions in S with the domain "truncated" to $[c, d]$. The subspaces and spaces with truncated domains inherit the topology and σ-algebra of $\mathcal{D}$. Given any space S, we assume that the k times product space S^k has the product topology and product σ-algebra defined in the natural way.

For any $s \geq 0$ and $h = (h_1, \ldots, h_k) \in \mathcal{D}^k$ [or $\zeta_c^d \mathcal{D}^k$]], we define the norm

$$||h||_s \doteq \max_{i=1,\ldots,k} \sup_{t \leq s} |h_i(s)|.$$

Thus the u.o.c. convergence in $\mathcal{D}^k$ [or $\zeta_c^d \mathcal{D}^k$] is equivalent to convergence in norm $|| \cdot ||_s$ for all $s > 0$. We define the scaling operator Γ^c, $c > 0$, for $h \in \mathcal{D}^k$ as follows:

$$(\Gamma^c h)(t) \doteq \frac{1}{c} h(ct) \ . \tag{2.1}$$

For a function $h \in \mathcal{D}$, we define the domain truncation operator ζ_c^d, for $0 \leq c < d$, in the natural way:

$$\zeta_c^d h \in \zeta_c^d \mathcal{D} \quad \text{and} \quad (\zeta_c^d h)(t) = h(t), \ c \leq t \leq d.$$

For $h \in \mathcal{D}$, and $0 \leq c < d$, we also define operator $\bar{\zeta}_c^d$ (which is a simultaneous domain truncation and shift, as well as recentering) as follows:

$$\bar{\zeta}_c^d h \in \zeta_0^{d-c} \mathcal{D} \quad \text{and} \quad (\bar{\zeta}_c^d h)(t) = h(c+t) - h(c).$$

For a set of functions, operators ζ_c^d and $\bar{\zeta}_c^d$ are applied componentwise.

We will write simply 0 for the zero element of $I\!\!R^k$ and for zero functions taking values in $\mathcal{D}^k$ [or $\zeta_c^d \mathcal{D}^k$].

Let $\mathbf{\Omega} \doteq (\Omega, \mathcal{F}, P)$ be a probability space. We assume that $\mathbf{\Omega}$ is large enough to support all the independent random processes that we use in the paper. Typically, we follow the convention of using bold font for stochastic processes and Roman font for deterministic functions, including realizations of random processes. Given any subset B of a topological space, we use $\bar{B}$ and B° to denote its closure and interior respectively. The following is the standard definition of a large deviation principle [6, p.5].

Definition 2.1 (LDP) *Let $\mathcal{X}$ be a topological space and $\mathcal{B}$ a σ-algebra on $\mathcal{X}$ (which is not necessarily the Borel σ-algebra). A sequence of random variables $\{\mathbf{X}_n\}$ on Ω taking values in $\mathcal{X}$ is said to satisfy the LDP with good rate function I if for all $B \in \mathcal{B}$,*

$$\limsup_{n\to\infty} \frac{1}{n} \log P(\mathbf{X}_n \in B) \leq - \inf_{x\in\bar{B}} I(x),$$

and

$$\liminf_{n\to\infty} \frac{1}{n} \log P(\mathbf{X}_n \in B) \geq - \inf_{x\in B^\circ} I(x),$$

where $I : \mathcal{X} \to I\!\!R_+ \cup \{\infty\}$ is a function with compact level sets.

3 The model and main results

3.1 The model

The system has N input flows, consisting of discrete *customers*, which need to be served by a single *channel* (or server). We will denote by N both the set of flows $\{1, \ldots, N\}$ and its cardinality. Each flow has its own queue where customers wait for service. (Sometimes, we use terms "flow" and "queue" interchangeably.)

The system operates in discrete time. A time interval $[t, t+1)$, with $t = 0, 1, 2, \ldots$, we will call the *time slot* t. In each time slot the channel can be in one out of the finite set $M = \{1, \ldots, M\}$ of *channel states*, and it can pick one of the flows for service. If in a given time slot the channel is in state $m \in M$ and flow $i \in N$ is chosen for service, then the integer number $\mu_i^m \geq 0$ of customers are served from the corresponding queue i (or the entire queue i content, if it is less than μ_i^m). Thus, associated with each channel state $m \in M$ is the fixed vector of service rates $(\mu_1^m, \ldots, \mu_N^m)$.

The channel state $m(t)$ in each time slot t is drawn independently according to some probability distribution $\pi = (\pi^1, \ldots, \pi^M)$. Without loss of generality, we can and will assume that $\pi_m > 0$ for all states m.

Denote by $A_i(t)$ the number of type i customers that arrived in time slot $t = 1, 2, \ldots$. We will adopt a convention that the customers arriving in slot t are immediately available for service in this slot. We will assume that all arrival processes are mutually independent, each sequence $A_i(t),\ t = 1, 2, \ldots$, is i.i.d., with finite exponential moments

$$Ee^{\theta A_i(1)} < \infty,\ \forall\theta \geq 0, \quad \forall i. \tag{3.1}$$

Let us denote by $\bar{\lambda}_i \doteq EA_i(1),\ i = 1, \ldots, N$, the mean arrival rate for flow i, and assume that $\bar{\lambda}_i > 0$ for all i.

The random process describing the behavior of the system is

$$Q(t) = (Q_i(t),\ i = 1, \ldots, N), \quad t = 0, 1, 2, \ldots$$

where $Q_i(t)$ is the type i queue length at time t.

For some (but not all) results in the paper we will need the following additional assumption:

$$A_i(1) < C < \infty, \quad \forall i. \tag{3.2}$$

3.2 Scheduling Rules. Stability

A *scheduling rule*, or a *queueing discipline*, is a rule that determines which flow to pick for service in a given time slot t, depending in general on the entire history of the process up to time t.

If we denote by $D_i(t)$, the number of type i customers served in the time slot $t-1$, then according to our conventions, for each $t = 1, 2, \ldots$,

$$Q_i(t) = Q_i(t-1) - D_i(t) + A_i(t), \ \forall i. \tag{3.3}$$

Note that $D_i(t) = \min\{Q_i(t-1), \mu_i^{m(t-1)}\}$ for the flow i chosen for service in slot $t-1$, and $D_i(t) = 0$ for all other flows (because in our model only one flow can be served in a slot).

If a scheduling rule is such that it picks a flow to be served in a given time slot t depending only on the current queue length vector $Q(t)$ and current channel state $m(t)$, then clearly $(Q(t), m(t)), t = 0, 1, 2, \ldots$, is a Markov chain with countable state space. (The EXP rule defined later is of this type.) We say that the system under a given scheduling rule of this type is *stable* if the Markov chain has a finite subset of states which is reachable from any other state with probability 1, and each state within the subset is positive recurrent. Stability implies existence of a stationary probability distribution. (If the Markov chain happens to be irreducible, stability is equivalent to ergodicity, and the stationary distribution is unique.)

Suppose a stochastic matrix $\phi = (\phi_{mi}, \ m \in M, \ i = 1, \ldots, N)$ is fixed, which means that $\phi_{mi} \geq 0$ for all m and i, and $\sum_i \phi_{mi} = 1$ for every m. Let Φ be the set of all such stochastic matrices ϕ. Given $\phi \in \Phi$ we define the vector $v = (v_1, \ldots, v_N) = v(\phi)$ as follows:

$$v_i = \sum \pi^m \phi_{mi} \mu_i^m \ , \ i \in N. \tag{3.4}$$

If each component ϕ_{mi} of matrix ϕ is interpreted as a "long-term" average fraction of time slots when flow i is chosen for service, out of those slots when the channel state is m, then $v(\phi)$ is simply the vector of average service rates which will be "given" to the flows. The set

$$V \doteq \{w \in R_+^N \mid w \leq v(\phi) \text{ for some } \phi \in \Phi\}$$

is called system *(service) rate region.*

It is well known (cf. [14] and references therein) that the condition $\bar{\lambda} \in V$ is necessary for stability. Throughout this paper we assume a slightly stronger condition:

$$\bar{\lambda} < v^* \quad \text{for some} \quad v^* \in V. \tag{3.5}$$

3.3 Exponential Scheduling Rule

Let a set of positive parameters $a_1, \ldots, a_N$ and $\eta \in (0,1)$ be fixed. The following scheduling rule is called Exponential [10], or EXP: it chooses for service in time slot t a single queue

$$i \in i(Q(t)) = \arg\max_i c_i \;\mu_i(t) \;\exp\left(\frac{a_i Q_i(t)}{c + [\overline{Q}(t)]^\eta}\right) \;, \tag{3.6}$$

where $\mu_i(t) \equiv \mu_i^{m(t)}$, $\overline{Q}(t) \doteq (1/N)\sum_i a_i Q_i(t)$, and $c, c_1, \ldots, c_N$, are some additional positive parameters. (Ties are broken in an arbitrary, but a priori fixed way, for example in favor of the smallest index within the set $i(Q(t))$.)

Proposition 3.1 [10] *If condition (3.5) holds, the system under the EXP rule is stable.*

Proposition 3.1 says that the EXP rule is *throughput optimal* in the sense that it makes the system stable as long as, essentially, stability is feasible at all.

In the rest of the paper, to simplify exposition, we assume that parameters $c, c_1, \ldots, c_N$ are all equal to 1. (Setting these parameters to arbitrary values does not affect main results, and it does not affect the proofs in any essential way.)

3.4 Main results

The function $Q_*(t) \doteq \max_i a_i Q_i(t)$ of the state $Q(t)$ will be called *maximal weighted queue length*. (The corresponding *random processes* are denoted by $\mathbf{Q}(t)$ and $\mathbf{Q}_*(t)$, $t = 0, 1, 2, \ldots$.) It will be convenient to extend the domain of $Q(\cdot)$ and $Q_*(\cdot)$ (as well as other functions introduced later in the paper), which are naturally defined in discrete time, to continuous time t by adopting the convention that the functions are constant within each time slot $[k, k+1)$, where k is integer. Now we are in position to formulate our main result.

Theorem 3.2 *Suppose condition (3.5) is satisfied. Then, the following holds.*

(i) There exists $T^0 \in (0,\infty)$ such that for any scheduling rule and any $t > T^0$, we have the following lower bound:

$$\liminf_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\mathbf{Q}_*(nt) > 1\right) \geq -J_* \;, \tag{3.7}$$

where $J_* > 0$ is defined and explained later in Section 6. [Additional condition (3.2) is not *required.]*

(ii) Suppose additional condition (3.2) holds. Consider the system under the EXP scheduling rule and the Markov chain $\mathbf{Q}(\cdot)$ being in a stationary regime (which exists by Proposition 3.1). Then, we have the following upper bound:

$$\limsup_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\mathbf{Q}_*(0) > 1\right) \leq -J_{**} \;, \tag{3.8}$$

where J_{**}, $0 \le J_{**} \le J_*$, is defined and explained later in Section 8 (see (8.7)).

(iii) Moreover, under the conditions of (ii),

$$J_{**} = J_* \tag{3.9}$$

and, therefore, under the EXP rule

$$\lim_{n\to\infty} \frac{1}{n} \log P\left(\frac{1}{n}\mathbf{Q}_*(0) > 1\right) = -J_* \ . \tag{3.10}$$

Theorem 3.2(iii) shows that the EXP rule is optimal in that it maximizes the exponential decay rate of the stationary distribution of the maximal weighted queue length $\mathbf{Q}_*(\cdot)$. We remark that the upper bound characterization given in Theorem 3.2(ii) is quite generic and can be used for other non-scaling-invariant scheduling rules, not just EXP; as such, we believe it is of independent interest.

4 Extended description of the system process. Sequence of fluid-scaled processes

As the formulation of Theorem 3.2 suggests (and is typical for this type of large deviations results for queueing systems), its proof involves considering a sequence of "fluid-scaled" versions of the queue length process $\mathbf{Q}$, namely the processes $\Gamma^n\mathbf{Q} = ((1/n)\mathbf{Q}(nt),\ t \ge 0)$, for $n = 1, 2, \ldots$. In this section we define this sequence formally. But first, we need to introduce additional functions associated with the system evolution.

For $t \ge 0$ let

$$F_i(t) \doteq \sum_{k=1}^{\lfloor t \rfloor} A_i(k) \quad \text{and} \quad \hat{F}_i(t) \doteq \sum_{k=1}^{\lfloor t \rfloor} D_i(k) \tag{4.1}$$

denote the total number of flow i customers, respectively arrived to and departed from the system by (and including) time t, that is in the time slots $1 \le k \le \lfloor t \rfloor$. (Recall our convention, introduced in Section 3.4, that we extend the domain of discrete time processes to continuous time $t \ge 0$.) Also, denote by $G_m(t)$ the total number of time slots $0 \le k \le \lfloor t-1 \rfloor$ when the channel was in state m; and by $\hat{G}_{mi}(t)$ the number of time slots $0 \le k \le \lfloor t-1 \rfloor$ when the server state was m and flow i was chosen for service.

The following set of functions describes the evolution of the system in time interval $[0, \infty)$:

$$(Q, Q_*, F, \hat{F}, G, \hat{G}),$$

where

$$Q = (Q(t) = (Q_1(t), \ldots, Q_N(t)), \quad t \ge 0),$$

$$Q_* = (Q_*(t) \equiv \max_i a_i Q_i(t), \quad t \ge 0),$$

$$F = (F(t) = (F_1(t), \dots, F_N(t)), \quad t \geq 0),$$
$$\hat{F} = (\hat{F}(t) = (\hat{F}_1(t), \dots, \hat{F}_N(t)), \quad t \geq 0),$$
$$G = ((G_m(t),\ m \in M),\ t \geq 0)\ ,$$
$$\hat{G} = ((\hat{G}_{mi}(t),\ m \in M,\ i \in N),\ t \geq 0).$$

The set of functions $(Q, Q_*, F, \hat{F}, G, \hat{G})$ clearly has redundancies. The entire set is uniquely determined by the initial state $Q(0)$, the realizations F and G of the input flow and channel state processes, which "drive" the system, and the realization $\hat{G}$, which determines the scheduling choices. In particular, the following basic relations (implied by (3.3), (4.1), and the definitions of $G_m(\cdot)$ and $\hat{G}_{mi}(\cdot)$) hold:

$$Q_i(t) = Q_i(0) + F_i(t) - \hat{F}_i(t),\ t \geq 0, \forall i, \tag{4.2}$$

$$G_m(t) = \sum_i \hat{G}_{mi}(t),\ t \geq 0, \forall m. \tag{4.3}$$

Also, from (3.3) and the observation following it, we see that, if for all t in an interval $[t_1 - 1, t_2]$ we have $Q_i(t) > \max_m \mu_i^m$ (which means that the service provided to queue i is not "wasted"), then

$$\hat{F}_i(t_2) - \hat{F}_i(t_1) = \sum_m \mu_i^m [\hat{G}_{mi}(t_2) - \hat{G}_{mi}(t_1)]. \tag{4.4}$$

In what follows, we will use bold font $(\mathbf{Q}, \mathbf{Q}_*, \mathbf{F}, \hat{\mathbf{F}}, \mathbf{G}, \hat{\mathbf{G}})$ when we view this set of functions as a random process, and use Roman font when we view it as a deterministic sample path.

For each index $n = 1, 2, \dots$, consider a (stochastically equivalent) version of our system, and denote by $(\mathbf{Q}^{(n)}, \mathbf{Q}_*{}^{(n)}, \mathbf{F}^{(n)}, \hat{\mathbf{F}}^{(n)}, \mathbf{G}^{(n)}, \hat{\mathbf{G}}^{(n)})$ the corresponding process. The corresponding sequence of fluid-scaled processes is defined as

$$(\mathbf{q}^{(n)}, \mathbf{q}_*{}^{(n)}, \mathbf{f}^{(n)}, \hat{\mathbf{f}}^{(n)}, \mathbf{g}^{(n)}, \hat{\mathbf{g}}^{(n)}) \doteq \Gamma^n(\mathbf{Q}^{(n)}, \mathbf{Q}_*{}^{(n)}, \mathbf{F}^{(n)}, \hat{\mathbf{F}}^{(n)}, \mathbf{G}^{(n)}, \hat{\mathbf{G}}^{(n)}),$$

with $n = 1, 2, \dots$.

5 Sample path large deviations principle: Mogulskii Theorem

The sequence of processes $(\mathbf{f}^{(n)}, \mathbf{g}^{(n)})$ is known to satisfy a sample path LDP, described in this section.

Assumption (3.1) on the input flows implies (cf. Chapter 2.2.1 in [6]) that, for each i, function

$$L_i(\xi) \doteq \sup_{\theta \geq 0} [\theta\xi - \log E e^{\theta A_i(1)}],\ \xi \geq 0,$$

which is the large deviations rate function of the sequence $(1/n)[A_i(1) + \ldots + A_i(n)]$, has the following properties. $L_i(\cdot)$ is a convex lower semi-continuous function on $[0, \infty)$, taking values in $[0, +\infty]$, attaining its unique minimum 0 at $\bar{\lambda}_i$, i.e.

$$L_i(\bar{\lambda}_i) = 0, \quad \text{and} \quad L_i(\xi) > 0 \text{ for } \xi \neq \bar{\lambda}_i,$$

and it is superlinear at infinity, i.e.

$$L_i(\xi)/\xi \to \infty, \quad \xi \to \infty.$$

(We adopt the convention that $L_i(\xi) = +\infty$ for $\xi < 0$.) In particular, if $A_i(1)$ is bounded, then $L_i(\cdot)$ is finite and continuous in $[C_i^{min}, C_i^{max}]$, where C_i^{min} and C_i^{max} are the minimum and maximum possible values, and is $+\infty$ elsewhere; if $A_i(1)$ is unbounded, $L_i(\cdot)$ is finite and continuous in $[C_i^{min}, \infty)$.

For a vector $y \in R^N$ we will use notation

$$L_{(f)}(y) \doteq \sum_i L_i(y_i).$$

(Subscript (f) indicates that this is the rate function associated with input flows.)

The relative entropy of a probability distribution $\gamma = (\gamma_1, \ldots, \gamma_M)$ with respect to the distribution π we denote by

$$L_{(g)}(\gamma) \doteq \sum_{m \in M} \gamma_m \log \frac{\gamma_m}{\pi_m}.$$

According to Sanov theorem (cf. Theorem 2.1.10 in [6]), $L_{(g)}(\cdot)$ is the large deviations rate function for the sequence of empirical distributions of the channel state over n trials (with $n \to \infty$). Function $L_{(g)}(\cdot)$ is (finite) continuous and convex on the simplex of probability distributions γ; we adopt the convention that $L_{(g)}(\cdot)$ is defined on R^M and is $+\infty$ outside the above simplex.

For a pair (f, g) of vector-functions $f \in \mathcal{D}^N$ and $g \in \mathcal{D}^M$, its *cost* $J_t(f, g)$ in time interval $[0, t]$ is defined as

$$J_t(f,g) \doteq \begin{cases} \int_0^t [L_{(f)}(f'(s)) + L_{(g)}(g'(s))]ds & \text{if } \zeta_0^t(f,g) \in \zeta_0^t \mathcal{A}_0^{N+M}, \\ \infty & \text{otherwise}, \end{cases} \tag{5.1}$$

More generally, if the functions f and g have a bounded domain $[0, d]$, that is $(f, g) \in \zeta_0^d \mathcal{D}^{N+M}$, the cost $J_t(f, g)$ is still defined by (5.1), as long as $t \leq d$.

The following is (a form of) Mogulskii theorem (cf. Theorem 5.1.2 in [6]).

Proposition 5.1 *Consider a sequence of scaled processes $(\mathbf{f}^{(n)},\ \mathbf{g}^{(n)})$, $n = 1, 2, \ldots$, as defined in Section 4. Then, for every $c \geq 0$ and $t \geq 0$, the sequence of processes $\bar{\zeta}_c^{c+t}(\mathbf{f}^{(n)},\ \mathbf{g}^{(n)})$ satisfies the LDP with good rate function $J_t(\cdot)$. In more detail, for any measurable $B \subseteq \zeta_0^t \mathcal{D}^{N+M}$, we have the following asymptotic (respectively lower and upper) bounds:*

$$\liminf_{n \to \infty} \frac{1}{n} \log P(\bar{\zeta}_c^{c+t}(\mathbf{f}^{(n)},\ \mathbf{g}^{(n)}) \in B) \geq -\ \inf\{J_t(h) \mid h \in B^\circ\}\ , \tag{5.2}$$

$$\limsup_{n \to \infty} \frac{1}{n} \log P(\bar{\zeta}_c^{c+t}(\mathbf{f}^{(n)},\ \mathbf{g}^{(n)}) \in B) \leq -\ \inf\{J_t(h) \mid h \in \bar{B}\}\ . \tag{5.3}$$

6 LD lower bound under any scheduling rule. Proof of Theorem 3.2(i)

6.1 Simple trajectories to raise maximal weighted queue length

Let $\gamma = \{\gamma_m, m \in M\}$ be some ("twisted") probability distribution on the set of channel states, not necessarily equal to the distribution π. We denote by V_γ the corresponding "twisted" rate region, defined the same way as V but with π replaced by γ. (Thus $V = V_\pi$.) In addition, for every non-zero subset $N' \subseteq N$, we denote by $V_\gamma(N')$ the projection of V_γ onto the corresponding subspace $R^{|N'|}$, where $|N'|$ is the cardinality of N'. We denote by $V_\gamma^*(N')$ the subset of maximal elements of $V_\gamma(N')$, that is

$$V_\gamma^*(N') \doteq \{v \in V_\gamma(N') \mid v \le w \in V_\gamma(N') \text{ implies } w = v\}.$$

For a fixed non-zero subset $N' \subseteq N$, consider pairs of a distribution γ and a vector $\lambda = \{\lambda_i,\ i \in N'\}$ such that there exists a vector $\mu = \{\mu_i,\ i \in N'\} \in V_\gamma^*(N')$ for which the following condition holds:

$$a_i(\lambda_i - \mu_i) = \ell > 0, \quad \forall i \in N'.$$

(Note that if such a vector μ exists, it is unique, because this is the point where the ray emanating from point λ in the direction given by $\{-1/a_i,\ i \in N'\}$, hits the region $V_\gamma(N')$.) Let us denote

$$J_*(N') \doteq \inf \frac{L_{(g)}(\gamma) + \sum_{i \in N'} L_i(\lambda_i)}{\ell},$$

where the inf is taken over all pairs of γ and λ, as specified above. Finally, we define

$$J_* = \min_{N' \subseteq N,\ N' \neq \emptyset} J_*(N').$$

We now give the interpretation of the above definitions. Let $N' = N$ for simplicity. Consider the process with large index n on a (large) time interval $[0, nt]$ for some fixed t. Suppose the empirical distribution of the channel states in this interval is a "twisted" distribution γ, possibly different from π. Moreover, we assume that the fluid-scaled channel state process trajectory is "close to" linear: $g^{(n)}(s) \approx g(s) \equiv \gamma s,\ 0 \le s \le t$. Suppose also that the fluid-scaled input flow trajectory is 'close to" linear: $f^{(n)}(s) \approx f(s) \equiv \lambda s,\ 0 \le s \le t$, for some vector λ not necessarily equal to the average rate vector $\bar{\lambda}$. The cost of this linear trajectory of the input and channel state processes is $J_t(f, g) = [L_{(f)}(\lambda) + L_{(g)}(\gamma)]t$. (In other words, the "probability" of $(\mathbf{f}^{(n)}, \mathbf{g}^{(n)})$ being close to (f, g) in the interval $[0, t]$ is roughly $\exp[-n(L_{(f)}(\lambda) + L_{(g)}(\gamma))t]$.) Suppose now that vectors γ and λ satisfy the conditions specified above, with the corresponding vector μ. Then, a scheduling rule can be chosen (at least in principle) such that the (scaled) service process trajectory is approximately linear: $\hat{f}^{(n)}(s) \approx \hat{f}(s) \equiv \mu s,\ 0 \le s \le t$. Then, if $q^{(n)}(0) = 0$, the queue length trajectory in $[0, t]$ is approximately linearly increasing as well, and moreover,

$$a_i q_i^{(n)}(s) \approx a_i q_i(s) = \ell s \text{ for each } i.$$

This means that for all flows, $a_i q_i^{(n)}(s)$ is approximately equal to their maximum $q_*^{(n)}(s)$ at any time s, and $q_*^{(n)}(s)$ reaches level ℓt at time t. Thus, the constructed linear *simple trajectory* (f,g,q), which is determined by the vectors λ, γ and μ, has the "unit cost of raising $q_*(s)$" equal to $[L_{(g)}(\gamma)+L_{(f)}(\lambda)]/\ell$. Therefore, the value J_* defined above in this section is the minimum possible unit cost of raising $q_*(s)$ along a simple trajectory.

The key property of the above construction of a simple trajectory (f,g,q) is as follows. Given vectors λ and γ, the corresponding vector of service rates μ is optimal in the sense that all $a_i q_i(s)$, and then $q_*(s)$, simultaneously reach level ℓt at time t. Using the condition that μ is a maximal element of $V_\gamma(N')$, it is easy to see that if (f,g) is the trajectory of input and channel state processes "offered" to the system, then *under any scheduling rule* and *for any initial condition* $q(0)$, at least one of the $a_i q_i(t)$, and then $q_*(t)$, is ℓt or greater. Thus, for any scheduling rule, J_* serves as an upper bound of the minimum possible cost of raising (scaled) maximal weighted queue length $q_*(\cdot)$ to level 1.

Our simple trajectory construction is in a sense analogous to, and serves the same purpose as, those in [12, 13]. It is however necessarily more involved, because in our case the rate region is more general convex, while in [12, 13] the outer boundary of the rate region is a hyperplane (which implies simple "work conservation" properties).

6.2 Proof of Theorem 3.2(i)

The proof formalizes the argument presented above in this section, using the construction of a simple trajectory and Mogulskii theorem (Proposition 5.1). This formalization is completely analogous to the proof of Theorem 3.2(ii) in [13] (or proof of Theorem 6.8(ii) in [12]). We omit details.

7 Refined Mogulskii Theorem (upper bound)

From this point on in the paper, we make assumption (3.2), with some fixed $C>0$.

From the standard Cramer theorem for scalar random variables (cf. Theorem 2.2.3 in [6]), we have the following bound, recorded here for future reference: for any interval $[\xi_1,\xi_2]$, where $0\le\xi_1<\xi_2\le C$, and any fixed $\delta>0$, there exists a sufficiently large $\tau>0$ such that, uniformly on non-negative $0\le t_1<t_2$ satisfying $t_2-t_1\ge\tau$:

$$\log P\{\frac{1}{t_2-t_1}[F_i(t_2)-F_i(t_1)]\in[\xi_1,\xi_2]\}\le-\left[\min_{[\xi_1,\xi_2]}L_i(\xi)-\delta\right](t_2-t_1). \tag{7.1}$$

If $B\subset I\!R_+^M$ is compact, then according to Sanov theorem (cf. Theorem 2.1.10 in [6]), we can record the following property analogous to (7.1): for any fixed $\delta>0$, there exists a sufficiently large $\tau > 0$ such that, uniformly on non-negative integers $0 \le t_1 < t_2$ satisfying $t_2 - t_1 \ge \tau$, we have

$$\log P\{\frac{1}{t_2 - t_1}[G(t_2) - G(t_1)] \in B\} \le - \left[\min_{\gamma \in B} L_{(g)}(\gamma) - \delta\right](t_2 - t_1). \tag{7.2}$$

Suppose we have an integer function $u(n) \uparrow \infty$ as $n \to \infty$, which is sublinear in n, i.e., $u(n)/n \downarrow 0$. (An example of such a function is $u(n) = \lceil n^\alpha \rceil$, with $0 < \alpha < 1$.) For any (non-decreasing) RCLL vector-function $h \in \mathcal{D}^{N+M}$, and each n, we denote by $U^n h$ the continuous piece-wise linear function obtained from h as follows: we divide the time interval $[0, \infty)$ into subintervals of equal length $u(n)/n$, that is $[0, u(n)/n]$, $[u(n)/n, 2u(n)/n], \ldots$, and linearize h within each subinterval.

Theorem 7.1 *Assume (3.2). Consider the sequence of scaled processes* $(\mathbf{f}^{(n)},\ \mathbf{g}^{(n)})$, $n = 1, 2, \ldots$, *as defined in Section 4. Let* $t > 0$ *be fixed. Suppose, for each* n *there is a fixed measurable* $B^{(n)} \subseteq \zeta_0^t \mathcal{D}^{N+M}$, *which is a subset of the set of feasible realizations of* $(\mathbf{f}^{(n)},\ \mathbf{g}^{(n)})$ *in* $[0, t]$. *(Here* feasible *simply means that each jump size of each component is such that it can occur with a positive probability.) Then, for any fixed function* $u(n)$ *as defined above, we have the following asymptotic upper bound:*

$$\limsup_{n \to \infty} \frac{1}{n} \log P(\zeta_0^t(\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B^{(n)}) \le - \liminf_{n \to \infty}\ \inf\{J_{\theta^{(n)}(t)}(U^n h) \mid h \in B^{(n)}\}\ , \tag{7.3}$$

where $\theta^{(n)}(t)$ *is the largest multiple of* $u(n)/n$ *not greater than* t, *i.e.,*

$$\theta^{(n)}(t) \doteq \frac{u(n)}{n} \lfloor \frac{t}{u(n)/n} \rfloor. \tag{7.4}$$

Proof. To avoid clogging notation, assume that $\theta^{(n)}(t) = t$ for each n, i.e., the time interval $[0, t]$ is divided into the integer number $tn/u(n)$ of $u(n)/n$-long subintervals. The proof is a fairly straightforward combinatorial estimate.

Let us fix a small $\delta > 0$. We choose a large integer $K > 0$ and divide interval $[0, C)$ into K subintervals, each $\epsilon = C/K$-long, namely $[(k-1)\epsilon, k\epsilon)$ with $k = 1, \ldots, K$. (Constant C is the upper bound in (3.2).) The k-th interval defined above we will call k-th "bin". We will choose K to be large enough so that

$$\max\{|L_i(y_1) - L_i(y_2)| : L_i(y_1), L_i(y_2) < \infty,\ y_1, y_2 \in [(k-1)\epsilon, k\epsilon)\} < \delta/(4N)$$

uniformly for all bins k. (Recall that each L_i is continuous in the compact domain where it is finite.) We will choose $\tau > 0$ such that (7.1) holds, with δ replaced by $\delta/(4N)$, for all i and for the intervals $[\xi_1, \xi_2]$ being closures of the bins.

Let us divide the simplex of all vectors representing probability distributions γ on the set of channel states M into K non-intersecting subsets ("bins"), such that the oscillation (difference between maximum and minimum) of $L_{(g)}(\gamma)$ within the closure of each bin is at most $\delta/4$. (The latter can always be achieved by making K larger, if necessary.) We also will increase τ, if necessary, to make sure that (7.2) holds for all such bins, with δ replaced by $\delta/4$.

Let $\hat{J}$ denote the lim inf in the RHS of (7.3). From now on in this proof we will only consider sufficiently large n such that $\inf\{J_t(U^n h) \mid h \in B^{(n)}\} > \hat{J} - \delta$, and $u(n) > \tau$, where τ is chosen above.

Consider a fixed n, a vector-function $h = (f_i,\ i = 1, \dots, N; g) \in B^{(n)}$, and its piece-wise linearization $U^n h = (U^n f_i,\ i = 1, \dots, N; U^n g)$. Recall that each component of $U^n h$ has a constant non-negative derivative in each of the $tn/u(n)$-long time-subintervals of $[0, t]$. Thus, the vector-function h can belong to one of the finite number $[K^{(N+1)}]^{tn/u(n)}$ of "aggregate bins," according to which bins the (constant) slopes of the components $U^n f_i$ and $U^n g$ belong to, in each of the time-subintervals.

Now, consider any fixed aggregate bin, let us call it B_{ab}, containing at least one function belonging to $B^{(n)}$, and let us pick a fixed representative element $h \in B^{(n)}$. (Recall that $J_t(U^n h) > \hat{J} - \delta$.) Consider one of the time-subintervals $[(j-1)u(n)/n, ju(n)/n]$ for some integer $1 \le j \le tn/u(n)$. By definition, the event $\{\zeta_0^t(\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B_{ab}\}$ implies, in particular, that

$$\frac{1}{u(n)/n}[\mathbf{f}_i^{(n)}(ju(n)/n) - \mathbf{f}_i^{(n)}((j-1)u(n)/n)] \in [\xi_1, \xi_2],$$

where $[\xi_1, \xi_2]$ is the bin uniquely determined by the aggregate bin B_{ab}. For the component f_i of the picked element h denote

$$y \equiv \frac{1}{u(n)/n}[f_i(ju(n)/n) - f_i((j-1)u(n)/n)],$$

and recall that $L_i(y) - \min_{[\xi_1,\xi_2]} L_i(\xi) \le \delta/(4N)$. Then from (7.1) (with δ replaced by $\delta/(4N)$) we have

$$\log P\{\frac{1}{u(n)/n}[\mathbf{f}_i^{(n)}(ju(n)/n) - \mathbf{f}_i^{(n)}((j-1)u(n)/n)] \in [\xi_1, \xi_2]\} \le$$
$$-[\min_{[\xi_1,\xi_2]} L_i(\xi) - \delta/(4N)]u(n) \le -[L_i(y) - \delta/(2N)]u(n) \le$$
$$-n \int_{(j-1)u(n)/n}^{ju(n)/n} \left[L_i(\frac{d}{ds}[U^n f_i(s)]) - \delta/(2N) \right] ds.$$

Analogous estimates hold for all subintervals $[(j-1)u(n)/n, ju(n)/n]$, all $\mathbf{f}_i^{(n)}$, and also (using (7.2)) for the process $\mathbf{g}^{(n)}$. Since all $\mathbf{f}_i^{(n)}$ and $\mathbf{g}^{(n)}$ have independent increments, we can combine the above estimates to obtain the following upper bound:

$$\log P\{\zeta_0^t(\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B_{ab}\} \le -[J_t(U^n h)n - \delta t n] \le$$
$$\le -\hat{J}n + \delta n + \delta t n \ .$$

The total number of aggregate bins is $\exp\{\frac{(N+1)(\log K)t}{u(n)} n\}$ with $[(N+1)(\log K)t]/u(n) \to 0$ as $n \to \infty$. This means that

$$P\{\zeta_0^t(\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B^{(n)}\} \le e^{\delta_n n} P\{\zeta_0^t(\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B_{ab}\}$$

with $\delta_n \to 0$. Therefore,

$$\limsup_{n\to\infty} \frac{1}{n} \log P(\zeta_0^t(\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B^{(n)}) \leq -\ \hat{J} + \delta(1+t)\ .$$

Since δ can be chosen arbitrarily small, the proof is complete. ∎

8 Large deviations upper bound via Refined Mogulskii Theorem. Proof of Theorem 3.2(ii)

8.1 Generalized fluid sample paths

From now on we specify the function $u(n)$, defined in Section 7, to be

$$u(n) = \lceil n^\alpha \rceil,\ n = 1, 2, \ldots,$$

for some fixed $\alpha \in (0, \eta)$.

Definition 8.1 *Suppose an increasing subsequence $\mathcal{N}$ of the sequence of positive integers is fixed, and the following conditions (i) and (ii) hold.*

(i) For each $n \in \mathcal{N}$, there is a fixed (feasible) realization $(q^{(n)}, q_*^{(n)}, f^{(n)}, \hat{f}^{(n)}, g^{(n)}, \hat{g}^{(n)})$ of the process $(\mathbf{q}^{(n)}, \mathbf{q}_*{}^{(n)}, \mathbf{f}^{(n)}, \hat{\mathbf{f}}^{(n)}, \mathbf{g}^{(n)}, \hat{\mathbf{g}}^{(n)})$.

(ii) As $n \to \infty$, we have the u.o.c. convergence

$$(q^{(n)}, q_*^{(n)}, f^{(n)}, \hat{f}^{(n)}, g^{(n)}, \hat{g}^{(n)}) \to (q, q_*, f, \hat{f}, g, \hat{g}) \tag{8.1}$$

for some set of Lipschitz continuous functions $(q, q_*, f, \hat{f}, g, \hat{g})$, and the u.o.c. convergence

$$\bar{J}^{(n)} = (\bar{J}_t^{(n)},\ t \geq 0) \doteq (J_{\theta^{(n)}(t)}[U^n(f^{(n)}, g^{(n)})],\ t \geq 0) \to \bar{J} = (\bar{J}_t,\ t \geq 0) \tag{8.2}$$

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

Then, the entire construction

$$\psi = [\mathcal{N};\quad (q^{(n)}, q_*^{(n)}, f^{(n)}, \hat{f}^{(n)}, g^{(n)}, \hat{g}^{(n)}),\ \bar{J}^{(n)},\ n \in \mathcal{N};\quad (q, q_*, f, \hat{f}, g, \hat{g});\quad \bar{J}]$$

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

Remark. A set of functions $(q, q_*, f, \hat{f}, g, \hat{g})$, defined as a limit of a sequence of fluid scaled trajectories of a process, is sometimes called a *fluid sample path* (FSP), cf. [13]. Therefore, the term "generalized" in the above definition of a GFSP refers to the fact that

GFSP contains not only the "fluid limit" of a (pre-limit) sequence, but the sequence itself. Moreover, the pre-limit sequence is required to satisfy condition (8.2).

Given that the (unscaled) functions $F_i^{(n)}$, $\hat{F}_i^{(n)}$, $G_m^{(n)}$, $\hat{G}_{mi}^{(n)}$ obviously have uniformly bounded increments within one time slot, the GFSP components $f, \hat{f}, g, \hat{g}$ are non-decreasing Lipschitz continuous (and then absolutely continuous) functions in $[0, \infty)$, with $f(0) = 0, \hat{f}(0) = 0, g(0) = 0, \hat{g}(0) = 0$; therefore, the components q and q_* are Lipschitz as well. Moreover, the functions $\bar{J}$ are uniformly Lipschitz across all GFSP, because each $\bar{J}^{(n)}$ is Lipschitz with the derivative upper bounded by the maximum of all possible *finite* values of $L_{(f)}(y) + L_{(g)}(\gamma)$.

Further, for any $0 \le t_1 < t_2 < \infty$,

$$\bar{J}_{t_2} - \bar{J}_{t_1} \ge J_{t_2}(f,g) - J_{t_1}(f,g). \tag{8.3}$$

Indeed, if f and g are linear in $[t_1, t_2]$, then (8.3) holds, because, for each n, by Jensen inequality,

$$\bar{J}_{t_2}^{(n)} - \bar{J}_{t_1}^{(n)} \ge (\theta_2 - \theta_1) L_{(f)} \left(\frac{f^{(n)}(\theta_2) - f^{(n)}(\theta_1)}{\theta_2 - \theta_1} \right) + (\theta_2 - \theta_1) L_{(g)} \left(\frac{g^{(n)}(\theta_2) - g^{(n)}(\theta_1)}{\theta_2 - \theta_1} \right), \tag{8.4}$$

where we denoted $\theta_2 = \theta^{(n)}(t_2)$ and $\theta_1 = \theta^{(n)}(t_1) + 1/n$; it remains to take $n \to \infty$ limit. This means that (8.3) holds for any piece-wise linear f and g (with finite number of pieces) in $[t_1, t_2]$. It remains to observe that, for actual functions f and g in $[t_1, t_2]$, we can construct a sequence of their piece-wise linearizations with the derivatives converging to the corresponding derivatives of f and g almost everywhere in $[t_1, t_2]$.

We will also need the following simple facts (in Lemmas 8.2 and 8.3).

Lemma 8.2 *Suppose there exists a sequence* $\{(f^{(n)}, g^{(n)}),\ n \in \mathcal{N}\}$ *of feasible realizations of the (scaled) input and channel state processes, such that*

$$(f^{(n)}, g^{(n)}) \to (f, g) \quad u.o.c. \tag{8.5}$$

Then there exists a GFSP, having this (f,g) *as its* f*- and* g*-components and such that its refined cost function* $\bar{J} = (\bar{J}_t,\ t \ge 0)$ *is equal to* $(J_t(f,g),\ t \ge 0)$. *(Note that the sequence* $\{(f^{(n)}, g^{(n)})\}$ *in (8.5) does* not *have to be a part of such GFSP.)*

Proof. First of all, (f,g) must be Lipschitz. (Because the corresponding unscaled realizations $F^{(n)}$ and $G^{(n)}$ have uniformly bounded increments within each time slot.) Let us fix an integer $T > 0$ and consider the function f_{lin}, which is the piece-wise linearized version of f in interval the $[0, T]$ over $1/T$-long subintervals (T^2 of them in total). For $t \ge T$, $f_{lin}(t) = f(t)$. Given (8.5), we can pick fixed n_1 such that $u(n_1)/n_1 \le 1/T^2$ and n_1 is sufficiently large so that we can construct a feasible realization $f_{lin}^{(n_1)}$ such that $\|f_{lin}^{(n_1)} - f_{lin}\|_T \le \bar{C}/n_1$, where the constant $\bar{C} > 0$ does not depend on n_1. (For example, let $0 \le C_i^{min} \le C_i^{max} < \infty$ denote the minimum and maximum possible values of $A_i(1)$. We can make each component $f_{i,lin}^{(n_1)}(\cdot)$

to "track" $f_{i,lin}(\cdot)$ recursively in time, by giving it the increment C_i^{max}/n_1 at time t (being multiple of $1/n_1$) if $f_{i,lin}^{(n_1)}(t-) < f_{i,lin}(t-)$, and increment it by C_i^{min} otherwise.) We repeat the construction with T replaced by kT, $k = 2, 3, \ldots$, to obtain sequence of corresponding piece-wise linearizations f_{lin} (depending on k) and corresponding realizations $f_{lin}^{(n_k)}$ with $n_k \to \infty$. We have $f_{lin}^{(n_k)} \to f$ u.o.c. and $[u(n_k)/n_k]/[1/(kT)] \to 0$ as $n_k \to \infty$. (This implies that "most" of the $u(n_k)/n_k$-long subintervals of $[0, kT]$ lie entirely within one of the $1/(kT)$-long subintervals.) "In parallel," we can also construct analogous sequence $g_{lin}^{(n_k)}$ approximating g. It remains to pick, if necessary, a subsequence of n_k to obtain a sequence of process realizations defining a GFSP. Obviously, this GFSP has (f, g) as its components. It is also easy to verify that the refined cost of this GFSP is indeed $(J_t(f, g),\ t \ge 0)$. Indeed, we can show that the difference between $\bar{J}_t^{(n_k)}$, defined for $(f_{lin}^{(n_k)}, g_{lin}^{(n_k)})$ as in (8.2), and the corresponding $J_t(f_{lin}, g_{lin})$ (recall that (f_{lin}, g_{lin}) also depends on k) vanishes u.o.c.; in turn, $J_t(f_{lin}, g_{lin})$ converges u.o.c. to $J_t(f, g)$, because $(d/dt)J_t(f_{lin}, g_{lin}) \to (d/dt)J_t(f, g)$ almost everywhere. ∎

Lemma 8.3 *Suppose, a sequence of GFSP* ${}^k\psi$, $k = 1, 2, \ldots$, *is such that the values of* $\|{}^kq(0)\|$ *are uniformly bounded. Then, there exists a GFSP* ψ *such that, along some subsequence of* k,

$$[({}^kq,\ {}^kq_*,\ {}^kf,\ {}^k\hat{f},\ {}^kg,\ {}^k\hat{g});\ \ {}^k\bar{J}] \to [(q, q_*, f, \hat{f}, g, \hat{g});\ \ \bar{J}]\ \ u.o.c., \tag{8.6}$$

for the corresponding components of ${}^k\psi$ *and* ψ.

Proof. We can find a subsequence along which uniform convergence (8.6) holds, because $\|{}^kq(0)\|$ are uniformly bounded and all component functions in the LHS are uniformly Lipschitz. Restricting ourselves from now on to this subsequence, for any $T > 0$ we can find a sufficiently large k so that $[({}^kq,\ {}^kq_*,\ {}^kf,\ {}^k\hat{f},\ {}^kg,\ {}^k\hat{g});\ \ {}^k\bar{J}]$ is uniformly close (say, within ϵ_k) to $[(q, q_*, f, \hat{f}, g, \hat{g});\ \ \bar{J}]$ in $[0, T]$. Then, we can find a sufficiently large n_k so that $[({}^kq^{(n_k)},\ {}^kq_*^{(n_k)},\ {}^kf^{(n_k)},\ {}^k\hat{f}^{(n_k)},\ {}^kg^{(n_k)},\ {}^k\hat{g}^{(n_k)});\ \ {}^k\bar{J}^{(n_k)}]$ is uniformly close (within ϵ_k) to $[({}^kq,\ {}^kq_*,\ {}^kf,\ {}^k\hat{f},\ {}^kg,\ {}^k\hat{g});\ \ {}^k\bar{J}]$ in $[0, T]$. Choosing a sequence of $[({}^kq^{(n_k)},\ {}^kq_*^{(n_k)},\ {}^kf^{(n_k)},\ {}^k\hat{f}^{(n_k)},\ {}^kg^{(n_k)},\ {}^k\hat{g}^{(n_k)});\ \ {}^k\bar{J}^{(n_k)}]$ with $T \to \infty$ and $\epsilon_k \to 0$, we see that this sequence has $[(q, q_*, f, \hat{f}, g, \hat{g});\ \ \bar{J}]$ as its u.o.c. limit, and thus we constructed a GFSP we seek. ∎

8.2 Definition of cost J_{**}. Equivalent form of Theorem 3.2(ii)

Let J_{**} denote the lowest refined cost of a GFSP which "brings" $q_*(t)$ to level 1 from the zero initial state $q(0) = 0$. Namely,

$$J_{**} \doteq \inf_{t \ge 0} J_{**,t}, \tag{8.7}$$

where

$$J_{**,t} \doteq \inf\{\bar{J}_t \mid \psi :\ q(0) = 0,\ q_*(t) \ge 1\}. \tag{8.8}$$

From the definition of J_* in Section 6 (via the construction of simple trajectories) and Lemma 8.2, we conclude that

$$J_{**} \leq J_*,$$

because we can always construct a GFSP for which (f, g) is a simple trajectory with $q_*(t) \geq 1$ for some $t > 0$, and with $J_t(f, g)$ being arbitrarily close to J_*.

The goal of this section is to establish the following fact, which is Theorem 3.2(ii) rephrased in terms of the sequence of fluid-scaled processes.

Theorem 8.4 *For each parameter $n = 1, 2, \ldots$, consider a version of the system under the EXP rule in a stationary regime. Then, the corresponding sequence of fluid-scaled processes is such that*

$$\limsup_{n\to\infty} \frac{1}{n} \log P\left(\mathbf{q}_*{}^{(n)}(0) > 1\right) \leq -J_{**} \ . \tag{8.9}$$

The proof of Theorem 8.4 will be carried out in Section 8.4, analogously to the proof of a similar result in Section 10 of [13], namely, by using classical Wentzell-Freidlin constructions [8] to "translate" large deviations results on a finite time interval into results in a stationary regime. But first, as in [13], we need to establish those finite time interval properties.

8.3 Large deviations properties on a finite time interval

Theorem 8.5 *For any fixed $T \geq 0$ and $0 \leq c < 1$, let us denote*

$$J_{**,\leq T,c} \doteq \inf\{\bar{J}_t \mid \psi : \ q_*(0) \leq c, \ q_*(t) \geq 1 \textit{ for some } t \leq T\}.$$

Then, we have:

$$\limsup_{n\to\infty} \frac{1}{n} \log \sup_{q_*^{(n)}(0)\leq c} P\left(\sup_{t\in[0,T]} \mathbf{q}_*{}^{(n)}(t) > 1\right) \leq -J_{**,\leq T,c} \ , \tag{8.10}$$

where the sup *over $q_*^{(n)}(0) \leq c$ is a supremum over all processes with non-random initial state satisfying this condition, and*

$$J_{**,\leq T,c} \downarrow J_{**,\leq T,0} \equiv \inf_{t\leq T} J_{**,t}, \quad \textit{as } c \downarrow 0. \tag{8.11}$$

Proof. We have

$$\sup_{q_*^{(n)}(0)\leq c} P\left(\sup_{t\in[0,T]} \mathbf{q}_*{}^{(n)}(t) > 1\right) \leq P\left((\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B^{(n,c)}\right),$$

where

$$B^{(n,c)} \doteq \{(f^{(n)}, g^{(n)}) \mid \exists \ q^{(n)}(0) \text{ and } t \leq T, \text{ such that } q_*^{(n)}(0) \leq c \text{ and } q_*^{(n)}(t) > 1\}.$$

By Theorem 7.1,

$$\limsup_{n\to\infty} \frac{1}{n} \log P\left((\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B^{(n,c)}\right) \le - \liminf_{n\to\infty} \inf\{J_{\theta^{(n)}(T)}(U^n h) \mid h \in B^{(n,c)}\} . \quad (8.12)$$

If we denote by Y the lim inf in the RHS of (8.12), we have $Y \ge J_{**,\le T,c}$. Indeed, we can choose a subsequence, on n, of trajectories $h^{(n)} \in B^{(n,c)}$, with $J_{\theta^{(n)}(T)}(U^n h^{(n)}) \to Y$, which converges and defines a GFSP with refined cost Y, and satisfies the conditions defining $J_{**,\le T,c}$. This proves (8.10). The cost $J_{**,\le T,c}$ is non-increasing and continuous in c. (The continuity easily follows from Lemma 8.3.) This proves (8.11). ∎

Theorem 8.6 *For any fixed $C^* > \delta > 0$, and $T > 0$, let us denote*

$$K(C^*, \delta, T) \doteq \inf\{\bar{J}_T \mid \psi: \ q_*(0) \le C^*, \ q_*(t) \ge \delta \text{ for all } t \in [0,T]\}.$$

Then, we have:

$$\limsup_{n\to\infty} \frac{1}{n} \log \sup_{q_*^{(n)}(0) \le C^*} P\left(\sup_{t\in[0,T]} \mathbf{q}_*^{(n)}(t) \ge \delta\right) \le -K(C^*, \delta, T) . \quad (8.13)$$

(The sup *over $q_*^{(n)}(0) \le C^*$ is supremum over all processes with non-random initial state satisfying this condition.)*

Proof of this fact is completely analogous to that of Theorem 8.5. ∎

Theorem 8.7 *Given fixed $C^* > 0$, the value of the cost $K(C^*, \delta, T)$, defined in Theorem 8.6, grows linearly with T, uniformly on $\delta > 0$. More precisely, for any $C^* > 0$ there exists $\Delta_1 > 0$ such that, for all sufficiently large T and all $\delta \in (0, C^*)$, $K(C^*, \delta, T) \ge \Delta_1 T$.*

Proof of this theorem is postponed until Section 10, because it requires the notion of a local fluid sample path, which we define and study in Section 9.

The proof of the process Q stability, carried out in [10], includes (as the key part) the following property of the scaled processes: for some $T_2 > 0$ and $\delta_2 > 0$, and all sufficiently large n, uniformly on the initial states with $q_*^{(n)}(0) \le 1$, we have

$$E\mathbf{q}_*^{(n)}(T_2) \le 1 - \delta_2.$$

In turn, this property along with Dynkin inequality implies (see Lemma 9.2 in [13]) the following

Lemma 8.8 *Let constants $C^* > \delta > 0$ be fixed. Consider the stopping time*

$$\beta_1^{(n)} = \inf\{t \ge 0 \mid \mathbf{q}_*^{(n)}(t) \le \delta\} .$$

Then, for all sufficiently large n, uniformly on the initial states with $q_*^{(n)}(0) \le C^*$, we have

$$E\beta_1^{(n)} \le \Delta C^* < \infty ,$$

for some finite $\Delta > 0$.

8.4 Proof of Theorem 8.4

Let us choose arbitrary $C^* > 1$ and choose $T > 0$ large enough so that $K(C^*, \delta, T) > J_{**}$ for any $\delta > 0$. (We can choose such T by Theorem 8.7.) Let us choose arbitrary small $\epsilon_2 > 0$ and then $\epsilon^* > 0$ small enough so that $J_{**,\le T,\epsilon^*} > J_{**} - \epsilon_2$. (See Theorem 8.5.) Let us choose arbitrary ϵ and δ such that $0 < \delta < \epsilon < \epsilon^*$.

Let us denote by $p^{(n)}$ the stationary distribution of the rescaled process $\mathbf{q}^{(n)}$. To prove Theorem 8.4 we need to show that

$$\limsup_{n\to\infty} \frac{1}{n} \log p^{(n)}(\mathbf{q}_*^{(n)}(t) > 1) \le -J_{**}. \tag{8.14}$$

Consider the sequence of scaled processes $\{\mathbf{q}^{(n)},\ n = 1, 2, \ldots\}$, and for each of them define the sequence of stopping times

$$0 = \eta_0^{(n)} \le \beta_1^{(n)} \le \eta_1^{(n)} \le \beta_2^{(n)} \le \ldots$$

as follows:

$$\beta_j^{(n)} = \inf\{t \ge \eta_{j-1}^{(n)} \mid \mathbf{q}_*^{(n)}(t) \le \delta\},\ j \ge 1,$$
$$\eta_j^{(n)} = \inf\{t \ge \beta_j^{(n)} \mid \mathbf{q}_*^{(n)}(t) \ge \epsilon\},\ j \ge 1\ .$$

Given that we consider an ergodic version of each process $\mathbf{q}^{(n)}$, it is easy to verify that this process, sampled at the stopping times $\eta_j^{(n)}$, $j = 1, 2, \ldots$, is also ergodic. Indeed, this is the Markov chain

$$\hat{\mathbf{q}}^{(n)} = \{\hat{\mathbf{q}}_j^{(n)} \doteq \mathbf{q}^{(n)}(\eta_j^{(n)}),\ j = 1, 2, \ldots\},$$

with finite number of states (because jumps of $\mathbf{q}^{(n)}$ are uniformly bounded); all states of $\hat{\mathbf{q}}^{(n)}$ are connected because so are the states of $\mathbf{q}^{(n)}$. Let us denote by $\hat{p}^{(n)}$ the unique stationary distribution of chain $\hat{\mathbf{q}}^{(n)}$. It is also easy to check that, if we consider an ergodic version of process $\mathbf{q}^{(n)}$, then $E[\eta_j^{(n)} - \eta_{j-1}^{(n)}] < \infty$ for any n. Given this, we have the following representation of the stationary distribution of $\mathbf{q}^{(n)}$ via that of $\hat{\mathbf{q}}^{(n)}$ (cf. Lemma 10.1 in [13] and references therein):

$$p^{(n)}(B) = \frac{\int_{\Xi^{(n)}} \hat{p}^{(n)}(dx) E_x \int_0^{\eta_1^{(n)}} I\{\mathbf{q}^{(n)}(t) \in B\} dt}{\int_{\Xi^{(n)}} \hat{p}^{(n)}(dx) E_x \eta_1^{(n)}}\ , \tag{8.15}$$

where B is a subset of states of $\mathbf{q}^{(n)}$ (the entire state space is countable, so there is no measurability issues), $\Xi^{(n)}$ is the (countable) state space of $\hat{\mathbf{q}}^{(n)}$, $\eta_1^{(n)}$ is the stopping time for process $\mathbf{q}^{(n)}$ as defined above; also, here and below (with a slight abuse of notation) we write P_x for the conditional distribution of process $\mathbf{q}^{(n)}$ given initial state $\mathbf{q}^{(n)}(0) = x$, and E_x for the corresponding expectation.

We are interested (see (8.14)) in estimating $p^{(n)}(B)$ for the set $B = \{q_* > 1\}$. We will evaluate the asymptotics of both the denominator and the numerator of (8.15) as $n \to \infty$.

First, it is easy to see that

$$\liminf_{n\to\infty} \inf_{x\in\Xi^{(n)}} E_x[\eta_1^{(n)} - \beta_1^{(n)}] > 0,$$

because the jumps of $\mathbf{Q}^{(n)}$ are uniformly bounded and then the increments of $\mathbf{q}^{(n)}$ within $1/n$-long intervals are bounded by $\hat{C}/n$ for some constant $\hat{C}$. Therefore, the lim inf of the denominator of (8.15) is bounded away from 0 for all large n.

To estimate the numerator of (8.15), consider the following additional stopping time:

$$\eta^{(n),\uparrow} = \inf\{t \geq 0 \mid \mathbf{q}_*^{(n)}(t) \geq 1\}.$$

Note that, for all sufficiently large n, at the stopping time $t = \eta_j^{(n)}$ we must have $\mathbf{q}_*^{(n)}(t) < \epsilon^*$ (recall that $\epsilon < \epsilon^*$), and therefore $q_* < \epsilon^*$ for any state in $\Xi^{(n)}$; similarly, at the stopping time $t = \eta^{(n),\uparrow}$ we have $\mathbf{q}_*^{(n)}(t) < C^*$.

For any fixed x such that $q_* \leq \epsilon^*$, we can write

$$E_x \int_0^{\eta_1^{(n)}} I\{\mathbf{q}^{(n)}(t) \in B\} dt \leq P_x\{\eta^{(n),\uparrow} \leq \beta_1^{(n)}\} \sup_{y:\ q_* \leq C^*} E_y \beta_1^{(n)} \ .$$

We know from Lemma 8.8 that

$$\limsup_{n\to\infty} \sup_{y:\ q_* \leq C^*} E_y \beta_1^{(n)} \leq \Delta C^*,$$

and have the estimate

$$P_x(\eta^{(n),\uparrow} \leq \beta_1^{(n)}) \leq P_x(\beta_1^{(n)} \geq T) + P_x(\eta^{(n),\uparrow} \leq T).$$

Finally, we have (since $K(C^*, \delta, T) > J_{**}$ by our choice of T)

$$\limsup_{n\to\infty} \frac{1}{n} \log \sup_{x:\ q_* \leq \epsilon^*} P_x(\beta_1^{(n)} \geq T) < -J_{**} \ ,$$

and (due to Theorem 8.5 and since $J_{**,\leq T,\epsilon^*} > J_{**} - \epsilon_2$ by our choice of ϵ^*)

$$\limsup_{n\to\infty} \frac{1}{n} \log[\sup_{x:\ q_* \leq \epsilon^*} P_x(\eta^{(n),\uparrow} \leq T)] \leq -(J_* - \epsilon_2) \ .$$

Since ϵ_2 can be chosen arbitrarily small, we have proved (8.14), and therefore the desired bound (8.9). ∎

9 Construction of a local fluid sample path

In this section, we introduce the notion of a local fluid sample path. This notion naturally arises if we want to study properties of GFSPs on infinitesimal intervals. Roughly speaking, to estimate the derivatives of a GFSP components, sometimes one needs to consider the behavior of the process on a time scale "finer" than fluid scale. (This may not be necessary under *any* scheduling rule, but typically is unavoidable for rules that are *not* asymptotically scaling-invariant, such as EXP.)

Consequently, the idea of a local FSP construction is roughly as follows. (The actual construction, given later in this section, is somewhat more involved.) For each n, we consider the fluid scaled functions, say function $f_i^{(n)}(\cdot)$ to be specific, in the $Sn^{\eta-1}$-long intervals $[\tau, \tau + Sn^{\eta-1}]$, where τ and $S > 0$ are fixed constants. We look at the increments of this function, "magnified in space and time" by factor $n^{1-\eta}$:

$$ {}_\diamond f_i^{(n)}(s) = n^{1-\eta}[f_i^{(n)}(\tau + n^{\eta-1}s) - f_i^{(n)}(\tau)], \quad s \in [0, S], $$

and take a limit, which is a function ${}_\diamond f_i(s)$ on $[0, S]$. Thus, in terms of the original - unscaled - time, a local FSP describes the evolution of the process over Sn^η-long intervals. This illustrates the reason for choosing (for the EXP rule) the specific rescaling described above: when the queue length vector $Q(\cdot)$, and therefore $\bar{Q}(\cdot)$, is of the order of n, then n^η is the relevant space and time scale order to consider the dynamics of the queue length. (Order n^η differences between queue lengths Q_i correspond to order 1 differences in the exponents in (3.6).)

The formal construction is as follows. Consider a fixed GFSP ψ. Let us introduce "recentered" weighted queue lengths (similar to those in [10])

$$ \tilde{Q}_i^{(n)}(t) \doteq a_i Q_i^{(n)}(t) - b_i[\bar{Q}^{(n)}(t)]^\eta, $$

where a fixed vector $b = (b_1, \dots, b_N) \in I\!R^N$ is chosen such that $e^b \in I\!R^N_{++}$ is an outer normal to the rate region V at some point $v^* > \bar{\lambda}$. (Note that necessarily v^* is a maximal element of V, that is $v^* \in V_\pi^*(N)$.) The fluid-scaled version $\tilde{q}_i^{(n)} = \Gamma^n \tilde{Q}_i^{(n)}$ is

$$ \tilde{q}_i^{(n)}(t) = a_i q_i^{(n)}(t) - b_i[\bar{q}^{(n)}(t)]^\eta n^{\eta-1}, \ t \geq 0, $$

and its fluid limit is $(a_i q_i(t), \ t \geq 0)$. Therefore, we have u.o.c. convergence:

$$ \tilde{q}_i^{(n)} = (\tilde{q}_i^{(n)}(t), \ t \geq 0) \to (a_i q_i(t), \ t \geq 0), \quad \forall i, $$

$$ \tilde{q}_*^{(n)} = (\tilde{q}_*^{(n)}(t), \ t \geq 0) \to q_*, $$

where

$$ \tilde{q}_*^{(n)}(t) \doteq \max_i \tilde{q}_i^{(n)}(t). $$

Let a time point $\tau_1 > 0$ be fixed, such that $q_*(\tau_1) > 0$. Suppose we have a sequence (in n) of time intervals $[t_1^{(n)}, t_2^{(n)}]$, such that the following condition is satisfied:

$$ t_2^{(n)} - t_1^{(n)} = S\sigma_n, $$

where $S > 0$ is a fixed constant and $\sigma_n = [\bar{q}^{(n)}(t_1^{(n)})]^\eta n^{\eta-1}$.

For each n, consider the following rescaled functions. For $s \in [0,S]$, let

$$
\begin{aligned}
{}_\diamond q_i^{(n)}(s) &\doteq \frac{1}{\sigma_n}[\tilde{q}_i^{(n)}(t_1^{(n)} + \sigma_n s) - \tilde{q}_*^{(n)}(t_1^{(n)})], \ i \in N, \\
{}_\diamond q_*^{(n)}(s) &\doteq \max_i {}_\diamond q_i^{(n)}(s) \equiv \frac{1}{\sigma_n}[\tilde{q}_*^{(n)}(t_1^{(n)} + \sigma_n s) - \tilde{q}_*^{(n)}(t_1^{(n)})], \\
{}_\diamond f_i^{(n)}(s) &\doteq \frac{1}{\sigma_n}[f_i^{(n)}(t_1^{(n)} + \sigma_n s) - f_i^{(n)}(t_1^{(n)})], \ i \in N, \\
{}_\diamond \hat{f}_i^{(n)}(s) &\doteq \frac{1}{\sigma_n}[\hat{f}_i^{(n)}(t_1^{(n)} + \sigma_n s) - \hat{f}_i^{(n)}(t_1^{(n)})], \ i \in N, \\
{}_\diamond g_m^{(n)}(s) &\doteq \frac{1}{\sigma_n}[g_m^{(n)}(t_1^{(n)} + \sigma_n s) - g_m^{(n)}(t_1^{(n)})], \ m \in M, \\
{}_\diamond \hat{g}_{mi}^{(n)}(s) &\doteq \frac{1}{\sigma_n}[\hat{g}_{mi}^{(n)}(t_1^{(n)} + \sigma_n s) - \hat{g}_{mi}^{(n)}(t_1^{(n)})], \ m \in M, \ i \in N.
\end{aligned}
\tag{9.1}
$$

Now, we can find a subsequence of n such that the following u.o.c. convergence holds:

$$
({}_\diamond q_*^{(n)}, {}_\diamond f^{(n)}, {}_\diamond \hat{f}^{(n)}, {}_\diamond g^{(n)}, {}_\diamond \hat{g}^{(n)}) \to ({}_\diamond q_*, {}_\diamond f, {}_\diamond \hat{f}, {}_\diamond g, {}_\diamond \hat{g}), \tag{9.2}
$$

where all vector-functions are considered in the interval $s \in [0,S]$, and the limit functions in the RHS are Lipschitz continuous. Note that expression (9.2) is for 5-tuples: the convergence ${}_\diamond q^{(n)} \to {}_\diamond q$ cannot be included in the proper sense, because it is possible that for some i, ${}_\diamond q_i^{(n)}(0) \to -\infty$. However, the uniform convergence (9.2) implies that we can choose a further subsequence of n along which, in addition, we have the uniform in $[0,S]$ convergence

$$
{}_\diamond q^{(n)} \to {}_\diamond q, \tag{9.3}
$$

in the following sense: each component ${}_\diamond q_i$, $i \in N$, of ${}_\diamond q$ is either a finite Lipschitz continuous function or ${}_\diamond q_i(s) \equiv -\infty$. (The latter is simply the convention for the case when ${}_\diamond q_i^{(n)}(0) \to -\infty$, implying by (9.2) that ${}_\diamond q_i^{(n)}(s) \to -\infty$ uniformly on $s \in [0,S]$.) Clearly, we have: ${}_\diamond q_*(s) = \max_i {}_\diamond q_i(s)$, $s \in [0,S]$, and ${}_\diamond q_*(0) = 0$.

It is easy to see that

$$
\liminf_{n\to\infty} \sigma_n^{-1}[\bar{J}^{(n)}_{t_2^{(n)}} - \bar{J}^{(n)}_{t_1^{(n)}}] \geq J_S({}_\diamond f, {}_\diamond g). \tag{9.4}
$$

Indeed, first, we can establish (9.4) for any piece-wise linearization of $({}_\diamond f, {}_\diamond g)$, using convexity of the rate functions $L_{(f)}(\cdot)$ and $L_{(g)}(\cdot)$, along with the fact that $n^\alpha/n = o(n^\eta/n)$ as $n \to \infty$; and then observe that such piece-wise linearizations can have the cost arbitrarily close to $J_S({}_\diamond f, {}_\diamond g)$. (The argument here is similar to that in the proof of (8.3).)

The 6-tuple $({}_\diamond q, {}_\diamond q_*, {}_\diamond f, {}_\diamond \hat{f}, {}_\diamond g, {}_\diamond \hat{g})$ constructed above is what we will call a *local fluid sample path* (LFSP). In what follows, we restrict ourselves to the subset N' of those i for which ${}_\diamond q_i(0)$ is finite, and simply exclude all components $i \notin N'$ from consideration. This in particular means that we consider $V_\gamma(N')$ in place of V_γ, and exclude $i \notin N'$ from the definition of the cost $J_s({}_\diamond f, {}_\diamond g)$. (If the set $N \setminus N'$ of flows with ${}_\diamond q_i(0) = -\infty$ is non-empty, the LFSP describes the behavior of the system in a time interval where flows $i \in N \setminus N'$ cannot "compete for service" with flows $i \in N'$. This means that, in any slot, the EXP scheduling rule will pick for service a flow $i \in N'$ as long as at least one of these flows can be served at non-zero rate in this slot.) However, to avoid clogging notation, and without loss of generality, let us assume that $N' = N$. The following lemma describes the basic dynamics of an LFSP.

Lemma 9.1 *For any LFSP, for almost all $s \in [0, S]$, the following (proper) derivatives exist and are finite:*

$$\lambda(s) \doteq \frac{d}{ds}{}_\diamond f(s), \quad \gamma(s) \doteq \frac{d}{ds}{}_\diamond g(s), \tag{9.5}$$

$$\mu(s) \doteq \frac{d}{ds}{}_\diamond \widehat{f}(s), \tag{9.6}$$

$$\frac{d}{ds}{}_\diamond q(s), \quad \frac{d}{ds}{}_\diamond q_*(s), \quad \frac{d}{ds}{}_\diamond \widehat{g}(s), \tag{9.7}$$

and, moreover, the following relations hold:

$$\frac{d}{ds}{}_\diamond q(s) = a \times [\lambda(s) - \mu(s)], \tag{9.8}$$

$${}_\diamond q_*(s) = \max_i {}_\diamond q_i(s), \tag{9.9}$$

$$\frac{d}{ds}{}_\diamond q_*(s) = \frac{d}{ds}{}_\diamond q_i(s) \;\; \textit{for each } i \textit{ such that } {}_\diamond q_i(s) = {}_\diamond q_*(s), \tag{9.10}$$

$$\mu_i(s) = \sum_m \mu_i^m \frac{d}{ds}{}_\diamond \widehat{g}_{mi}(s), \quad \forall i, \tag{9.11}$$

$$\gamma_m(s) = \sum_i \frac{d}{ds}{}_\diamond \widehat{g}_{mi}(s), \quad \forall m, \tag{9.12}$$

$$\mu(s) \in \arg\max_{v \in V_{\gamma(s)}} e^{{}_\diamond q(s) + b} \cdot v. \tag{9.13}$$

Proof. Consider a pre-limit sequence of functions $({}_\diamond q^{(n)}, {}_\diamond q_*^{(n)}, {}_\diamond f^{(n)}, {}_\diamond \widehat{f}^{(n)}, {}_\diamond g^{(n)}, {}_\diamond \widehat{g}^{(n)})$, all defined in $[0, S]$, uniformly converging to $({}_\diamond q, {}_\diamond q_*, {}_\diamond f, {}_\diamond \widehat{f}, {}_\diamond g, {}_\diamond \widehat{g})$. (Such sequence exists by (9.2) and (9.3) in the LFSP definition. Recall that we only consider flows $i \in N'$ for which ${}_\diamond q_i(\cdot)$ are finite, and we assume $N' = N$ to simplify notation.) The derivatives (9.5)-(9.7) exist a.e. because all functions are Lipschitz. Relations (9.8), (9.11), (9.12) (their integral forms, to be precise) are a simple consequence of the corresponding relations (4.2), (4.4) and (4.3) for the pre-limit functions; (9.9) follows from the limit version of (9.1). Almost all s are such that all derivatives $(d/ds){}_\diamond q_i(s)$ and $(d/ds){}_\diamond q_*(s)$ exist, and for such s relation (9.10) must hold. To prove the key relation (9.13), consider the behavior of ${}_\diamond q$ in a small interval $[s, s+\delta]$. For the unscaled process, this corresponds to the interval $[nt_1^{(n)} + n\sigma_n s, nt_1^{(n)} + n\sigma_n s + n\sigma_n \delta]$. (Its length $n\sigma_n\delta$ is of the order n^η.) Setting $t = nt_1^{(n)} + \sigma_n s$ as $n \to \infty$, it is easy to see (analogously to the argument in Section 4.3, p. 198, of [10]) that the ratios of $\exp(\cdot)$ terms (for different i) in EXP rule definition (3.6) converge to the ratios of the numbers $e^{\diamond q_i(s)+b_i}$. This means that, if δ is small enough, the unscaled process in $[nt_1^{(n)}+n\sigma_n s, nt_1^{(n)}+n\sigma_n s+n\sigma_n\delta]$ is such that at any time when channel state is m, only flows $i \in \arg\max \mu_i^m e^{\diamond q_i(s)+b_i}$ can be chosen for service. This easily implies (9.13); we omit details which are almost identical to those in the proof of Lemma 5(ii) in [14]. ∎

The time points $s \in [0, S]$ for which the derivatives (9.5)-(9.7) exist and relations (9.8)-(9.13) hold, are called *regular*. According to Lemma 9.1, almost all points in $[0, S]$ are regular. In what follows, we adopt the convention that when we write any expression or condition involving any of the derivatives (9.5)-(9.7) at point s, we always mean that it holds under the additional condition that s is regular, even if we do not state this explicitly.

The simple property formulated in the next lemma is nevertheless very important for our analysis. It says that if the derivative of the cost $J_s(\diamond f, \diamond g)$ at some point s is small enough, then $\diamond q_*(s)$ must decrease at that point. Intuitively, this is because small $J_s'(\diamond f, \diamond g)$ implies that the instantaneous input rates $\diamond f_i'(s)$ are close to the average input rates $\bar{\lambda}_i$ and the "instantaneous distribution of channel states" $\diamond g'(s)$ is close to the stationary distribution π; therefore, instantaneously, we have a "non-overloaded" system, in which the queues have a "tendency" to decrease.

Lemma 9.2 *There exist fixed constants $\epsilon_1 > 0$ and $\delta_1 > 0$ such that the following holds: at (almost) any time $s \in [0, S]$,*

$$\frac{d}{ds} J_s(\diamond f, \diamond g) \le \epsilon_1 \quad \textit{implies} \quad \frac{d}{ds} \diamond q_*(s) \le -\delta_1. \tag{9.14}$$

Proof. Let ϕ^* be a stochastic matrix such that $v^* = v(\phi^*)$, where v^* is the vector (on the outer boundary of V) from the definition of b. (That is, (3.4) holds for v^* and ϕ^*.) Let us denote by N_* the subset of those i in N for which $\diamond q_i(s) = \diamond q_*(s)$. Consider a fixed state m, such that $\gamma_m(s) > 0$, and denote $\phi_{mi}(s) = \diamond \hat{g}_{mi}'(s)/\gamma_m(s)$. Then, using an argument analogous to that in [10] (at the end of Section 4.3), we can establish the following fact:

$$\sum_{i \in N_*} \phi_{mi}(s) \ge \sum_{i \in N_*} \phi_{mi}^*,$$

and therefore, since the value of $e^{b_i}\mu_i^m$ is the same for all i such that $\phi_{mi}^* > 0$ (because e^b is the outer normal to V at v^*, also see [10]),

$$\sum_{i \in N_*} e^{b_i} \phi_{mi}(s) \mu_i^m \ge \sum_{i \in N_*} e^{b_i} \phi_{mi}^* \mu_i^m. \tag{9.15}$$

Now, we can proceed with the proof of the lemma statement. Small $\frac{d}{ds} J_s(\diamond f, \diamond g)$ implies that $\lambda(s)$ is close to $\bar{\lambda}$, and $\gamma(s)$ is close to π. Therefore, uniformly on all sufficiently small values of the cost derivative, we have:

$$\lambda(s) < \bar{\lambda}^* < v^*, \quad \text{for some fixed vector } \bar{\lambda}^* \text{ close to } \bar{\lambda},$$

and (summing up (9.15), weighted by $\gamma_m(s)$, over m)

$$\sum_{i\in N_*} e^{b_i}\mu_i(s) \geq \sum_{i\in N_*} e^{b_i} v_i^* - \epsilon_3 \quad \text{for some fixed small } \epsilon_3 > 0.$$

This means that, for any fixed $\epsilon_4 > 0$, uniformly on all sufficiently small values of $\frac{d}{ds} J_s({}_\diamond f, {}_\diamond g)$,

$$\sum_{i\in N_*} e^{b_i}[\lambda_i(s) - \mu_i(s)] \leq \sum_{i\in N_*} e^{b_i}[\bar{\lambda}_i^* - v_i^*] + \epsilon_4 < 0.$$

Using the fact that at the regular point s we have ${}_\diamond q_*'(s) = {}_\diamond q_i'(s) = a_i(\lambda_i(s) - \mu_i(s))$ for all $i \in N_*$, the lemma statement follows. ∎

Lemma 9.3 *There exist fixed constants $\epsilon_2 > 0$ and $\delta_2 > 0$ such that the following holds for any LFSP:*

$$J_S({}_\diamond f, {}_\diamond g) - J_0({}_\diamond f, {}_\diamond g) \leq \epsilon_2 S \quad \textit{implies} \quad {}_\diamond q_*(S) - {}_\diamond q_*(0) \leq -\delta_2 S. \tag{9.16}$$

Proof. Let us pick positive $\epsilon_2 < \epsilon_1$. (Constants ϵ_1 and δ_1 are those from Lemma 9.2.) Denote by B_1 the subset of $s \in [0, S]$ where $J_s'({}_\diamond f, {}_\diamond g) \geq \epsilon_2$, and let $B_2 = [0, S] \setminus B_1$. These sets are Lebesgue measurable, with the Lebesgue measure $\Lambda(B_1) \leq (\epsilon_2/\epsilon_1)S$. Obviously,

$$\int_{B_1} {}_\diamond q_*'(s) ds \leq \ell \frac{\epsilon_2}{\epsilon_1} S, \tag{9.17}$$

where $\ell > 0$ is the Lipschitz constant for ${}_\diamond q_*(\cdot)$. From Lemma 9.2 and bound $\Lambda(B_2) \geq [1 - (\epsilon_2/\epsilon_1)]S$, we have

$$\int_{B_2} {}_\diamond q_*'(t) \leq -\delta_1 S[1 - (\epsilon_2/\epsilon_1)]. \tag{9.18}$$

Summing up (9.17) and (9.18), we obtain

$${}_\diamond q_*(S) - {}_\diamond q_*(0) \leq -\delta_1 S[1 - (\epsilon_2/\epsilon_1)] + \ell \frac{\epsilon_2}{\epsilon_1} S. \tag{9.19}$$

Choosing any $\delta_2 < \delta_1$ and sufficiently small ϵ_2, we complete the proof. ∎

10 Proof of Theorem 8.7

Let $T > 0$ be fixed. Consider a GSFP ψ such that $\min_{[0,T]} q_*(t) > 0$. Let us fix constant positive $\epsilon_3 < \epsilon_2$ (this ϵ_2 and δ_2 are those from Lemma 9.3), and assume that

$$\bar{J}_T - \bar{J}_0 \leq \epsilon_3 T.$$

Let us fix arbitrary $S > 0$, and for each n consider the set of non-overlapping subintervals of $(t, t+T]$, having the form $(h, h+d]$ and constructed as follows. The left boundary h of the first interval is $h = 0$ and the length is $d = S[\bar{q}^{(n)}(h)]^\eta n^{\eta - 1}$. The second interval is obtained by "resetting" h to $h+d$, and so on. As soon as we obtain an interval with the right boundary being greater than T, we do not consider this last interval, and stop iterating. Choose the subset of those subintervals, for which we have

$$\bar{J}^{(n)}_{h+d} - \bar{J}^{(n)}_{h} \le \epsilon_2 d.$$

We call the chosen set of intervals A_2, and their union $B_2 \subseteq (0,T]$. (Both A_2 and B_2 depend on n.) We also denote $B_1 = (0,T] \setminus B_2$.

The following property holds: *for all sufficiently large n, for all intervals $(h, h+d] \in A_2$, we have*

$$\tilde{q}^{(n)}_*(h+d) - \tilde{q}^{(n)}_*(h) \le -(1/2)\delta_2 d. \tag{10.1}$$

Indeed, if this were not true, we could pick one violating interval $[h, h+d]$ for each n (let's call this interval $[t_1^{(n)}, t_2^{(n)}]$), and then choose a subsequence of such intervals, which leads to an LFSP, violating Lemma 9.3.

Now, we can essentially repeat the argument of the proof of Lemma 9.3, but look at increments rather than derivatives. Namely, for all large n, the Lebesgue measure of B_1 is upper bounded by $C_2[2\epsilon_3/\epsilon_2]T$, and the total increment of $\tilde{q}^{(n)}_*(\cdot)$ over the set B_1 is upper bounded by $C_3[2\epsilon_3/\epsilon_2]T$, where C_2 and C_3 are constants. From property (10.1) we obtain that the total increment of $\tilde{q}^{(n)}_*(\cdot)$ over the set B_2 is upper bounded by $-\delta_2 T\{1 - C_3[2\epsilon_3/\epsilon_2]\}$. Choosing ϵ_3 small enough and taking the limit as $n \to \infty$, we see that $q_*(T) - q_*(0) \le -\delta_3 T$ for sufficiently small fixed $\delta_3 > 0$. (Compare to (9.19).)

Thus, we have established the existence of $\epsilon_3 > 0$ and $\delta_3 > 0$ such that for any $T > 0$ and any GFSP with $q_*(t)$ not hitting 0 in $[0,T]$, condition $\bar{J}_T - \bar{J}_0 \le \epsilon_3 T$ implies $q_*(T) - q_*(0) \le -\delta_3 T$. Then the statement of Theorem 8.7 follows. ∎

11 Construction of an LFSP from a "low cost" GFSP

In this section we start working towards the proof of Theorem 3.2(iii). Namely, we will show that if we have a GFSP ψ such that $q_*(0) = 0$ (which is equivalent to $q(0) = 0$) and $q_*(T) = 1$ for some finite $T > 0$, and its refined cost $\bar{J}_T = J_{***} < \infty$, then we can construct an LFSP with the "unit cost of raising q_*" being at most $J_{***} + \epsilon$, with arbitrarily small $\epsilon > 0$.

Suppose such a GFSP and arbitrary $\epsilon > 0$ are fixed. Then there exists a time point $\tau \in (0,T)$ such that $q_*(\tau) > 0$, $q_*'(\tau) > 0$, $\bar{J}_\tau' > 0$, and

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

This means that we can find a finite interval $[t_1, t_2]$ (containing τ in its interior - but this fact will not be important for our purposes) such that $0 < t_1 < t_2 < T$, $q_*(t) > q_*(t_1) > 0$ for all $t \in (t_1, t_2]$, and finally

$$\frac{\bar{J}_{t_2} - \bar{J}_{t_1}}{q_*(t_2) - q_*(t_1)} < J_{***} + 2\epsilon. \tag{11.1}$$

Let us fix arbitrary $S > 0$. We claim that for every sufficiently large n, we can find an interval $[t_1^{(n)}, t_2^{(n)}]$ within $[t_1, t_2]$, satisfying the following conditions:

$$t_2^{(n)} - t_1^{(n)} = S[\bar{q}^{(n)}(t_1^{(n)})]^\eta n^{\eta-1} \equiv S\sigma_n,$$
$$\tilde{q}_*^{(n)}(t_2^{(n)}) - \tilde{q}_*^{(n)}(t_1^{(n)}) > 0,$$
$$\frac{\bar{J}^{(n)}_{t_2^{(n)}} - \bar{J}^{(n)}_{t_1^{(n)}}}{\tilde{q}_*^{(n)}(t_2^{(n)}) - \tilde{q}_*^{(n)}(t_1^{(n)})} < J_{***} + 3\epsilon. \tag{11.2}$$

To show this, we use essentially the same construction as in the proof of Theorem 8.7 in Section 10. Namely, we subdivide the interval $[t_1, t_2]$, for each n, as follows. The left boundary h of the first interval is $h = t_1$ and the length is $d = S[\bar{q}^{(n)}(h)]^\eta n^{\eta-1}$. The second interval is obtained by "resetting" h to $h + d$, and so on. As soon as we obtain an interval with the right boundary being greater than t_2, we do not consider this last interval, and stop iterating. It is easy to see that one of the constructed intervals must satisfy the conditions specified above, because otherwise (11.1) could not hold.

For each n let us pick the interval constructed just above, and then choose a subsequence of them such that, for some fixed $\tau_1 \in [t_1, t_2]$, we have $t_1^{(n)} \to \tau_1$ (and then $t_2^{(n)} \to \tau_1$ as well). Then, we can choose a further subsequence such that $({}_\diamond f^{(n)}, {}_\diamond \widehat{f}^{(n)}, {}_\diamond g^{(n)}, {}_\diamond \widehat{g}^{(n)}, {}_\diamond q_*^{(n)}, {}_\diamond q^{(n)})$ converges to an LFSP $({}_\diamond f, {}_\diamond \widehat{f}, {}_\diamond g, {}_\diamond \widehat{g}, {}_\diamond q_*, {}_\diamond q)$, in the sense specified in Section 9. In addition, we obviously have ${}_\diamond q_*(S) - {}_\diamond q_*(0) \geq 0$. Recall that $S > 0$ is a constant which can be chosen arbitrarily. This means that we can construct the above LFSP in an arbitrarily long time interval $[0, S]$.

From Lemma 9.3, there exists $\epsilon_3 > 0$, such that

$$J_S({}_\diamond f, {}_\diamond g) - J_0({}_\diamond f, {}_\diamond g) \geq \epsilon_3 S. \tag{11.3}$$

This and the uniform on n bound (11.2) imply that for some fixed $\epsilon_4 > 0$ we have

$${}_\diamond q_*(S) - {}_\diamond q_*(0) \geq \epsilon_4 S. \tag{11.4}$$

Then, (11.2) implies the key cost estimate:

$$\frac{J_S({}_\diamond f, {}_\diamond g) - J_0({}_\diamond f, {}_\diamond g)}{{}_\diamond q_*(S) - {}_\diamond q_*(0)} \leq J_{***} + 3\epsilon. \tag{11.5}$$

Since both ${}_\diamond q_*(s)$ and $J_s({}_\diamond f, {}_\diamond g)$ are Lipschitz, for some $\epsilon_5 > 0$, we have the upper bounds complementing (11.3) and (11.4):

$${}_\diamond q_*(S) - {}_\diamond q_*(0) \leq \epsilon_5 S, \quad J_S({}_\diamond f, {}_\diamond g) - J_0({}_\diamond f, {}_\diamond g) \leq \epsilon_5 S. \tag{11.6}$$

12 Proof of Theorem 3.2(iii)[1]

To prove (3.9) we need to show that J_{**} in (3.8) (defined via GFSPs in (8.7)), which we know is no greater than J_*, is in fact equal to J_*. Recall that J_{**} is the infimum of refined costs of those GFSPs taking $q_*(t)$ from 0 to 1.

Thus, it will suffice to show that if we have a GFSP ψ, as defined in Section 11, with refined cost $\bar{J}_T = J_{***} < \infty$, taking $q_*(t)$ from 0 to 1 in interval $[0, T]$, then $J_{***} \geq J_*$. Consider such a GFSP. From the construction and estimates of Section 11 we have the following

Assertion 1. *For an arbitrarily large $S > 0$, and arbitrarily small $\epsilon > 0$, we can construct an LFSP on the time interval $[0, S]$, such that bounds (11.4) (for some fixed $\epsilon_4 > 0$) and (11.5) hold. (The latter bound basically says that, for the LFSP, the "unit cost of raising ${}_\diamond q_*(s)$" is close to J_{***}.)*

Now, for a general LFSP, consider the following function of its ${}_\diamond q$-state, analogous to that in [15] (section 9.2):

$$\Psi({}_\diamond q) = \sum_i \frac{1}{a_i} e^{{}_\diamond q_i + b_i}. \tag{12.1}$$

We will also use its logarithm:

$$\Phi({}_\diamond q) = \log \Psi({}_\diamond q). \tag{12.2}$$

By convention, these definitions include the cases when some of the components ${}_\diamond q_i = -\infty$. In particular, if some of the queues i are *not* included in the LFSP (see LFSP definition), then those i are not included in the summation, which is equivalent to assuming ${}_\diamond q_i = -\infty$.

The function $\Phi({}_\diamond q)$ uniformly "approximates" ${}_\diamond q_*$ in the sense that, $\|\Phi({}_\diamond q) - {}_\diamond q_*\| \leq \Delta$ for some fixed $\Delta > 0$. Combining this fact with Assertion 1, in which we choose S to be sufficiently large, we obtain

Assertion 2. *For an arbitrarily small $\epsilon > 0$, there exists an LFSP on a time interval $[0, S]$, such that the following bounds hold (with some $\epsilon_4 > 0$):*

$$\Phi({}_\diamond q_*(S)) - \Phi({}_\diamond q_*(0)) \geq (\epsilon_4/2)S, \tag{12.3}$$

$$\frac{J_S({}_\diamond f, {}_\diamond g) - J_0({}_\diamond f, {}_\diamond g)}{\Phi({}_\diamond q_*(S)) - \Phi({}_\diamond q_*(0))} \leq J_{***} + 4\epsilon. \tag{12.4}$$

Since ϵ in Assertion 2 can be chosen arbitrarily small, to prove $J_{***} \geq J_*$, it suffices to show that the LHS of (12.4) is $\geq J_*$. We have

$$\frac{J_S({}_\diamond f, {}_\diamond g) - J_0({}_\diamond f, {}_\diamond g)}{\Phi({}_\diamond q_*(S)) - \Phi({}_\diamond q_*(0))} \geq \frac{\int_B [L_{(f)}(\lambda(s)) + L_{(g)}(\gamma(s))] ds}{\int_B [(d/ds)\Phi({}_\diamond q(s))] ds}, \tag{12.5}$$

[1]In the original version of the paper (see also [16]), the proof in this section was given for the special case of two flows. The current - general case - proof of Theorem 3.2(iii) was added during Dec. 2007 revision, and was motivated in part by Lemma 4 in the more recent paper [17].

where B is the subset of regular time points s within $[0, S]$, such that $(d/ds)\Phi({}_\diamond q(s)) > 0$. (Regularity of point s means that derivatives of all components of the LFSP exists as well as $(d/ds)\Phi({}_\diamond q(s)) > 0$. Almost all points s are regular w.r.t. Lebesgue measure, and the set B is Lebesgue measurable.) The RHS of (12.5) is lower bounded by

$$\inf_{s \in B} \frac{L_{(f)}(\lambda(s)) + L_{(g)}(\gamma(s))}{(d/ds)\Phi({}_\diamond q(s))}.$$

Notice that the condition $(d/ds)\Phi({}_\diamond q(s)) > 0$ (or, equivalently, $(d/ds)\Psi({}_\diamond q(s)) > 0$) at a regular point, implies that $\lambda(s)$ and $\gamma(s)$ must be such that $\lambda(s) \notin V_{\gamma(s)}$. Then, the following Lemma 12.1 completes the proof of Theorem 3.2(iii).

Lemma 12.1 *For any LFSP, at any regular point $s \in B$,*

$$\frac{L_{(f)}(\lambda(s)) + L_{(g)}(\gamma(s))}{(d/ds)\Phi({}_\diamond q(s))} \geq J_*. \tag{12.6}$$

Proof of Lemma 12.1 is given later in this section, because it in turn follows from Lemma 12.2, given next. The meaning of Lemma 12.2 is analogous to that of Lemma 4 in [17]. (Paper [17] considers FSPs under MaxWeight scheduling rule, and the corresponding potential function Ψ of the form $\sum_i q_i^{\kappa+1}$ with $\kappa > 0$.) Namely, it says the following: if $\lambda(s)$ and $\gamma(s)$ in (12.6) are fixed and we can *choose* ${}_\diamond q(s)$ resulting in the largest derivative of $(d/ds)\Phi({}_\diamond q(s))$ (and then the smallest LHS in (12.6)), among those ${}_\diamond q(s)$ with a given value of $\Phi({}_\diamond q(s))$, then an optimal ${}_\diamond q(s)$ is such that it leads to a simple (linear) trajectory. This then allows us to lower bound the LHS of (12.6) by the unit cost of a simple trajectory, which is at least J_* (by the definition of J_*).

We remark that the proof of Lemma 12.2, although related to, is different from that of Lemma 4 in [17], even besides the fact that our potential function is different. More important, our proof does *not* require the region V_γ to be necessarily polyhedral. Non-polyhedral rate regions arise, for example, in many models of scheduling in wireless systems.

Lemma 12.2 *Suppose vectors λ and γ, such that $\lambda \notin V_\gamma$, are fixed. Consider the following auxiliary optimization problem:*

$$\max_{{}_\diamond q} \min_{v \in V_\gamma} \sum_i \frac{1}{a_i} e^{{}_\diamond q_i + b_i}[a_i(\lambda_i - v_i)], \tag{12.7}$$

over vectors ${}_\diamond q$ with components $-\infty \leq {}_\diamond q_i < \infty$, subject to the constraint

$$\Psi({}_\diamond q) \leq A, \text{ where } A > 0 \text{ is a fixed constant.} \tag{12.8}$$

(The min *in (12.7) is nothing else but the time derivative $(d/ds)\Psi({}_\diamond q(s))$, given ${}_\diamond q(s) = {}_\diamond q$.) A solution ${}_\diamond q$ to (12.7)-(12.8) always exists, because the set of vectors $e^{{}_\diamond q + b}$ satisfying (12.8) is compact. Then:*

(i) There exists $\ell > 0$, such that any solution ${}_\diamond q$ to (12.7)-(12.8) has the following structure. There exists

$$\mu \in \arg\max_{v \in V_\gamma} e^{\diamond q + b} \cdot v, \tag{12.9}$$

such that

$$a_i(\lambda_i - \mu_i) = \ell, \quad \text{if} \quad e^{\diamond q_i} > 0, \tag{12.10}$$

and $a_i(\lambda_i - \mu_i) \le \ell$, otherwise.

Consequently,
(ii) the value of problem (12.7) is $A\ell$,
(iii) there exists a simple (linear) trajectory (see definition in Section 6.1) with the unit cost equal

$$\frac{\sum_{i \in N'} L_i(\lambda_i) + L_{(g)}(\gamma)}{\ell} \ge J_*, \tag{12.11}$$

where N' is the subset of i for which ${}_\diamond q_i > -\infty$, for a solution ${}_\diamond q$ to (12.7)-(12.8).

Proof of Lemma 12.2. If we change variables, $e^{\diamond q_i + b_i} = y_i$, problem (12.7)-(12.8) can be rewritten as follows

$$\max_{y \in I\!R^N} \min_{v \in V_\gamma} \sum_i y_i(\lambda_i - v_i) \tag{12.12}$$

subject to

$$\sum_i \frac{1}{a_i} y_i \le A, \tag{12.13}$$

$$y_i \ge 0, \quad \forall i. \tag{12.14}$$

Note the following properties of the function

$$H(y) = \max_{v \in V_\gamma} y \cdot v, \quad y \in I\!R^N.$$

It is convex, because it is the Legendre transform of the indicator function of convex set V_γ (see [9]). Or, it is easy to see directly that the function $H(y)$ is the maximum of linear functions $y \cdot v$, $y \in I\!R^N$, with parameters $v \in V_\gamma$. Also, a vector μ is a subgradient of H at point y if and only if $\mu \in \arg\max_{v \in V_\gamma} y \cdot v$.

The min in (12.12) is then the concave function $\hat{H}(y) \equiv y \cdot \lambda - H(y)$, and thus we have the convex problem

$$\max_{y \in I\!R^N} \hat{H}(y) \tag{12.15}$$

subject to (12.13)-(12.14). The Lagrangian for this problem is

$$\hat{H}(y) - \ell(\sum_i y_i/a_i - A) + \sum_i \beta_i y_i,$$

where $\ell \ge 0$ and $\beta_i \ge 0$ are Lagrange multipliers. For any optimal solution y^* of problem (12.15)-(12.13)-(12.14), there exist such fixed ℓ and β_i's, for which the zero vector is a supergradient of the Lagrangian at point y^*, and moreover (recall properties of $H(y)$) this supergradient must have form $\lambda - \mu + (-\ell/a_1 + \beta_1, \dots, -\ell/a_N + \beta_N) = 0$ for some

$$\mu \in \arg\max_{v \in V_\gamma} y^* \cdot v.$$

Fixing such μ, for each i we have

$$\lambda_i - \mu_i - \ell/a_i + \beta_i = 0.$$

We see that ℓ must be positive (because otherwise $\mu \leq \lambda$), and $\beta_i = 0$ if $y_i^* > 0$ by complimentary slackness. This completes the proof of statement (i), except we need to show uniqueness of ℓ across all possible optimal points y^*; the uniqueness follows from the fact that the value of problem (12.12)-(12.14) is

$$\sum_i y_i^*(\lambda_i - \mu_i) = \sum_{i:\ y_i^*>0} \frac{1}{a_i} y_i^* a_i(\lambda_i - \mu_i) = A\ell.$$

Thus, both (i) and (ii) are proved. To prove (iii), we pick any solution ${}_\diamond q$ to (12.7)-(12.8). Then, given conditions (12.9) and (12.10), the construction of a simple path in Section 6.1 applies. This in turn implies the inequality (12.11), by the J_* definition. ∎

Proof of Lemma 12.1. If we denote $\Psi({}_\diamond q(s)) = A$, then, according to Lemma 12.2(ii), $(d/ds)\Psi({}_\diamond q(s)) \leq A\ell$, and then $(d/ds)\Phi({}_\diamond q(s)) \leq \ell$. This, along with (12.11), gives (12.6). ∎

Aknowledgments. I would like to thank the referee and associate editor for very thorough reviews, which, in particular, helped in improving the exposition of the paper.

References

[1] M. Andrews, K. Kumaran, K. Ramanan, A. L. Stolyar, R. Vijayakumar, P. Whiting. Providing Quality of Service over a Shared Wireless Link. *IEEE Communications Magazine*, 2001, Vol.39, No.2, pp.150-154.

[2] M. Andrews, K. Kumaran, K. Ramanan, A. L. Stolyar, R. Vijayakumar, P. Whiting. Scheduling in a Queueing System with Asynchronously Varying Service Rates. *Probability in Engineering and Informational Sciences*, 2004, Vol. 18, pp. 191-217.

[3] P.Bender, P.Black, M.Grob, R.Padovani, N.Sindhushayana, A.Viterbi, "CDMA/HDR: A Bandwidth Efficient High Speed Wireless Data Service for Nomadic Users," *IEEE Communications Magazine*, July 2000.

[4] D. Bertsimas, I. C. Paschalidis, J. N. Tsitsiklis. Asymptotic Buffer Overflow Probabilities in Multiclass Multiplexers: An Optimal Control Approach. *IEEE Trans. Automat. Control*, 43:315-335, 1998.

[5] P. Billingsley. Convergence of Probability Measures. Wiley, 1968.

[6] A.Dembo, O.Zeitouni. Large Deviations Techniques and Applications. Springer, 1998. (2nd edition)

[7] W. Feller. *An Introduction to Probability Theory and its Applications*, Wiley, 1950.

[8] M.I.Freidlin, A.D.Wentzell. Random Perturbations of Dynamical Systems. Springer, 1998. (2nd edition)

[9] R.T.Rockafellar. *Convex Analysis.* Princeton Univ. Press, 1970.

[10] S. Shakkottai and A. L. Stolyar. Scheduling for Multiple Flows Sharing a Time-Varying Channel: The Exponential Rule. *Analytic Methods in Applied Probability. In Memory of Fridrih Karpelevich. Yu. M. Suhov, Editor.* American Mathematical Society Translations, Series 2, Volume 207, pp. 185-202. American Mathematical Society, Providence, RI, 2002.

[11] S. Shakkottai, R. Srikant, and A. L. Stolyar. Pathwise Optimality of the Exponential Scheduling Rule for Wireless Channels. *Advances in Applied Probability*, 2004, Vol. 36, No. 4, pp. 1021-1045.

[12] A. L. Stolyar and K. Ramanan. Largest Weighted Delay First Scheduling: Large Deviations and Optimality. *Annals of Applied Probability*, 2001, Vol. 11, pp. 1-48.

[13] A.L. Stolyar. Control of End-to-End Delay Tails in a Multiclass Network: LWDF Discipline Optimality. *Annals of Applied Probability*, 2003, Vol.13, No.3, pp.1151-1206.

[14] A. L. Stolyar. MaxWeight Scheduling in a Generalized Switch: State Space Collapse and Workload Minimization in Heavy Traffic. *Annals of Applied Probability*, 2004, Vol.14, No.1, pp.1-53.

[15] A. L. Stolyar. Dynamic Distributed Scheduling in Random Access Networks. *Journal of Applied Probability*, 2008, Vol.45, No.2, to appear.

[16] A. L. Stolyar. Large Deviations of Queues under QoS Scheduling Algorithms. *Proceedings of the 44th Annual Allerton Conference*, September 2006.

[17] V. J. Venkataramanan and X. Lin. Structural Properties of LDP for Queue-Length Based Wireless Scheduling Algorithms. *Proceedings of the 45th Annual Allerton Conference*, September 2007.

[18] P. Viswanath, D. Tse, and R. Laroia. Opportunistic Beamforming using Dumb Antennas. *IEEE Transactions on Information Theory*, 2002, Vol. 48(6), pp. 1277-1294.

[19] L. Ying, R. Srikant, A. Eryilmaz, and G. E. Dullerud. A Large Deviations Analysis of Scheduling in Wireless Networks. *IEEE Transactions on Information Theory*, 2006, Vol.52, No.11, pp. 5088-5098.