Pathwise Optimality of the Exponential Scheduling Rule for Wireless Channels

Sanjay Shakkottai

Wireless Networking and Communications Group

Department of Electrical and Computer Engineering

The University of Texas at Austin

shakkott@ece.utexas.edu

R. Srikant*

Coordinated Science Laboratory

and

Department of Electrical and Computer Engineering

University of Illinois at Urbana-Champaign

rsrikant@uiuc.edu

Alexander L. Stolyar

Mathematical Sciences Research Center

Bell Labs Lucent Technologies

600 Mountain Avenue

Murray Hill, NJ 07974

stolyar@research.bell-labs.com

August 13, 2004

Abstract

We consider the problem of scheduling transmissions of multiple data users (flows) sharing the same wireless channel (server). The unique feature of this problem is the fact that the capacity (service rate) of the channel varies randomly with time and asynchronously for different users.

We study a scheduling policy called Exponential scheduling rule, which was introduced in an earlier paper. Given a system with N users, and any set of positive numbers $\{a_n\}, n = 1, 2, \ldots, N$, we show that in a heavy-traffic limit, under a non-restrictive *complete resource pooling* condition, this algorithm has the property that, for each time t, it (asymptotically) minimizes $\max_n a_n \tilde{q}_n(t)$, where $\tilde{q}_n(t)$ is user n queue length in the heavy traffic regime.

Key words and phrases: Wireless networks, scheduling, quality of service, Exponential rule, heavy traffic limit, queueing networks, complete resource pooling, workload, state space collapse, optimality

Abbreviated Title: Exponential Rule Optimality

AMS Subject Classification: 60K25, 90B15, 60J70

1 Introduction

The task of supporting data flows over wireless channels is currently of great interest. This involves scheduling transmissions of multiple data flows over a shared wireless channel. The unique feature of this problem is the fact that the capacity (service rate) of the channel varies randomly with time and asynchronously for different users. The variations of the channel capacity are due to different (and random) interference levels observed by different users, and also due to *fast fading* [6], of a radio signal received by a moving user.

Such a scheduling problem arises, for example, in the 3G CDMA High Data Rate (HDR) system [7], where multiple mobile users in a cell share the same CDMA wireless channel. Packets destined for different users arrive to the cell base-station, are temporarily queued (a separate queue is maintained for each user), and need to be delivered to the mobile users. On the downlink (the link from cell base station to users), time is divided into fixed size (1.67 msec) time slots. This slot size is short enough so that each user's channel quality stays approximately constant within one time slot. In each time slot, one user is scheduled for transmission. At the beginning of each time slot, each user reports to the base station the rate at which data can be transmitted if this user is scheduled for transmission. In a HDR system (and in the generic variable channel model as well) a "good" scheduling algorithm should take advantage of channel variations by giving some form of priority to users with instantaneously better channels.

In this paper, we study the Exponential scheduling algorithm which explicitly uses information on the state of the channel and the queues. This rule was introduced in [17], where it

is shown that this rule is *throughput optimal* (TO), i.e., it renders queues stable in any system for which stability is feasible at all, with any other rule.

While this is a nice property to have, one must note that there are other throughput-optimal policies (see [1]). If a given load is stabilizable by some policy and a TO-policy is used for scheduling, then the queue lengths will not "blow up"; but this property does not indicate how large they can get. For supporting QoS requirements, one would like to have small queue-lengths (or small delays). In this paper, for a system with N users, given any set of positive numbers $\{a_n\}, n = 1, \ldots, N$, we consider the performance criterion $\max_i a_n Q_n(t)$, where $Q_n(t)$ is the queue length of the nth user at the base-station. An interesting question to ask with such a measure is the following: Is it possible to design a "statistics-oblivious" scheduling algorithm that is the "best" with respect to the above criterion, i.e., minimize $\max_n a_n Q_n(t)$, where a "statistics-oblivious" scheduling algorithm is one which *does not have any knowledge of arrival rates or channel statistics*.

In this paper, we consider this problem for a heavily-loaded system, under a (non-restrictive in applications) *complete resource pooling* (CRP) condition. Associated with CRP condition is the notion of *workload* $X(t) = \sum_n \alpha_n Q_n(t)$, where α_n's are some fixed positive constants, called *workload contributions*. (Workload contributions do *not* depend on a scheduling rule.) We show that in the heavy traffic limit, under the CRP condition and the EXP rule, the (appropriately rescaled) queue length process is such that, first, the workload is minimized and converges to a one-dimensional reflected Brownian motion, and second, a *state space collapse* (see [15, 16, 20, 3, 21]) occurs such that the queue length vector is always proportional to $[a_1^{-1}, \ldots, a_N^{-1}]$, i.e., all (rescaled) $a_n Q_n(t)$ are equal. Therefore, the EXP rule is *optimal* in the sense that the limiting process is such that *for each time t, it minimizes (rescaled)* $\max_n a_n Q_n(t)$. (As we comment in Section 7, this optimality property is *pathwise*, i.e., holds with probability 1 in an appropriately defined probability space.)

Related previous work includes that of [19] where the *MaxWeight* scheduling rule is analyzed for a *generalized switch* model (which includes our variable channel model as a special case). It is proved in [19] that, in the heavy traffic limit and under the CRP condition, the MaxWeight rule minimizes the (rescaled) system workload $X(t)$ (which converges to a one-dimensional reflected Brownian motion) and causes the state space collapse such that the queue length vector is always proportional to $[\alpha_1/\gamma_1, \ldots, \alpha_N/\gamma_N]$, where γ_n's are the algorithm parameters. Thus, if the problem is to minimize $\max_n a_n Q_n(t)$, then in order to solve it with the MaxWeight rule, the workload contributions α_n *which depend on the arrival rates and channel statistics* would need to be either "precomputed" or estimated from "real-time" queue length measurements, which may be impossible or impractical in many cases. In contrast, *the EXP rule solves this problem without any knowledge of input rates and channel statistics.* We note however that an additional (large deviations) condition on the arrival process is needed to prove our result.

In [19], the proof uses properties of the system on a fluid-time-scale as well as on a diffusion time-scale. Our proof differs from this in that we need to use *three* different time-scales, the diffusion time-scale, a "conventional" fluid time-scale (defined in Section 6), and a slower fluid time-scale (which was introduced in [17]). The properties on this slower time-scale (introduced

in Section 6.1) are key to proving our result.

A multiuser variable channel scheduling model in heavy traffic has also been considered in [5]. However, the controls proposed in [5] require that most of the service resources be "preallocated" based on the arrival rates, and only a small portion of the resources is used for dynamic control.

The outline of the rest of this paper is as follows. In Section 2, we introduce the model and basic notation. Next, in Section 3, we describe the Exponential rule. Then, in Section 4, we describe the Static Service Split (SSS) rule, a rule which "preallocates" resources based on arrival rates and channel statistics. We then state the main result of this paper in Section 5. In Section 6, we study the properties of the system on two fluid time-scales under critical loading, i.e., when the arrival rate is on the boundary of the stability region. Using these properties, in Section 7, we finally prove the main result of this paper that the EXP rule is *optimal* in the sense that, in the heavy traffic scaling limit, it minimizes $\max_i a_i Q_i(t)$ for each t.

2   The Model and Notations

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

The channel operates in discrete time. The time interval $[t, t+1)$, with $t = 0, 1, 2, \ldots$ will be called the *time slot t*. There is a finite set of channel states $M \equiv \{1, \ldots, M\}$, and the channel state is constant within each time slot. Associated with each state $m \in M$ is a fixed vector of data rates $(\mu_1^m, \ldots, \mu_N^m)$, where all μ_n^m are strictly positive integers. The meaning of μ_n^m is as follows. If in a given time slot t the channel is in state m and all service (in this time slot) is allocated to queue n, then μ_n^m type n customers are served from those already present at time t (or the entire content of queue n at t, whichever is less). Note that what we call a "channel state" here is actually a collection of channel states with respect to individual flows.

However, the *service in any time slot may be split* according to a (generally speaking random) stochastic vector $\bar{\sigma} = (\bar{\sigma}_1, \ldots, \bar{\sigma}_N)$, $\bar{\sigma}_n \geq 0, \forall n, \sum_n \bar{\sigma}_n = 1$. If in a given time slot t the channel is in state m and a "split" vector $\bar{\sigma}$ is chosen, then for each queue n, $\lfloor \bar{\sigma}_n \mu_n^m \rfloor$ type n customers are served from those already present at time t (or the entire queue n content at t, whichever is less). Here and below $\lfloor \cdot \rfloor$ denotes the integer part, and $\lceil \cdot \rceil$ - the ceiling of a number, where the ceiling is the smallest integer greater than or equal to the argument of the ceiling function.

The random channel state process is assumed to be an irreducible discrete time Markov chain with the (finite) state space M. The (unique) stationary distribution of this Markov chain we denote by $\pi \equiv (\pi^1, \ldots, \pi^M)$.

---

[1]The words flow and user are used interchangeably throughout this paper.

Denote by $A_n(t)$ the number of type n customers arrived in time slot t. We will adopt a convention that all arrivals in the time slot t actually happen at time t, but those arrivals are *not* available for service until time slot $t + 1$. We assume that the input processes $A_n(t)$ are mutually independent processes. Let us denote by λ_n, $n = 1, \ldots, N$, the mean arrival rate for flow n, i.e., the mean number of type n customers arriving in one time slot. Additional assumptions on the arrival process pertaining to the heavy traffic limit will be introduced in later sections.

The random process describing the behavior of the entire system is $S = (S(t), \ t = 0, 1, 2, \ldots)$, where

$$S(t) = \{(U_{n1}(t), \ldots, U_{nQ_n(t)}(t)), \ n = 1, \ldots, N; \ m(t) \},$$

$Q_n(t)$ is the type n queue length at time t, and $U_{nk}(t)$ is the current *delay* of the k-th type-n customer present in the system at time t. In other words, $U_{nk}(t)$ is the age (i.e., the time difference between the current time and the arrival time) of the k-th type-n customer in the queue (thus, $U_{n1}(t)$ is the current age of the type-n customer that arrived least recently). (Within each type, the customers are numbered in the order of their arrivals.) In our heavy traffic asymptotic regime, this delay is in fact (asymptotically) equivalent to end-to-end delay seen be a customer of a particular type. By convention, if $Q_n(t)$ is zero, the vector $(U_{n1}(t), \ldots, U_{nQ_n(t)}(t))$ is replaced by a special symbol (such as the null symbol). Since the possible values of $S(t)$ is countable, it follows that S is a process with a countable state-space. We will denote by $W_n(t) \equiv U_{n1}(t)$ the delay of flow n at time t (with $W_n(t) = 0$ if $Q_n(t) = 0$ by convention).

A measurable mapping H which takes a system state $S(t)$ in a time slot into a fixed probability distribution $H(S(t))$ on the set of stochastic vectors $\bar{\sigma}$, will be called a *scheduling rule*, or a *queueing discipline*. In this paper, we consider the Exponential rule described in Section 3. (We also consider a Static Service Split rule, described in Section 4, which helps to define the system stability region and serves as a tool in the proof.) We denote by $D_n(t)$ the number of type n customers served in the time slot t, then according to our conventions, for each time t,

$$Q_n(t + 1) = Q_n(t) - D_n(t) + A_n(t), \ \forall n,$$

where $D_n(t) = \min\{Q_n(t), \lfloor \bar{\sigma}_n(t)\mu_n^{m(t)} \rfloor\}$ and $\bar{\sigma}(t)$ is chosen randomly according to the distribution $H(S(t))$.

3   The Exponential Rule

Let an arbitrary set of positive constants $\gamma_1, \ldots, \gamma_N, a_1, \ldots, a_N$, and positive constants β and $\eta \in (0, 1)$ be fixed[2].

---

[2]These parameters remain fixed for the rest of this paper.

The following two related rules we will call Exponential. The Exponential (Queue length) rule (EXP-Q) chooses for service in time slot t a single queue

$$n \in n(S(t)) = \arg\max_i \gamma_i \ \mu_i(t) \ \exp\left(\frac{a_i Q_i(t)}{\beta + [\overline{Q}(t)]^\eta}\right) \ ,$$

where $\mu_n(t) \equiv \mu_n^{m(t)}$, $\overline{Q}(t) \equiv (1/N) \sum_j a_j Q_j(t)$. and the arg max function is defined by

$$\arg\max_{i \in \mathcal{A}} f_i(t) = \{j \in \mathcal{A} : f_j(t) = \max_{i \in \mathcal{A}} f_i(t)\}$$

Similarly, the Exponential (Waiting time) rule (EXP-W) chooses for service a queue

$$n \in n(S(t)) = \arg\max_i \gamma_i \ \mu_i(t) \ \exp\left(\frac{a_i W_i(t)}{\beta + [\overline{W}(t)]^\eta}\right) \ ,$$

where $\overline{W}(t) \equiv (1/N) \sum_j a_j W_j(t)$.

Remark 3.1 *1. Formally speaking, in the definition of the EXP rule, we also need to specify a "tie-breaking" convention. For example, we can assume that the queue $i = \max\{j \ : \ j \in n(S(t))\}$ is chosen.*
2. In the rest of this paper, without loss of generality we assume $\gamma_1 = 1$.
3. In this paper, for notational simplicity, we will consider only the EXP-Q rule. The results in this paper can be generalized to handle the EXP-W rule as well. Also, just to make the ensuing arguments more readable, we put $\eta = 1/2$. (The proof for any $0 < \eta < 1$ is obtained by trivial modifications.)

A useful property of this rule which has been proved in [17] is the following:

Theorem 3.1 *An EXP rule (either EXP-Q or EXP-W), with any fixed set of positive parameters β, $\eta \in (0,1)$, and $\gamma_n, a_n, \ n \in N$, is throughput optimal.*

4   The Static Service Split Rule

Suppose a (sub)stochastic matrix $\phi = (\phi_{mn}, m \in M, n = 1, \ldots, N)$ is fixed, which means that $\phi_{mn} \geq 0$ for all m and n, and $\sum_n \phi_{mn} \leq 1$ for every m. Consider a *Static Service Split* (SSS) scheduling rule [1], parameterized by the matrix ϕ. When the channel is in state m, the SSS rule chooses for service a (single) queue n with probability ϕ_{mn}. Clearly, the vector $v = (v_1, \ldots, v_N) = v(\phi)$, where

$$v_n = \sum_m \pi^m \phi_{mn} \mu_n^m \ ,$$

gives the long term average service rates allocated to different flows under ϕ. An SSS rule associated with a stochastic matrix ϕ^* we will call *maximal* if the vector $v(\phi^*)$ is not dominated

by $v(\phi)$ for any other stochastic matrix ϕ. (We say that a vector $v^{(1)}$ is dominated by a vector $v^{(2)}$ if $v_n^{(1)} \leq v_n^{(2)}$ for all n, and the strict inequality $v_n^{(1)} < v_n^{(2)}$ holds for at least one n.) We will denote by V, the set of all service rate vectors $v = v(\phi)$ corresponding to all SSS rules. It is clear to see that V is a convex polyhedron. For a maximal ϕ^*, for each $m \in M$, $\sum_n \phi_{mn}^* = 1$. Also, let $V^* \subset V$ be the subset of *maximal* $v \in V$, i.e., $v^* \in V^*$ is a service rate vector corresponding to a maximal SSS rule. We now present a very useful characterization of a maximal SSS rule. The following result is proved in [1].

Theorem 4.1 *Consider a maximal SSS rule associated with a stochastic matrix ϕ^*. Suppose in addition that all components of $v^* = v(\phi^*)$ are strictly positive. Then there exists a set of strictly positive constants α_n, $n = 1, 2, \ldots, N$ such that for each m, n,*

$$\phi_{mn}^* > 0 \text{ implies } n \in \arg\max_j \alpha_j \mu_j^m . \tag{1}$$

The theorem says that a maximal SSS rule simply chooses for service at any time t a queue n for which $\alpha_n \mu_n^{m(t)}$ is maximal. It does not specify what to do in case of a tie (when $\alpha_j \mu_j^m$ is the same for multiple queues); as a result the same set of $\{\alpha_n\}$ may (and typically will) correspond to different maximal SSS rules. In the rest of this paper, without loss of generality, we will normalize $\alpha_1 = 1$. Let us define b_i by

$$\gamma_n e^{b_n} \equiv \alpha_n , \ n \in N . \tag{2}$$

Note that $b_1 = 0$ (since $\alpha_1 = 1$ by the convention adopted).

Definition 4.1 *We will say that a Complete Resource Pooling (CRP) condition holds for a vector v^*, if it satisfies the following conditions:*

(i) *$v^* \in V^*$, i.e., it is the service rate vector corresponding to a maximal SSS rule.*

(ii) *all components of v^* are strictly positive, and moreover, v^* lies in the interior of one of the faces of V.*

The above definition of CRP is the same as that in [19]. However, for our model (which is a special case of the model in [19]) it is easy to to see that, given condition (i) in the above definition, condition (ii) is equivalent to the following one:

(ii') There exists an SSS rule ϕ^* corresponding to the vector v^* which satisfies the following condition. Choose any non-empty $\mathcal{A} \subseteq N$ such that $N \setminus \mathcal{A}$ is not empty. Then, there exists some channel state $m = m(\mathcal{A}) \in M$, users $i = i(\mathcal{A}) \in \mathcal{A}$ and $j = j(\mathcal{A}) \notin \mathcal{A}$ such that $\phi_{mi}^* > 0$ and $\phi_{mj}^* > 0$.

In other words, for our model,

The CRP condition holds if and only if conditions (i) and (ii') hold.

The notion of Complete Resource Pooling was introduced in earlier papers, in particular in the work on parallel server systems [10, 12, 2, 22], and is closely related to the equivalent workload formulations of Brownian control problems [13, 11]. From a physical viewpoint, CRP means that "service resources" (in our case - channel states) are in a certain way "connected" so that when the system is "heavily loaded," all the resources *can be* "pooled together" and operate as a single aggregate resource.

Consider any v^* (and a corresponding SSS rule ϕ^*) satisfying the CRP condition. From Theorem 4.1, it follows that there exists a set of strictly positive constants $\{\alpha_n,\ n = 1, 2, \ldots, N\}$ such that $\phi^*_{mn} > 0$ implies $n \in \arg\max_j \alpha_j \mu^m_j$. As an aside, it can be shown (using (ii') of the CRP condition) that $\{\alpha_n\}$ corresponding to v^* are unique up to scaling. Let us next define

$$\bar{\mu}^m \equiv \max_{n \in N} \alpha_n \mu^m_n \tag{3}$$

$$K^*(m) \equiv \arg\max_{n \in N} \alpha_n \mu^m_n \tag{4}$$

Lemma 4.1 *(i) For any v^* satisfying the CRP condition, and a corresponding $\{\alpha_n\}$ (defined in Theorem 4.1), we have*

$$\sum_{n \in N} \alpha_n v^*_n = \sum_{m \in M} \pi^m \bar{\mu}^m.$$

(ii) Further, for any other $v \in V$, and with the $\{\alpha_n\}$ defined in (i),

$$\sum_{n \in N} \alpha_n v_n \leq \sum_{m \in M} \pi^m \bar{\mu}^m.$$

Proof: We have

$$\sum_{n \in N} \alpha_n v^*_n = \sum_{n \in N} \alpha_n \sum_{m \in M} \pi^m \phi^*_{mn} \mu^m_n$$

$$= \sum_{m \in M} \pi^m \sum_{n \in N} \alpha_n \mu^m_n \phi^*_{mn}$$

$$= \sum_{m \in M} \pi^m \left[\sum_{\{n : \phi^*_{mn} > 0\}} \alpha_n \mu^m_n \phi^*_{mn} \right],$$

with $\sum_{\{n : \phi^*_{mn} > 0\}} \phi^*_{mn} = 1$. By definition of $\{\alpha_n\}$, for each m and for all n such that $\phi^*_{mn} > 0$, it follows that $\alpha_n \mu^m_n = \bar{\mu}^m$. Thus,

$$\sum_{n \in N} \alpha_n v^*_n = \sum_{m \in M} \pi^m \bar{\mu}^m \left[\sum_{n : \phi^*_{mn} > 0} \phi^*_{mn} \right]$$

$$= \sum_{m \in M} \pi^m \bar{\mu}^m$$

8

The proof of (ii) is straightforward. For brevity, we skip the proof. ∎

To illustrate the CRP condition, we provide the following example.

Example 4.1 *Suppose we have a two user system, and there are two (vector) channel states, each occuring independently, and with equal probability. In channel state 1, the user rates supported by the channel are given by the vector $(1, 3)$ and in state 2, the corresponding vector is given by $(3, 1)$. For such a channel, V is given by the rates corresponding to the convex polyhedron with vertices at $(0, 0)$, $(2, 0)$, $(1.5, 1.5)$ and $(0, 2)$. The rates in V^* are given by the vectors corresponding to the faces $(2, 0) - (1.5, 1.5)$ and $(1.5, 1.5) - (0, 2)$. Consider $v^* = (1.5, 1.5)$ (the vertex). For this vector, we see that (ii') is not satisfied. The rule in such a case is to serve user 1 in channel state 2 and user 2 in channel state 1. On the other hand, for any other $v^* \in V^*$ not on a vertex, we would require that in one of the states, rates need to be allocated to both users (i.e., "split" the service in that channel state), and the assumption will be satisfied.*

5   Heavy Traffic Regime for the Exponential Rule

Consider a fixed vector $v^* \in V^*$ satisfying the CRP condition. Associated with this vector is a maximal SSS rule ϕ^* and a set of coefficients $\{\alpha_i\}$ from Theorem 4.1, with $\alpha_1 = 1$. Let us define the quantity

$$X(t) \equiv \sum_{n=1}^{N} \alpha_n Q_n(t)$$

We will refer to $X(t)$ as the *workload.*

Now, consider a sequence of systems indexed by $r \in \mathcal{R} \equiv \{r_1, r_2, \ldots\}$, where $r_i > 0$ for all i, and $r_i \nearrow \infty$ as $i \to \infty$, and let $A_n^r = (A_n^r(t), \ t = 1, 2, \ldots)$, $n = 1, \ldots, N$, be the associated arrival processes.

Assumption 5.1 *We assume that the arrival processes satisfy the following properties. First, for each r, the sequences A_n^r are mutually independent. In addition, for each n and r,*

$$A_n^r(t), t = 0, 1, 2, \ldots, \ \text{are i.i.d.,} \tag{5}$$

and

$$\mathrm{Var}[\mathrm{A_n^r}(0)] \ \to \ \sigma_n^2 \geq 0, \ \text{as r} \to \infty, \tag{6}$$

Next, we assume that for each $n \in N$, the mean input rate parameter λ_n^r is such that

$$r(\lambda_n^r - v_n^*) \ \to \ c_n, \tag{7}$$

where $c_n \in R$ is a fixed constant. We note that v^* is the limiting mean arrival rate, and for this reason, from now on we will denote $\lambda \equiv v^*$.

Finally, for any $n \in N$ and any $\varepsilon > 0$, there exists a constant $a = a(\varepsilon) > 0$ such that the following (large deviations) estimate holds. Uniformly in $r \in \mathcal{R}$, for all sufficiently large k,

$$P\{|\frac{1}{k}\sum_{t=0}^{k-1} A_n^r(t) - \lambda_n^r| \geq \varepsilon\} < e^{-ak} . \tag{8}$$

Remark 1. The property (8) holds for example when, as $r \to \infty$, the distributions of $A_n^r(1)$ are stochastically dominated by and converge to the distribution of a random variable $A_n(1)$ with finite exponential moment $Ee^{\nu A_n(1)} < \infty$ for some fixed $\nu > 0$.

Under Assumptions (5), (6) and (7), it follows that a functional central limit (FCLT) holds for each input flow:

$$\{r^{-1}(F_n^r(r^2 t) - \lambda_n^r r^2 t), t \geq 0\} \quad \overset{w}{\to} \quad \{\sigma_n B_n(t), t \geq 0\} \tag{9}$$

where

$$F_n^r(t) \quad \equiv \quad \sum_{l=0}^{\lfloor t \rfloor} A_n^r(l)$$

is the cumulative number of flow n customers arrived in the interval $[0, t]$, B_n is a standard Brownian motion and $\overset{w}{\to}$ denotes weak convergence of processes in the Skorohod space $D([0, \infty), R)$.

Remark 2. It will be clear later that, as far as assumptions on the input flows are concerned, our proofs rely only on the FCLT (9) and the following property.
Uniformly in $r \in \mathcal{R}$ and integers $l \geq 0$, for all sufficiently large k,

$$P\{|\frac{1}{k}\sum_{t=0}^{k-1} A_n^r(l+t) - \lambda_n^r| \geq \varepsilon\} < e^{-ak} . \tag{10}$$

Thus, (10) can replace the combination of the iid assumption (5) and assumption (8). In other words, properties (9) and (10) hold in many cases when the input processes are not necessarily i.i.d. For example, they hold in the case when, for each n, the processes $A_n^r(\cdot)$ are Markov chains with the same *finite* state space, and transition probabilities of $A_n^r(\cdot)$ converge to those of an irreducible Markov chain $A_n(\cdot)$.

The Markov chain describing the channel process does not change with r, and is assumed to be independent of the arrival processes. Let us denote

$$\bar{\mu} \quad \equiv \quad \sum_{m \in M} \pi^m \bar{\mu}^m$$

10

where $\bar{\mu}^m$ was defined in (3). From Lemma 4.1, and the fact that $\lambda = v^*$, it follows that

$$\bar{\mu} \;=\; \sum_{n \in N} \alpha_n \lambda_n$$

We note that $\bar{\mu}$ is the maximum possible (average) rate at which the workload can be processed. Then, from the FCLT for Markov chains [9], for any initial state of the Markov chain, we have as $r \to \infty$,

$$\{r^{-1}(\sum_{l=1}^{\lfloor r^2 t \rfloor} \bar{\mu}^{m(l-1)} - \bar{\mu} r^2 t), \; t \ge 0\} \;\; \xrightarrow{w} \;\; \{\sigma_{N+1}^2 B_{N+1}(t), \; t \ge 0\}, \tag{11}$$

where $B_{N+1}(t)$ is a standard Brownian motion, with variance

$$\sigma_{N+1}^2 \;\equiv\; \lim_{l \to \infty} l^{-1} E(\sum_{t=1}^{l} \bar{\mu}^{m(t-1)} - \bar{\mu} l)^2.$$

The superscript of $\bar{\mu}^{m(l-1)}$, i.e. $m(l-1)$, denotes the channel state at (discrete) time $(l-1)$.

We will denote

$$\sigma^2 \;\equiv\; \sigma_{N+1}^2 + \sum_{n \in N} \alpha_n^2 \sigma_n^2$$

For each value of the scaling parameter r, let $Q^r(\cdot)$ and $X^r(\cdot)$ be the corresponding (vector) queue length and workload processes. Let us apply the diffusion scaling to $Q^r(\cdot)$ and $X^r(\cdot)$ to define the following scaled processes:

$$\begin{aligned} \tilde{q}^r(t) &\equiv r^{-1}Q^r(r^2 t), \; t \ge 0 \\ \tilde{x}^r(t) &\equiv r^{-1}X^r(r^2 t), \; t \ge 0. \end{aligned}$$

In general, in the rest of this paper, for a process $Z^r(t)$, we denote

$$\tilde{z}^r(t) \;\equiv\; r^{-1}Z^r(r^2 t), \; t \ge 0. \tag{12}$$

We assume that the initial states of the scaled processes converge in such a way that $\tilde{q}^r(0) \xrightarrow{w} \tilde{q}(0)$ as $r \to \infty$, and $\tilde{q}(0)$ satisfies $a_i \tilde{q}_i(0) = a_j \tilde{q}_j(0)$ for all $1 \le i, j \le N$. Equivalently, we define for each $n \in N$,

$$\alpha_n^0 \;\equiv\; \frac{a_n^{-1}}{\sum_{i=1}^{N} \alpha_i a_i^{-1}}$$

and let α^0 be the corresponding vector. Then, we have $\tilde{q}_n(0) = \alpha_n^0 \tilde{x}(0)$, where $\tilde{x}(0) \ge 0$ is a fixed constant. Finally denote

$$c \;\equiv\; \sum_n \alpha_n c_n,$$

and consider the following one dimensional reflected Brownian motion process $(\tilde{x}(t), t \ge 0)$:

$$\begin{aligned} \tilde{x}(t) &\equiv \tilde{x}(0) + ct + \sigma B(t) + \tilde{y}(t) \tag{13} \\ \tilde{y}(t) &\equiv -[\min(0, \inf_{0 \le u \le t} \{\tilde{x}(0) + cu + \sigma B(u)\})]. \tag{14} \end{aligned}$$

11

Theorem 5.1 *Consider the system in heavy traffic.*

(i) *Suppose the scheduling rule is EXP-Q. Then, as $r \to \infty$,*

$$\tilde{x}^r \quad \xrightarrow{w} \quad \tilde{x}, \tag{15}$$

and moreover,

$$\tilde{q}^r \quad \xrightarrow{w} \quad \tilde{q} \equiv \tilde{x}\alpha^0. \tag{16}$$

(ii) *The EXP-Q rule is asymptotically optimal in that it minimizes the workload process. More precisely, the workload process $\tilde{x}_G$ corresponding to a scheduling policy G is such that for any time $t \geq 0$ and any $u \geq 0$,*

$$\liminf_{r \to \infty} P(\tilde{x}_G^r(t) > u) \quad \geq \quad P(\tilde{x}(t) > u). \tag{17}$$

Let us define $\tilde{q}_{n,G}^r$ to be the (diffusion)-scaled queue length of flow n under a scheduling policy G. Then, in addition to (17), the following properties hold: for any time $t \geq 0$ and $u \geq 0$,

$$\liminf_{r \to \infty} P(\max_n a_n \tilde{q}_{n,G}^r(t) > u) \quad \geq \quad P(\max_n a_n \tilde{q}_n(t) > u), \tag{18}$$

and for any continuous non-decreasing cost function $C(\xi)$, $\xi \geq 0$, and any $t \geq 0$,

$$\liminf_{r \to \infty} E\left[\int_0^t C\left(\max_n a_n \tilde{q}_{n,G}^r(s)\right) ds\right] \quad \geq \quad E\left[\int_0^t C\left(\max_n a_n \tilde{q}_n(s)\right) ds\right]. \tag{19}$$

Remark. Just as in [19], our proof of the main result uses Skorohod representation, i.e. the proof is pathwise. Consequently, statement (ii) of the above theorem can be reformulated as *pathwise optimality* of the EXP rule. Namely, the sequence of processes corresponding to arbirary scheduling rule G and to the EXP rule can be constructed on a common probability space, such that the following holds for any $t \geq 0$:

$$\liminf_{r \to \infty} \tilde{x}_G^r(t) \quad \geq \quad \tilde{x}(t), \tag{20}$$

$$\liminf_{r \to \infty} \max_n a_n \tilde{q}_{n,G}^r(t) \quad \geq \quad \max_n a_n \tilde{q}_n(t), \tag{21}$$

and for any continuous non-decreasing cost function $C(\cdot)$,

$$\liminf_{r \to \infty} \int_0^t C\left(\max_n a_n \tilde{q}_{n,G}^r(s)\right) ds \quad \geq \quad \int_0^t C\left(\max_n a_n \tilde{q}_n(s)\right) ds. \tag{22}$$

6   The Fluid Limit

Let $\lambda = v^*$ satisfy the CRP condition, and consider the sequence of processes indexed by $\{r\}$ as introduced in the previous section. In this section, we study this sequence of processes under the fluid limit scaling and the EXP-Q discipline.

First, we need to define the following additional functions associated with the system for each value of the scaling parameter r. We recall that the functions $F^r(t), Q^r(t)$ and $X^r(t)$ have already been defined. Let

$$\hat{F}_n^r(t) \equiv \sum_{l=1}^{\lfloor t \rfloor} D_n^r(l)$$

denote the number of type-n customers that were served and have departed by time $t \geq 0$. Also, denote by $G_m^r(t)$ the total number of time slots before time t (i.e., among the slots $0, 1, \ldots, t-1$), when the channel was in state m; and by $\hat{G}_{mi}^r(t)$ the number of time slots before time t when the channel state was m and the channel was allocated to serve queue i. We will also denote for $t \geq 0$,

$$H^r(t) \equiv \sum_{l=1}^{\lfloor t \rfloor} \bar{\mu}^{m(l-1)}$$

$$Y^r(t) \equiv H^r(t) - \sum_n \alpha_n \hat{F}_n^r(t)$$

$$= \sum_{l=1}^{\lfloor t \rfloor} [\bar{\mu}^{m(l-1)} - \sum_n \alpha_n D_n^r(l)]$$

and[3]

$$W^r(t) \equiv X^r(0) + \sum_n \alpha_n F_n^r(t) - H^r(t).$$

Then the following relations obviously hold:

$$F_n^r(0) = \hat{F}_n^r(0) = 0, \tag{23}$$

$$Q_n^r(t) = Q_n^r(0) + F_n^r(t) - \hat{F}_n^r(t), \tag{24}$$

$$X^r(t) = W^r(t) + Y^r(t). \tag{25}$$

Finally, let

$$\widehat{Q}_n^r(t) \equiv a_n Q_n^r(t) - b_n \sqrt{\frac{1}{N} \sum_j a_j Q_j^r(t)},$$

---

[3]Note that $W^r(\cdot)$ is not to be confused with the waiting time of user i ($W_i(t)$) that is used in the EXP-W rule definition.

where b_n was defined in (2). We will consider the process $Z^r = (Q^r, \widehat{Q}^r, X^r, W^r, Y^r, F^r, \hat{F}^r, G^r, H^r, \hat{G}^r)$, where

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

$$\widehat{Q}^r = (\widehat{Q}_i^r(t), \quad t \geq 0, \quad i = 1, 2, \ldots, N),$$

$$X^r = (X^r(t), \quad t \geq 0),$$

$$W^r = (W^r(t), \quad t \geq 0),$$

$$Y^r = (Y^r(t), \quad t \geq 0),$$

$$F^r = (F_i^r(t), \quad t \geq 0, \quad i = 1, 2, \ldots, N),$$

$$\hat{F}^r = (\hat{F}_i^r(t), \quad t \geq 0, \quad i = 1, 2, \ldots, N),$$

$$G^r = (G_m^r(t), \quad t \geq 0, \quad m \in M),$$

$$\hat{G}^r = (\hat{G}_{mi}^r(t), \quad t \geq 0, \ m \in M, \ i = 1, 2, \ldots, N),$$

$$H^r = (H^r(t), \quad t \geq 0).$$

Recall that our convention allows us to view the above functions as continuous-time processes defined for all $t \geq 0$, but having constant values in each interval $[t, t+1)$.

Now consider the scaled process $z^r = (q^r, \hat{q}, x^r, w^r, y^r, f^r, \hat{f}^r, g^r, h^r, \hat{g}^r)$, where

$$q^r = (q_i^r(t), \quad t \geq 0, \quad i = 1, 2, \ldots, N),$$

$$\hat{q}^r = (\hat{q}_i^r(t), \quad t \geq 0, \quad i = 1, 2, \ldots, N),$$

$$x^r = (x^r(t), \quad t \geq 0),$$

$$w^r = (w^r(t), \quad t \geq 0),$$

$$y^r = (y^r(t), \quad t \geq 0),$$

$$f^r = (f_i^r(t), \quad t \geq 0, \quad i = 1, 2, \ldots, N),$$

$$\hat{f}^r = (\hat{f}_i^r(t), \quad t \geq 0, \quad i = 1, 2, \ldots, N),$$

$$g^r = (g_m^r(t), \quad t \geq 0, \quad m \in M),$$

$$h^r = (h^r(t), \quad t \geq 0),$$

$$\hat{g}^r = (\hat{g}_{mi}^r(t), \quad t \geq 0, \ m \in M, \ i = 1, 2, \ldots, N),$$

and the scaling is defined as

$$z^r(t) = \frac{1}{r} Z^r(rt) \ .$$

Note that from (24) we get:

$$q_i^r(t) \equiv q_i^r(0) + f_i^r(t) - \hat{f}_i^r(t), \quad t \geq 0, \quad i = 1, 2, \ldots, N. \tag{26}$$

Next, using the large deviations condition (8) from Assumption 5.1, we present a useful result which shows that the arrival rates and channel statistics converge to their respective

14

mean values uniformly over $r^{\frac{1}{4}}$ sized partitions of the interval $[0, rT]$, where $T > 0$. For each r, let us cover the interval $[0, rT]$ with $r^{\frac{3}{4}}T$ equal non-overlapping $r^{\frac{1}{4}}$-long intervals $[(i-1)r^{\frac{1}{4}}, ir^{\frac{1}{4}})$, $1 \le i \le r^{\frac{3}{4}}T$. (Since $r^{\frac{3}{4}}T$ may not be an integer, we should divide into, say, $\lfloor rT/r^{\frac{1}{4}} \rfloor + 1$ intervals. To avoid trivial complications and heavy notation, we assume that $r^{\frac{3}{4}}T$ *is* integer. It will be clear that we do not lose the correctness of the argument.)

Lemma 6.1 *With probability 1, for any rational $T > 0$ and $\varepsilon > 0$, there exists finite $R_0 \equiv R_0(T, \varepsilon)$ such that for all $r > R_0$,*

$$\max_{n \in N,\ 1 \le i \le r^{\frac{3}{4}}T} \left| \frac{F_n^r\left(ir^{\frac{1}{4}}\right) - F_n^r\left((i-1)r^{\frac{1}{4}}\right)}{r^{\frac{1}{4}}} - \lambda_n^r \right| < \varepsilon \ , \tag{27}$$

$$\max_{m \in M,\ 1 \le i \le r^{\frac{3}{4}}T} \left| \frac{G_m^r\left(ir^{\frac{1}{4}}\right) - G_m^r\left((i-1)r^{\frac{1}{4}}\right)}{r^{\frac{1}{4}}} - \pi^m \right| < \varepsilon \ . \tag{28}$$

Proof: The proof is presented in [17]. It uses the large deviations estimate (8) and a large deviations estimate for the finite state Markov chains. ∎

Definition 6.1 *A fixed set of functions $z = (q, \hat{q}, x, w, y, f, \hat{f}, g, h, \hat{g})$ is called a fluid sample path (FSP) if there exists a sequence $\mathcal{R}_f$ of values of r, and a sequence of paths $\{z^r\}$ such that as $r \to \infty$ along sequence $\mathcal{R}_f$,*

$$z^r \quad \to \quad z, \quad u.o.c.,$$

and the following properties hold:

$$\|q(0)\| \quad < \quad \infty,$$
$$(\hat{q}_n^r(t), t \ge 0) \quad \to \quad (a_n q_n(t), t \ge 0) \quad u.o.c.,$$
$$(f_n^r(t), t \ge 0) \to (\lambda_n t, t \ge 0) \quad u.o.c., \tag{29}$$
$$(g_m^r(t), t \ge 0) \to (\pi_m t, t \ge 0) \quad u.o.c., \tag{30}$$

and, finally, for all rational $\varepsilon > 0$ and $T > 0$, for all sufficiently large r (which may depend on T, ε),

$$\text{conditions (27) and (28) hold.} \tag{31}$$

We now state a basic result, whose proof we omit. The details of the proof is analogous to that of Lemma 1 in [1].

Lemma 6.2 *Suppose a sequence $\{r\}$, and the associated sequence of paths $\{z^r\}$, is fixed, such that $\sup \|q^r(0)\| < \infty$, and properties (29),(30), and (31) hold. Then, there exists a subsequence $\{k\} \subseteq \{r\}$ such that the paths $\{z^k\}$ converges uniformly on compact sets (u.o.c.) to an FSP.*

The following lemma states some basic properties of a fluid sample path. We omit the proof for brevity.

Lemma 6.3 *For any fluid sample path z, all its components are Lipschitz continuous, and in addition,*

$$
\begin{aligned}
f_n(t) &= \lambda_n t, \quad t \geq 0, \\
g_m(t) &= \pi^m t, \quad t \geq 0, \\
q_n(t) &= f_n(t) - \hat{f}_n(t), \quad t \geq 0, \\
w(t) &= w(0) = x(0), \quad t \geq 0, \\
x(t) &= \sum_n \alpha_n q_n(t) = x(0) + y(t).
\end{aligned}
$$

Since all the components of the FSP are Lipschitz in $[0, \infty)$, they are absolutely continuous. Therefore, at almost all points $t \in [0, \infty)$ (with respect to Lebesgue measure), the derivatives of all those functions exist. We will call such points *regular*. The vector $q(t)$ corresponding to an FSP we will call its state at time t. The dynamics of $q(t)$ is governed by the differential equation

$$
\frac{d}{dt} q(t) = \lambda - v(t),
$$

which holds at every regular point t, and where the notation $v(t) \equiv \hat{f}'(t)$ is used.

6.1 A "Local" Fluid Limit

In this section, we study convergence of the process sample paths on a "finer" time-scale. The paths we obtain in the limit will be called the "local-fluid-limits." We consider a time interval $[t, t + \delta]$ for some t and $\delta > 0$. For each r, we study the (unscaled) time interval $[rt, rt + r\delta]$. We divide this interval into $\sqrt{r}$ sub-intervals, each of length $\sqrt{r}\delta$, and study appropriately rescaled processes on such sub-intervals. The convergence properties over this time-scale, as we will see later, are essential to prove the main result of this paper.

Let us recall the notation

$$
\hat{q}_n^r(t) \equiv \frac{1}{r} \widehat{Q}_n^r(rt) \equiv \frac{1}{r}[a_n Q_n^r(rt) - b_n \sqrt{\bar{Q}^r(rt)}\,],
$$

where

$$
\bar{Q}^r(t) \equiv \frac{1}{N} \sum_j a_j Q_j^r(t).
$$

Let us fix some $\delta > 0$ and a sequence $\{s^r\}$ of non-negative numbers. Now consider the sequence of unscaled intervals $[rs^r, rs^r + \sqrt{r}\delta]$ and the following processes:

$$
\diamond q_i^r(x) \equiv \sqrt{r}[\hat{q}_i^r(s^r + x/\sqrt{r}) - \max_{k \in N} \hat{q}_k^r(s^r)], \quad x \in [0, \delta], \, i \in N, \tag{32}
$$

16

and denote
$$\diamond\zeta^r(x) \equiv \max_i \diamond q_i^r(x), \ x \in [0, \delta] \ .$$

We note that the process $\diamond q_i^r(\cdot)$ is obtained from the process $Q_i^r(\cdot)$ by the time "speedup" of $\sqrt{r}$ and the "space" scaling by the factor $1/\sqrt{r}$ (in addition to the centering). Thus, it is very natural that as $r \to \infty$, if the sequence $\{s^r\}$ converges to s, and the initial conditions converge in a suitable manner, then the sequence of processes $\{\diamond q_i^r(\cdot), \ i \in N\}$, all defined in the interval $[0, \delta]$, should converge over a subsequence of $\{r\}$ to another - "local" - fluid limit. The proposition formulated below formalizes this observation.

Let us define the following functions, all defined for $x \in [0, \delta]$, and associated with the time interval $[rs^r, rs^r + \sqrt{r}\delta]$:
$\sqrt{r}\diamond f_i^r(x)$ is the number of type i customers arrived in the (unscaled) interval $[rs^r, rs^r + \sqrt{r}x]$;
$\sqrt{r}\diamond \widehat{f}_i^r(x)$ is the number of type i customers served in the (unscaled) interval $[rs^r, rs^r + \sqrt{r}x]$;
$\sqrt{r}\diamond g_i^r(x)$ is the number of (complete) time slots in the (unscaled) interval $[rs^r, rs^r + \sqrt{r}x]$, in which the channel was in state m;
$\sqrt{r}\diamond \widehat{g}_{mi}^r(x)$ is the number of (complete) time slots in the (unscaled) interval $[rs^r, rs^r + \sqrt{r}x]$, in which the channel was in state m and the service was allocated to queue i.

Finally, let us define
$$\epsilon^{**} \equiv (\min_{n,m} \mu_n^m)(\min_n \alpha_n)(\min_m \pi^m) \min\{\phi_{mi}^*, m \in M, i \in N : \phi_{mi}^* > 0\}.$$

and let
$$\epsilon^* \equiv \frac{\epsilon^{**}}{N \max_n \alpha_n} > 0. \tag{33}$$

Proposition 6.1 *Consider a sequence $r \in \mathcal{R}_f$ such that the associated sequence of paths $\{z^r\}, r \in \mathcal{R}_f$ converges to an FSP. Fix any regular point $t > 0$ (for the FSP) and any $\delta > 0$ such that*
$$\max_i a_i q_i(s) > 0, \quad s \in [t, t + \delta]. \tag{34}$$

Suppose that a further subsequence $\mathcal{R}_{f1} \subseteq \mathcal{R}_f$ is fixed, and associated paths $\{z^r, r \in \mathcal{R}_{f1}\}$, and sub-intervals $[rs^r, rs^r + \sqrt{r}\delta]$ satisfy the following conditions:

(i) For some fixed $s \in [t, t + \delta]$,
$$s^r \to s, \tag{35}$$

(ii) The ordering of values of $\hat{q}_i^r(s^r), i \in N$, remains the same. To be definite, without loss of generality, let
$$\hat{q}_1^r(s^r) \geq \ldots \geq \hat{q}_N^r(s^r) \ .$$

(iii) For each $i \in N$,
$$\diamond q_i^r(0) \to \diamond q_i(0) \ ,$$

where $_\diamond q_1(0) = 0$ (by our construction), and each other $_\diamond q_i(0)$ is either finite non-positive or $-\infty$.

Then, the following properties (A) and (B) hold:

Property (A): *There exists a further subsequence $\mathcal{R}_{f2} \subseteq \mathcal{R}_{f1}$ such that along this subsequence, the following additional properties hold for each $i \in N$ and $m \in M$:*

$$(_\diamond f_i^r(x),\ 0 \le x \le \delta) \to (\lambda_i x,\ 0 \le x \le \delta) \quad u.o.c. \tag{36}$$

$$(_\diamond g_i^r(x),\ 0 \le x \le \delta) \to (\pi^m x,\ 0 \le x \le \delta) \quad u.o.c. \tag{37}$$

$$(_\diamond \widehat{g}_{mi}^r(x),\ 0 \le x \le \delta) \to (_\diamond \widehat{g}_{mi}(x),\ 0 \le x \le \delta) \quad u.o.c. \tag{38}$$

$$(_\diamond \widehat{f}_i^r(x),\ 0 \le x \le \delta) \to (_\diamond \widehat{f}_i(x),\ 0 \le x \le \delta) = (\sum_m \mu_i^m {}_\diamond \widehat{g}_{mi}(x),\ 0 \le x \le \delta) \quad u.o.c. \tag{39}$$

$$(_\diamond q_i^r(x),\ 0 \le x \le \delta) \to (_\diamond q_i(x),\ 0 \le x \le \delta) = (_\diamond q_i(0) + a_i(\lambda_i x - {}_\diamond \widehat{f}_i(x)),\ 0 \le x \le \delta) \quad u.o.c. , \tag{40}$$

$$(_\diamond \zeta^r(x),\ 0 \le x \le \delta) \to (_\diamond \zeta(x),\ 0 \le x \le \delta) \quad u.o.c. , \tag{41}$$

where all the functions $_\diamond \widehat{g}_{mi}(\cdot)$ and $_\diamond \widehat{f}_i(\cdot)$ are non-decreasing Lipschitz continuous with value 0 at $x = 0$, and all functions $_\diamond q_i(\cdot)$ $_\diamond \zeta(\cdot)$ are Lipschitz continuous, and in addition

$$\sum_{i=1}^N {}_\diamond \widehat{g}_{mi}(x) = \pi^m x .$$

As before, we will call a point $x \in (0, \delta)$ regular if the derivatives of all the functions $_\diamond \zeta(\cdot)$, $_\diamond q_i(\cdot)$, $_\diamond \widehat{g}_{mi}(\cdot)$, and $_\diamond \widehat{f}_i(\cdot)$, exist in this point. Almost all points (with respect to Lebesgue measure) of the interval $(0, \delta)$ are regular.

Property (B): *At every regular point $x \in (0, \delta)$,*

$$(_\diamond \zeta(x))' \le 0. \tag{42}$$

Suppose in addition, $\{\arg\max_i {}_\diamond q_i(x)\} \ne \{1, 2, \ldots, N\}$. Then,

$$_\diamond \zeta'(x) \le -\epsilon^* . \tag{43}$$

Finally, suppose that for all $n \in N$, $_\diamond q_n(0) > -\infty$. Then, at every regular point $x \in (0, \delta)$,

$$(\min_i {}_\diamond q_n(x))' \ge 0. \tag{44}$$

Proof: The proof of (A) is completely analogous to the proof of Lemma 6.3, which in turn is presented in [1]. The convergence properties (36) and (37) trivially follow from the fact that we consider a sequence of paths that converge to an FSP (and thus, properties (27) and (28) hold).

18

We will now prove (43) of (B). Equations (42) and (44) will follow from a similar reasoning.

For each $r \in \mathcal{R}_{f2}$ consider the (unscaled) time interval $[rs^r, rs^r + \sqrt{r}\delta]$, and consider how the coefficient of $\mu_i(t)$ (in the EXP-Q rule) behaves in this interval. Obviously, multiplying the coefficients of all $\mu_i(t)$ by the same positive function of time (not necessarily a constant), does not change the EXP-Q scheduling rule. Therefore, the following functions $\hat{\gamma}_i^r(\cdot)$ can be regarded as the coefficients of $\mu_i(t)$. For every r, $i \in N$, and $x \in [0, \delta]$, we have

$$\hat{\gamma}_i^r(x) \equiv \gamma_i \exp\left(\frac{a_i Q_i^r(rs^r + \sqrt{r}x) - r\max_{k \in N} \hat{q}_k^r(s^r)}{\beta + \sqrt{\overline{Q}^r(rs + \sqrt{r}x)}} \right)$$

$$= \gamma_i \exp\left(\frac{a_i Q_i^r(rs^r + \sqrt{r}x) - b_i\sqrt{\overline{Q}^r(rs^r + \sqrt{r}x)} - r\max_{k \in N} \hat{q}_k^r(s^r) + b_i\sqrt{\overline{Q}^r(rs^r + \sqrt{r}x)}}{\sqrt{r}\left(\frac{\beta}{\sqrt{r}} + \sqrt{\frac{\overline{Q}^r(rs + \sqrt{r}x)}{r}} \right)} \right)$$

$$= \gamma_i \exp\left(\frac{b_i\sqrt{\frac{\overline{Q}^r(rs^r + \sqrt{r}x)}{r}}}{\frac{\beta}{\sqrt{r}} + \sqrt{\frac{\overline{Q}^r(rs^r + \sqrt{r}x)}{r}}} \right) \exp\left(\frac{{}_\diamond q_i^r(x)}{\frac{\beta}{\sqrt{r}} + \sqrt{\frac{\overline{Q}^r(rs^r + \sqrt{r}x)}{r}}} \right) .$$

Note that as $r \to \infty$ the convergence

$$\frac{\overline{Q}^r(rs^r + \sqrt{r}x)}{r} \longrightarrow \overline{q}(s)$$

is uniform on $x \in [0, \delta]$, where (from (34))

$$\overline{q}(s) \equiv \frac{1}{N} \sum_i a_i q_i(s) > 0 .$$

Therefore, we have

$$(\hat{\gamma}_i^r(x), \ 0 \le x \le \delta) \to (\alpha_i \exp({}_\diamond q_i(x)/\sqrt{\overline{q}(s)}), \ 0 \le x \le \delta) \quad u.o.c. . \tag{45}$$

Now, consider any regular point $x \in (0, \delta)$ such that $\{\arg\max_i {}_\diamond q_i(x)\} \ne \{1, 2, \ldots, N\}$. Consider the subset of flows $I^{**} \subseteq N$ for which ${}_\diamond q_i(x)$ is maximal, i.e. ${}_\diamond q_i(x) = {}_\diamond \zeta(x)$. By assumption, there is at least one flow $j \notin I^{**}$. Also, let us denote

$$M^{**} \equiv \{m \in M \mid \phi_{mi}^* > 0 \text{ for at least one } i \in I^{**}\} .$$

For every $m \in M^{**}$, let us pick an element $i(m) \in I^{**}$ for which $\phi_{mi}^* > 0$. Observe that the value of $\alpha_{i(m)}\mu_{i(m)}^m$ will be the same regardless of which of those elements we pick. From the form of the EXP-Q scheduling rule and the uniform convergence (45), we can make the following observation. There exists a small $\epsilon_1 > 0$ such that for any $z \in (x, x + \epsilon_1)$ and any $\epsilon_2 > 0$, we have the following estimate for all sufficiently large r,

$$\sum_{i \in I^{**}} \alpha_i({}_\diamond \widehat{f}_i^r(z) - {}_\diamond \widehat{f}_i^r(x)) =$$

19

$$\sum_{i\in I^{**}} \alpha_i \sum_m \mu_i^m (_\diamond g_{mi}^r(z) - _\diamond g_{mi}^r(x)) \geq$$

$$\sum_{m\in M^{**}} \sum_{i\in I^{**}} \alpha_i \mu_i^m (_\diamond g_{mi}^r(z) - _\diamond g_{mi}^r(x)) \geq$$

$$\sum_{m\in M^{**}} \pi^m (1-\epsilon_2)(z-x)\alpha_{i(m)}\mu_{i(m)}^m =$$

$$\sum_{m\in M^{**}} \pi^m (1-\epsilon_2)(z-x)(\sum_{i\in I^{**}} \phi_{mi}^* \alpha_i \mu_i^m + (1-\sum_{i\in I^{**}} \phi_{mi}^*)\alpha_{i(m)}\mu_{i(m)}^m) =$$

$$\sum_{m\in M^{**}} \pi^m (1-\epsilon_2)(z-x)\sum_{i\in I^{**}} \phi_{mi}^* \alpha_i \mu_i^m + \sum_{m\in M^{**}} \pi^m (1-\epsilon_2)(z-x)(1-\sum_{i\in I^{**}} \phi_{mi}^*)\alpha_{i(m)}\mu_{i(m)}^m$$

Now, by definition of $v_i(\phi^*)$, it follows that

$$\sum_{m\in M^{**}} \pi^m (1-\epsilon_2)(z-x)\sum_{i\in I^{**}} \phi_{mi}^* \alpha_i \mu_i^m = (1-\epsilon_2)(z-x)\sum_{i\in I^{**}} \alpha_i v_i(\phi^*) .$$

From the CRP condition, it follows that for some $m \in M^{**}$, $j \notin I^{**}$, $\phi_{mj}^* > 0$. Thus, it follows that

$$\sum_{m\in M^{**}} \pi^m (1-\sum_{i\in I^{**}} \phi_{mi}^*)\alpha_{i(m)}\mu_{i(m)}^m \geq \epsilon^{**}$$

Since $\epsilon_2 > 0$ is arbitrary, this implies that for any $z \in (x, x+\epsilon_1)$,

$$\sum_{i\in I^{**}} \alpha_i (_\diamond \widehat{f}_i(z) - _\diamond \widehat{f}_i(x)) \geq (z-x)\sum_{i\in I^{**}} \alpha_i v_i(\phi^*) + (z-x)\epsilon^{**} ,$$

and therefore

$$\sum_{i\in I^{**}} \alpha_i {}_\diamond \widehat{f}_i'(x) \geq \sum_{i\in I^{**}} \alpha_i v_i(\phi^*) + \epsilon^{**} .$$

Then, $_\diamond \widehat{f}_i'(x) \geq v_i(\phi^*) + \epsilon^*$ for at least one $i \in I^{**}$, and therefore $_\diamond q_i'(x) \leq -\epsilon^*$ holds for this i. Point x is regular. By definition, in any regular point the derivatives $_\diamond \zeta'(\cdot)$ and $_\diamond q_i'(\cdot)$ for all $i \in I^{**}$ are all equal. Thus

$$_\diamond \zeta'(x) \leq -\epsilon^* , \tag{46}$$

and we are done.

A similar method can be adopted to show (42) and (44). For brevity, we will skip the details. ∎

6.2 Properties of the FSP

Finally, we present the main result of this section. Consider a fixed FSP with non-zero initial state $q(0)$.

Theorem 6.1 *At every regular point $t > 0$,*

$$\max_i a_i q_i(t) > 0 \quad implies \quad (\max_i a_i q_i(t))' \leq 0. \tag{47}$$

Further, if $\{\arg\max_i a_i q_i(t)\} \neq \{1, 2, \ldots, N\}$, then

$$\max_i a_i q_i(t) > 0 \quad implies \quad (\max_i a_i q_i(t))' \leq -\epsilon^*. \tag{48}$$

Consequently, there exists $K > 0$ such that with $T = K\|q(0)\| > 0$, $a_i q_i(t) = a_j q_j(t)$ for all $i, j \in N$ for any $t \geq T$.

Finally, for all regular $t \geq T$,

$$(\max_i a_i q_i(t))' \;=\; 0. \tag{49}$$

Proof: We will now prove (48). Equations (47) and (49) will follow from a similar reasoning.

Fix an FSP z, and fix any subsequence $\mathcal{R}_f \subseteq \mathcal{R}$ and the corresponding sequence $\{z_r, r \in \mathcal{R}_f\}$ converging to z. Now, let us assume that for some fixed regular point $t > 0$ and a constant $\eta_1 < \epsilon^*$, $\zeta(t) \equiv \max_i \tilde{q}_i(t) \equiv \max_i a_i q_i(t) > 0$ and $\zeta'(t) > -\eta_1$. Further, assume that there is some flow j such that $j \notin \arg\max_i a_i q_i(t)$. Let us prove that this assumption leads to a contradiction.

If the assumption holds, then there exist rational constants $\delta > 0$, $\xi > 0$, and $\eta_2 \in (\eta_1, \epsilon^*)$, such that

$$\zeta(s) > \xi, \quad \forall s \in [t, t + \delta],$$

and

$$\frac{\zeta(t + \delta) - \zeta(t)}{\delta} \;>\; -\eta_2 \;.$$

and $j \notin \arg\max_i a_i q_i(s) \;\; \forall s \in [t, t + \delta]$.

For each r, let us now divide the interval $[t, t + \delta]$ into $\sqrt{r}$ intervals, each of length $\frac{\delta}{\sqrt{r}}$. (As before, to avoid trivial complications, we assume that $\sqrt{r}$ *is* integer. It will be clear that we do not lose the correctness of the argument.) Note that in the "unscaled time" (i.e. on the time scale of the original process S), each sub-interval is of length $\delta\sqrt{r}$.

As before, let us denote

$$\zeta^r(t) \equiv \max_i \hat{q}_i^r(t),$$

and fix any constant $\eta \in (\eta_2, \epsilon^*)$. From the Dirichlet principle[4], for all sufficiently large r, in at least one of the sub-intervals (of length $\frac{\delta}{\sqrt{r}}$), the average rate of change of $\zeta^r(.)$ is greater than or equal to $(-\eta)$. We pick such a sub-interval $[s^r, s^r + \delta/\sqrt{r}]$ for each r. Let us choose a further subsequence $\mathcal{R}_{f1} \subseteq \mathcal{R}_f$ such that, along this subsequence, $s^r \to s$ for some fixed $s \in [t, t+\delta]$. Obviously, the right end-point $s^r + \delta/\sqrt{r}$ of the sub-interval also converges to s.

We choose a further subsequence $\mathcal{R}_{f2} \subseteq \mathcal{R}_{f1}$ such that the order of values of $\hat{q}_i^r(s^r), i \in N$, remains the same. Without loss of generality let us assume that

$$\hat{q}_1^r(s^r) \geq \ldots \geq \hat{q}_N^r(s^r) \ .$$

Finally, for each $i \in N$, consider the following processes:

$$_\diamond q_i^r(x) \equiv \sqrt{r}[\hat{q}_i^r(s^r + x/\sqrt{r}) - \hat{q}_1^r(s^r)], \ x \in [0, \delta] \ ,$$

$$_\diamond \zeta^r(x) \equiv \max_i {}_\diamond q_i^r(x), \ x \in [0, \delta] \ ,$$

and choose a subsequence $\mathcal{R}_{f3} \subseteq \mathcal{R}_{f2}$ such that for each i,

$$_\diamond q_i^r(0) \to {}_\diamond q_i(0) \ ,$$

where $_\diamond q_1(0) = 0$ (by our construction), and each other $_\diamond q_i(0)$ is either finite non-positive or $-\infty$. Indeed by our construction, it follows that $_\diamond q_j(0) = -\infty$, where $j \notin \arg\max_i a_i q_i(t)$.

As we have satisfied conditions (i)-(iii) of Proposition 6.1, it follows that there exists a further subsequence $\mathcal{R}_{f4} \subseteq \mathcal{R}_{f3}$ along which

$$(_\diamond \zeta^r(x), \ 0 \leq x \leq \delta) \ \to \ (_\diamond \zeta(x), \ 0 \leq x \leq \delta) \quad u.o.c. \ ,$$

and $_\diamond \zeta(.)$ is differentiable almost everywhere (w.r.t. the Lebesgue measure). Further, from our construction, it follows that

$$_\diamond \zeta(\delta) - {}_\diamond \zeta(0) \ \geq \ -\eta\delta \ . \tag{50}$$

Also, it is clear that for all $x \in [0, \delta]$, $j \notin \arg\max_i {}_\diamond q_i(x)$. Thus, from (43) of Proposition 6.1, it follows that at each regular point $x \in (0, \delta)$, $_\diamond \zeta'(x) \leq -\epsilon^*$. This implies that

$$_\diamond \zeta(\delta) - {}_\diamond \zeta(0) \ \leq \ -\epsilon^*\delta \tag{51}$$

which contradicts (50). Thus, we are done. ∎

---

[4]This is also known as the pigeonhole principle, or the Dirichlet Box principle. In the discrete case, this principle states given n boxes and $m > n$ objects, at least one box must contain more than one object.

7 The Heavy Traffic Regime: Proof of Theorem 5.1

For each $r \in \mathcal{R}$ consider the following process, obtained by a diffusion scaling (defined by (12)):
$$(\tilde{q}^r, \tilde{x}^r, \tilde{w}^r, \tilde{y}^r, \tilde{f}^r, \tilde{g}^r, \tilde{h}^r) \ .$$
To prove properties (15)-(17), we will use Skorohod's representation theorem (see, for example, [14]), and show that on the appropriately chosen probability space, (15)-(17) hold with probability 1.

According to Skorohod representation theorem for each n, the sequence of the input processes $\{F_n^r\}$ can be constructed on a probability space such that the convergence in (9) holds u.o.c. with probability 1 (w.p.1). Similarly, the sequence of channel state processes (Markov chains) $\{m^r\}$ can be constructed on a probability space such that the convergence in (11) holds u.o.c. w.p.1., which can be written as

$$(\tilde{h}^r(t) - \bar{\mu}rt, \ t \geq 0) \ \stackrel{u.o.c.}{\to} \ (\sigma_{N+1}^2 B_{N+1}(t), \ t \geq 0) \ . \tag{52}$$

We can and do assume that the underlying probability space $(\Omega, \mathcal{F}, P)$ is a direct product of those $N + 1$ probability spaces. (Here, we use the fact that all arrival processes and channel state processes are mutually independent.) Without loss of generality we can also assume that this probability space is complete. By ω we will denote elements ("elementary outcomes") of Ω.

From Assumption 5.1 and the fact that the channel process is a finite state Markov chain, it follows (the proof is essentially the same as that of Lemma 6.1, which in turn is presented in [17]) that the following properties also hold with probability 1, for any fixed rational $T_3 > 0$:

$$\max_{n \in N, \ 1 \leq i \leq T_3 r^{7/4}} \left| \frac{F_n^r\left(ir^{\frac{1}{4}}\right) - F_n^r\left((i-1)r^{\frac{1}{4}}\right)}{r^{\frac{1}{4}}} - \lambda_j^r \right| \ \to \ 0, \tag{53}$$

$$\max_{m \in M, \ 1 \leq i \leq T_3 r^{7/4}} \left| \frac{G_m^r\left(ir^{\frac{1}{4}}\right) - G_m^r\left((i-1)r^{\frac{1}{4}}\right)}{r^{\frac{1}{4}}} - \pi^m \right| \ \to \ 0, \tag{54}$$

$$\max_{1 \leq i \leq T_3 r^{7/4}} \left| \frac{H^r\left(ir^{\frac{1}{4}}\right) - H^r\left((i-1)r^{\frac{1}{4}}\right)}{r^{\frac{1}{4}}} - \bar{\mu} \right| \ \to \ 0, \tag{55}$$

where (55) follows from (54).

We also have
$$\tilde{x}^r(t) = \tilde{w}^r(t) + \tilde{y}^r(t), \ t \geq 0$$
and
$$(\tilde{w}^r(t), \ t \geq 0) \ \stackrel{u.o.c.}{\to} \ (\tilde{w}(t), \ t \geq 0) \ ,$$
where
$$\tilde{w}^r(t) \equiv \tilde{x}^r(0) + \sum_n \alpha_n \tilde{f}_n^r(t) - \tilde{h}^r(t)$$

and

$$\tilde{w}(t) \equiv \tilde{x}(0) + at + \sigma B(t)$$

is a continuous function.

In the rest of this section, we restrict ourselves to a (measurable, probability 1) subset of elementary outcomes $\omega \in \Omega$, such that all the probability 1 properties specified above hold.

For each $r \in \mathcal{R}$, $\tilde{y}^r$ is a non-decreasing RCLL function. Therefore, for any fixed $\omega \in \Omega$, from any subsequence $\mathcal{R}_3(\omega) \subseteq \mathcal{R}$ (which may depend on ω!) it is always possible to find a further subsequence $\mathcal{R}_4(\omega) \subseteq \mathcal{R}_3(\omega)$ such that

$$\tilde{y}^r \Rightarrow \tilde{y} , \tag{56}$$

where $\tilde{y}$ is a non-decreasing RCLL function, which may take the values $+\infty$ (in other words, $\tilde{y} \in D([0, \infty), \bar{R})$, "$\Rightarrow$" denotes convergence in every point of continuity of the limit function (except maybe point 0) and $\bar{R}$ denotes the extended real line, that is, $\bar{R} \equiv R \bigcup \{-\infty, +\infty\}$. We note that (56) implies that

$$\tilde{x}^r \Rightarrow \tilde{x} \equiv \tilde{w} + \tilde{y} , \tag{57}$$

and therefore $\tilde{x}(t) < \infty$ if and only if $\tilde{y}(t) < \infty$.

The following lemma contains the key observation which will be used in the proof of Theorem 5.1. The statement of this lemma is analogous to Lemma 7 in [19]. However, our proof is quite different from that in [19]. In our proof, we resort to the additional local fluid limit on the $\sqrt{r}$ time-scale, in addition to the conventional fluid time-scale. Thus, we use *three* time-scales, a diffusion time-scale, and two fluid time-scales here.

Lemma 7.1 *Suppose, $\omega \in \Omega$ and a subsequence $\mathcal{R}_4(\omega) \subseteq \mathcal{R}$ are fixed such that, along this subsequence, (56) holds. Suppose, a sequence $\{\tilde{t}^r, \ r \in \mathcal{R}_4(\omega)\}$ is fixed such that*

$$\tilde{t}^r \to t' \geq 0 ,$$

and

$$\tilde{x}(\tilde{t}^r) \to C > 0 .$$

Let a rational $\delta > 0$ be fixed, and

$$\epsilon = \epsilon(\delta, t') = \sup_{\xi_1, \xi_2 \in [t'-\delta, t'+\delta] \cap R_+} |\tilde{w}(\xi_1) - \tilde{w}(\xi_2)| < C .$$

Then,
(a) $\tilde{y}$ (and $\tilde{x}$) is finite in $[0, t' + \delta)$;
(b) $\tilde{y}$ does not increase in $(t', t' + \delta)$, i.e., $\tilde{y}(t' + \delta-) - \tilde{y}(t') = 0$;
(c)   the following bound holds

$$C - \epsilon \leq \tilde{x}(t) \leq CC_2 + \epsilon , \quad \forall t \in [t', t' + \delta) ,$$

with $C_2 = 1 + \left\lceil \frac{2 \max_i a_i}{\epsilon^* \min_i \alpha_i} \right\rceil (\sum_n \alpha_n \lambda_n)$;

(d)   for any $\eta > 0$,

$$\left(\tilde{q}^r(t), \ t \in [t' + \eta, t' + \delta - \eta]\right) \overset{u.o.c.}{\to} \left(\tilde{q}(t), \ t \in [t' + \eta, t' + \delta - \eta]\right),$$

where $\tilde{q}(t) = \tilde{x}(t)\alpha^0$.

If, in addition, $\tilde{t}^r = t'$ for all r, and $\tilde{q}^r(t') \to \tilde{q}(t')$, then

(b') $\tilde{y}$ does not increase in $[t', t' + \delta)$, i.e., $\tilde{y}(t' + \delta-) - \tilde{y}(t'-) = 0$, where $\tilde{y}(0-) = 0$ by convention;

(d')   the following holds:

$$\left(\tilde{q}^r(t), \ t \in [t', t' + \delta]\right) \overset{u.o.c.}{\to} \left(\tilde{q}(t), \ t \in [t', t' + \delta]\right).$$

Proof. The key construction in this proof, namely the construction of a set of processes $\diamond q^{r,l,p}(\cdot)$ (see below) on a slower - fluid - time scale. This is similar to Bramson's construction in Section 5 of [3], where a slower-fluid scale was used. In our proof, we use two fluid time-scales.

Let us choose a fixed $T > 0$ as follows. Define $C_1 \equiv \frac{\max_i a_i}{\min_i \alpha_i} C$. Then define $T_1 \equiv \frac{C_1}{\epsilon^*}$, where ϵ^* is defined in Theorem 6.1. Let $T \equiv \lceil 2T_1 \rceil$. Now consider for each r, the unscaled time-interval $(r^2 \tilde{t}^r, r^2(\tilde{t}^r + \delta))$. We partition this interval into $\frac{r\delta}{T}$ intervals, each of length rT. We index these intervals by $l \in [0, \frac{r\delta}{T} - 1]$. We further partition each such interval into $\sqrt{r}$ sub-intervals, each of length $\sqrt{r}T$. We index any sub-interval by the pair (l, p), where l corresponds to the index of the interval of length rT and $p \in [0, \sqrt{r} - 1]$ corresponds to the pth sub-interval of the lth interval. Thus, (l, p) corresponds to a sub-interval of length $\sqrt{r}T$. We say $(l_1, p_1) > (l_2, p_2)$ if $l_1 > l_2$ or $l_1 = l_2$ and $p_1 > p_2$.

For each $n \in N, r, (l, p)$, define for $u \in [0, T]$,

$$\diamond \bar{q}_n^{r,l,p}(u) \equiv \frac{1}{\sqrt{r}} \left[\begin{array}{c} \left(a_n Q_n^r(r^2 \tilde{t}^r + rlT + \sqrt{r}pT + \sqrt{r}u) - b_n \sqrt{\bar{Q}^r(r^2 \tilde{t}^r + rlT + \sqrt{r}pT + \sqrt{r}u)}\right) \\ - \max_{k \in N} \left(a_k Q_k^r(r^2 \tilde{t}^r + rlT + \sqrt{r}pT) - b_n \sqrt{\bar{Q}^r(r^2 \tilde{t}^r + rlT + \sqrt{r}pT)}\right) \end{array} \right]$$

where

$$\bar{Q}^r(.) = \frac{1}{N} \sum_{j \in N} a_j Q_j^r(.)$$

We observe that this definition is analogous to (32) in Section 6.1. Let us denote $\diamond \bar{q}^{r,l,p}(.) = [\diamond \bar{q}_1^{r,l,p}(.), \ldots, \diamond \bar{q}_N^{r,l,p}(.)]$. Finally, define for $u \in [0, T]$,

$$G\left(\diamond \bar{q}^{r,l,p}(u)\right) \equiv \max_n \diamond \bar{q}_n^{r,l,p}(u) - \min_n \diamond \bar{q}_n^{r,l,p}(u)$$

$$\diamond \bar{\zeta}^{r,l,p}(u) \equiv \max_n \diamond \bar{q}_n^{r,l,p}(u)$$

Now fix any $\epsilon_1 > 0$. The following properties then hold.

Property 1. For all $r \in \mathcal{R}_4(\omega)$ large enough, for $l = 0$, there exists some $p_1^r \in [0, \sqrt{r} - 1]$ and $u \in [0, T]$, such that

$$G\left({}_\diamond \bar{q}^{r,0,p_1^r}(u)\right) \leq \epsilon_1 \tag{58}$$

$$G\left({}_\diamond \bar{q}^{r,0,p_1^r}(T)\right) \leq 2\epsilon_1 \tag{59}$$

We first show (58) by contradiction. Suppose it does not hold. Then, there exists a subsequence $\mathcal{R}_5(\omega) \subset \mathcal{R}_4(\omega)$ such that for all $r \in \mathcal{R}_5(\omega)$, for all $p \in [0, \sqrt{r} - 1]$ and $u \in [0, T]$, $G({}_\diamond \bar{q}^{r,0,p}(u)) > \epsilon_1$. For this subsequence, define for each $n \in N$,

$$\bar{q}_n^r(t) = \frac{1}{r} Q^r(r^2 \tilde{t}^r + rt) \quad t \in [0, T] \tag{60}$$

It is easily shown that we can choose a further subsequence $\mathcal{R}_6(\omega) \subseteq \mathcal{R}_5(\omega)$ along which the paths $\{\bar{q}_n^r(t), n \in N, t \in [0, T]\}$ converge to an FSP $\{\bar{q}_n(t), n \in N, t \in [0, T]\}$. From Theorem 6.1, and by the definition of T, it follows that for $t > \frac{T}{2}$, $a_i \bar{q}_i(t) = a_j \bar{q}_j(t)$ for all $i, j \in N$, and for all regular $t \geq \frac{T}{2}$,

$$(\max_i a_i \bar{q}_i(t))' = 0 \tag{61}$$

Next, by assumption, along the subsequence $r \in \mathcal{R}_5(\omega)$, we have

$$G({}_\diamond \bar{q}^{r,0,p}(u)) > \epsilon_1, \ p \in [0, \sqrt{r} - 1]. \tag{62}$$

Thus, for any sequence of sub-intervals (each of length $\sqrt{r}T$) indexed by $\{p^r\}$ which are chosen such that conditions (i)-(iii) of Proposition 6.1 hold, we have that along a further subsequence, ${}_\diamond \bar{\zeta}^{r,0,p^r}(.)$ converges to a "local" fluid limit ${}_\diamond \bar{\zeta}(.)$, which from (43) and (62) satisfies

$$_\diamond \bar{\zeta}'(x) \leq -\epsilon^* . \tag{63}$$

An argument analogous to Theorem 6.1 based on the Dirichlet principle can now be used to show that (61) and (63) lead to a contradiction. Finally, equation (59) now follows from (42), (44) and (41). We skip the details for brevity. $\diamond$

Next, it can be easily shown that for any fixed $\epsilon_2 > 0$, for r large enough,

$$\frac{1}{r} X^r(r^2 \tilde{t}^r + p_1^r T \sqrt{r}) \leq (CC_2 + \epsilon_2)$$

with

$$C_2 \equiv 1 + \left\lceil \frac{2 \max_i a_i}{\epsilon^* \min_i \alpha_i} \right\rceil \left(\sum_n \alpha_n \lambda_n \right).$$

This bound follows from the definition of T and assuming (this is the worst case) that no service is given to *any* queue over the first sub-interval of length rT.

Now, define for $u \in [0, T]$,

$$\bar{y}^{r,l,p}(u) = \frac{1}{r}\left(Y^r(r^2\tilde{t}^r + rlT + \sqrt{r}pT + \sqrt{r}u) - Y^r(r^2\tilde{t}^r + rlT + \sqrt{r}pT)\right) \tag{64}$$

Property 2. For all $r \in \mathcal{R}_4(\omega)$ large enough, for all $(l, p) > (0, p_1^r)$, where p_1^r is defined in Property 1, and $u \in (0, T)$,

$$G\left(_\diamond\bar{q}^{r,l,p}(0)\right) \leq 2\epsilon_1 \tag{65}$$
$$G\left(_\diamond\bar{q}^{r,l,p}(u)\right) \leq 3\epsilon_1 \tag{66}$$

We will prove this by contradiction. Let us assume Property 2 does not hold. Then, we can choose a further subsequence $\mathcal{R}_7(\omega) \subseteq \mathcal{R}_4(\omega)$, and for each $r \in \mathcal{R}_5(\omega)$, we define (l_2^r, p_2^r) to be the first index after $(0, p_1^r)$ such that at least one of the following conditions occur:

(i) $G\left(_\diamond\bar{q}^{r,l,p}(0)\right) > 2\epsilon_1$, or

(ii) $G\left(_\diamond\bar{q}^{r,l,p}(u)\right) > 3\epsilon_1$ for some $u \in (0, T)$.

If (i) occurs, then, we let (l_3^r, p_3^r) be the index of the interval *preceding* (l_2^r, p_2^r), else, we let $(l_3^r, p_3^r) = (l_2^r, p_2^r)$. This construction implies the following three properties:

(p1) For all $r \in \mathcal{R}_7(\omega)$, we have

$$G\left(_\diamond\bar{q}^{r,l_3^r,p_3^r}(0)\right) \leq 2\epsilon_1 \tag{67}$$

(p2) For all $r \in \mathcal{R}_7(\omega)$, one of the following two cases occur:

$$G\left(_\diamond\bar{q}^{r,l_3^r,p_3^r}(T)\right) > 2\epsilon_1 \quad \text{or} \tag{68}$$
$$G\left(_\diamond\bar{q}^{r,l_3^r,p_3^r}(u)\right) > 3\epsilon_1 \quad \text{for some } u \in (0, T) \tag{69}$$

(p3) For $(l, p) > (0, p_1^r)$ and $(l, p) < (l_3^r, p_3^r)$, and all $u \in [0, T]$, we have

$$G\left(_\diamond\bar{q}^{r,l_3^r,p_3^r}(u)\right) \leq 3\epsilon_1 \tag{70}$$

Now, we observe that from *(p3)*, it follows that for ϵ_1 small enough, and $(0, p_1^r) < (l, p) < (l_3^r, p_3^r)$, $\bar{y}^{r,l,p}(T) - \bar{y}^{r,l,p}(0) = 0$. This follows from the fact that if $G(_\diamond\bar{q}^{r,l,p}(.))$ is small enough, then, for each channel state m, only decisions from $K^*(m)$ (recall that $K^*(m) = \arg\max_n \alpha_n \mu_n^m$) can be chosen. Thus, it easily follows that for r large enough,

$$C - \epsilon < \frac{1}{r}X^r(r^2\tilde{t}^r + l_3^r Tr + p_3^r T\sqrt{r}) \leq CC_2 + \epsilon_2 \tag{71}$$

For the sequence of processes $\{_\diamond\bar{q}^{r,l_3^r,p_3^r}(.)\}$, is is easy to show that properties analogous to Proposition 6.1 hold. In particular, we can choose a subsequence $\mathcal{R}_8(\omega) \subseteq \mathcal{R}_7(\omega)$, such that the left-end points of the intervals (l_3^r, p_3^r) converge, and satisfy (i)-(iii) of Proposition 6.1. Then, along a further subsequence, we can show that $_\diamond\bar{q}^{r,l_3^r,p_3^r}(.)$ converges uniformly over

27

$[0, T]$ to $_\diamond \bar{q}(.)$, a "local" fluid limit, and this limit satisfies the analog of Property A and Property B of Proposition 6.1.

Also, it follows from *(p1)* and *(p2)* that this limit satisfies $G\left(_\diamond \bar{q}(0)\right) \leq 2\epsilon_1$ and one of

$$
\begin{aligned}
G\left(_\diamond \bar{q}(T)\right) &\geq 2\epsilon_1 \quad \text{or} \\
G\left(_\diamond \bar{q}(u)\right) &\geq 3\epsilon_1 \quad \text{for some } u \in [0, T]
\end{aligned}
$$

which contradicts the analog of (42), (43) and (44) in Proposition 6.1. Thus, we are done. $\diamond$

Property 2 implies that for all large r, for all $l \geq 1$, for all $p \in [0, \sqrt{r} - 1]$

$$
\bar{y}^{r,l,p}(T) - \bar{y}^{r,l,p}(0) = 0. \tag{72}
$$

This implies (a)-(c). Statement (d) follows from Property 2 and the fact that ϵ_1 can be chosen to be arbitrarily small.

To prove properties (b') and (d'), we use the same exact construction. ∎

7.1   Proof of Theorem 5.1

To prove the convergences (15) and (16) in Statement (i) of the theorem, it will suffice to prove the following

Proposition 7.1 *As $r \to \infty$ (along $\mathcal{R}$), with probability 1 we have the following convergences*

$$
(\tilde{y}^r(t), \ t \geq 0) \overset{u.o.c.}{\to} (\tilde{y}(t), \ t \geq 0) , \tag{73}
$$

where $\tilde{y}$ is defined by (14), and

$$
(\tilde{q}^r(t), \ t \geq 0) \overset{u.o.c.}{\to} (\tilde{q}(t), \ t \geq 0) , \tag{74}
$$

where $\tilde{q} = \tilde{x}\alpha^0$.

Proof of Proposition 7.1: The proof uses Lemma 7.1. The details are identical to those in the proof of Theorem 1 (statement 1) in [19]. ∎

Proof of property (17) in Statement (ii) of Theorem 5.1 is identical to the reasoning in the proof of Theorem 1 (statement 2) in [19], which basically follows from the properties of the solution to the one-dimensional Skorohod problem (see Proposition B.1 of [2]). Property (18) follows directly from (17). Finally, (19) follows from (18) by using Fatou's lemma.

8 Conclusion

In this paper, we have studied a scheduling algorithm (the EXP rule) for a wireless channel. This is an "on-line" algorithm, which only uses the information on the current state of the channel and the queues. We have studied the EXP rule for a heavily-loaded system, under a complete resource pooling (CRP) condition. We show that, in the heavy traffic limit, this algorithm minimizes the workload and *induces a state space collapse* such that the queue length vector is always proportional to the vector $[a_1^{-1}, \ldots, a_N^{-1}]$, where a_i's are positive parameters which can be chosen arbitrarily. (As a result, the algorithm minimizes the maximum of the weighted queue lengths $\max_i a_i Q_i(t)$.)

The fact that workload minimization and a state space collapse can be achieved by an on-line rule has been demonstrated in previous work. We believe an important contribution of this paper is that we show the existence an on-line scheduling rule which, in addition, induces state space collapse *along the desired line*. This latter feature of the algorithm is important for applications. Our technique of using a "local fluid limit" in the proof of state space collapse may be of independent interest.

References

[1] 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.

[2] S. L. Bell, R. J. Williams. Dynamic Scheduling of a System with Two Parallel Servers in Heavy Traffic with Complete Resource Pooling: Asymptotic Optimality of a Continuous Review Threshold Policy. *Annals of Applied Probability*, Vol. 11, (2001), pp. 608-649.

[3] M. Bramson. State Space Collapse with Applications to Heavy Traffic Limits for Multiclass Queueing Networks. *Queueing Systems*, Vol. 30, (1998), pp. 89-148.

[4] M. Bramson, J. G. Dai. Heavy Traffic Limits for Some Queueing Networks. *Annals of Applied Probability*, Vol. 11, (2001), pp. 49-88.

[5] R. Buche, H. J. Kushner. Control of Mobile Communications with Time-Varying Channels in Heavy Traffic. *IEEE Transactions on Automatic Control*, Vol. 47, No. 6, June 2002.

[6] W. C. Jakes, "*Microwave Mobile Communications*," Wiley-Interscience, 1974.

[7] 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.

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

[9] P. W. Glynn. Diffusion Approximations. *Handbooks on OR and MS,* Vol. 2, D. P. Heyman and M. J. Sobel, Eds., Elsevier Science Publishers, 1990.

[10] J. M. Harrison. Heavy Traffic Analysis of a System with Parallel Servers: Asymptotic Optimality of Discrete Review Policies. *Annals of Applied Probability*, Vol. 8, (1998), pp. 822-848.

[11] J. M. Harrison. Brownian Models of Open Processing Networks: Canonical Representation of Workload. *Annals of Applied Probability*, Vol. 10, (2000), pp. 75-103.

[12] J. M. Harrison, M. J. Lopez. Heavy Traffic Resource Pooling in Parallel-Server Systems. *Queueing Systems*, Vol. 33, (1999), pp. 339-368.

[13] J. M. Harrison, J.A.Van Mieghem. Dynamic Control of Brownian Networks: State Space Collapse and Equivalent Workload Formulations. *Annals of Applied Probability*, Vol. 7, (1997), pp. 747-771.

[14] S. N. Ethier and T. G. Kurtz. *Markov Process: Characterization and Convergence*. John Wiley and Sons, New York, 1986.

[15] M. I. Reiman. Some Diffusion Approximations with State Space Collapse. In *Proc. of the Internat. Seminar on Modeling and Performance Evaluation Methodology*, Lecture Notes in Control and Information Sciences. Springer, New York, 1984, pp. 209-240.

[16] M. I. Reiman. A Multiclass Feedback Queue in Heavy Traffic. *Advances in Applied Probability*, Vol. 20, (1988), pp. 179-207.

[17] S. Shakkottai and A. L. Stolyar. Scheduling for Multiple Flows Sharing a Time-Varying Channel: The Exponential Rule. *American Mathematical Society Translations, Series 2, A volume in memory of F. Karpelevich, Yu. M. Suhov, Editor*, vol. 207, 2002, pp. 185-202. American Mathematical Society, Providence, RI, 2002.

[18] A. L. Stolyar. On the Stability of Multiclass Queueing Networks: A Relaxed Sufficient Condition via Limiting Fluid Processes. *Markov Processes and Related Fields*, 1(4), 1995, pp.491-512.

[19] 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.

[20] W. Whitt. Weak Convergence Theorems for Priority Queues: Preemptive Resume Discipline. *Journal of Applied Probability*, Vol. 8, (1971), pp. 74-94.

[21] R. J. Williams. Diffusion Approximations for Open Multiclass Queueing Networks: Sufficient Conditions Involving State Space Collapse. *Queueing Systems*, Vol. 30, (1998), pp. 27-88.

[22] R. J. Williams. On Dynamic Scheduling of a Parallel Server System with Complete Resource Pooling. *Fields Institute Communucations*, (1998).