Optimality and Large Deviations of Queues under the pseudo-Log Rule Opportunistic Scheduling

Bilal Sadiq and Gustavo de Veciana
Dept. of Electrical and Computer Engineering,
The University of Texas at Austin

***Abstract*— We consider a wireless node shared by multiple user flows where the channel capacity available to each user varies randomly with time. A scheduling rule in this context selects which flow to serve based on the current channel state and user queues. This involves a tradeoff between *maximizing current service rate* (being opportunistic) versus *balancing unequal queues* (enhancing user-diversity to enable future high capacity opportunities). We propose a throughput-optimal scheduling rule, called the pseudo-Log (p-Log) rule, and show that in the case of two users, it maximizes the asymptotic exponential decay rate of the sum-queue distribution. The proof relies on the *radial sum-rate monotonicity* (RSM) property satisfied by the p-Log rule, whereby as the queues scale up linearly, the scheduling rule de-emphasizes queue-balancing in favor of greedily maximizing the service rate. It also relies on refined sample path large deviation principle recently introduced by Stolyar to study such non-homogenous schedulers.**

In a companion paper we demonstrate via further analysis and simulations other virtues of RSM opportunistic schedulers (in particular the Log rule) in terms of minimizing overall mean delay, robustness to uncertainty in the traffic and channel statistics etc. The p-Log rule is a slight modification of the Log rule, for the sake of analytical convenience.

I. Introduction

We consider a two-queue-single-server system, where the queues are fed by exogenous arrivals and the server– having asynchronously time-varying capacity for each queue– can be dynamically allocated to one queue or the other. This models a wireless channel shared by heterogeneous users. The asynchronously time-varying nature of server capacity for each queue provides an opportunity to exploit favorable server states, e.g., by scheduling the queue that currently has a higher service rate – this is referred to as opportunistic scheduling. An opportunistic or channel-aware scheduler, however, may not even be stable, (i.e., keep the queues bounded) unless it is chosen carefully, e.g., using prior knowledge of mean arrival rates. Except in some degenerate cases, in order to ensure stability (for all possibly stabilizable arrival processes), an opportunistic scheduler must be both channel- and queue-aware, and must tradeoff *maximizing current service rate* versus *balancing unequal queues*. Note that by balancing queues, one can enhance subsequent user diversity, i.e., ensure more queues are non-empty, so as to achieve higher service rates in the future. Moreover, the *performance* of any scheduler depends on how this tradeoff is made. The queue-and-channel-aware schedulers that can achieve stability without knowledge of arrival or channel statistics, if stability is at all feasible, are called throughput-optimal. Examples are MaxWeight [1], Exp rule [2], and Log rule [3].

Stability, however, is a weak form of optimality, and it is of interest to study schedulers that are *delay-optimal*, e.g., schedulers that minimize the overall average delay (per data unit) seen by the users; or ones which minimize the probability that either the sum-queue or the largest queue overflows a large buffer. These schedulers are harder to characterize for servers with time-varying capacity, but some results are available that we briefly discuss next.

In [4] and [5] the Longest-Connected-Queue (LCQ) and Longest-Queue-Highest-Possible-Rate (LCQHPR) scheduling policies are introduced. Strong results are shown for these policies; they stochastically minimize the max and sum queue process, and thus also the max and sum queue tails and mean delay. However, in addition to assuming certain symmetry conditions on arrival and channel statistics, [4] is limited to on-off server capacities where only a single queue can be scheduled per time slot, and [5] assumes that the scheduler can allocate service rates from the information theoretic multiuser capacity region associated with the current channel state. In both cases, the above-mentioned tradeoff between queue balancing and throughput maximization is absent. Indeed in [4], all policies that pick a connected queue result in the same service rate, whereas, in case of [5], all policies that pick a service vector from the maximal points of the current capacity region, i.e. points on the max-sum-rate face, result in the same overall service rate. Thus one can achieve the *queue balancing* goal, without ever compromising *service rate*. Not surprisingly, in both cases the optimal policy turns out to be greedy, in that it allocates as much service rate as possible to the longest/longer queues.

A related server allocation problem is studied in [6]. The paper considers minimizing mean delay in a two queue system where each queue has a dedicated server and a third server can be dynamically shared between them. Without the underlying symmetry assumptions of [4] and [5], and using a dynamic programming approach, the existence of a monotone increasing switching curve on the state space of queue process is shown; (recall that for a two queue system, the switching curve under LCQ and LCQHPR policies lies along the line where both queues are equal.)

Finally, relaxing the symmetry assumptions of [4] and [5], [7] and [8] consider the asymptotic probability of max-queue overflow. The server capacity in [7], though time-varying, is identical for all users at any given time, thus the contention

This research was supported in part by grants AFOSR FA9550-07-1-0428 and NSF CNS-0721532.

between queue-balancing and service rate-maximization is again absent. In fact, the sum-queue process in [7] is identical for all work conserving schedulers. However [8] considers a server with *asynchronously* time-varying capacity across users and shows that the Exp rule minimizes the asymptotic probability of max-queue overflow. Indeed the model in [8] accurately captures a wireless channel shared by heterogenous users, and exhibits the tradeoff between queue-balancing and rate-maximization. Existence of this tradeoff also implies that, unlike the LCQ and LCQHPR policies, the asymptotic optimality of Exp rule does not translate to minimizing asymptotic probability of sum-queue overflow or the mean delay. In order to minimize the asymptotic probability of max-queue overflow, the desirable mode of overflow is the one where all queues (or, more precise, the set of overflowing queues, which then exclusively share the server,) grow at the same rate and overflow at the same time. This constraints the system throughput, while, of course, aggressively balancing the queues across users. In this paper we show that a radial sum-rate monotone scheduler (see [3] or Section III for definitions of radial sum-rate monotonicity), called the pseudo-Log rule, minimizes the asymptotic probability of sum-queue overflow. As we shall see, in order to minimize the asymptotic probability of sum-queue overflow, the desirable mode of overflow is the one where the system throughput is the highest possible and queues may build up at different rates. Although our focus is on overflows of the sum-queue (instead of overflows of the max-queue as in [8]), the general proof technique in [8] lends itself well to our problem and we rely heavily on the results developed therein.

Organization: The rest of the chapter is organized as follows. The system model is described in Section II. Queue-and-channel aware schedulers of interest and the property of radial sum-rate monotonicity are reviewed in Section III, followed by the introduction of pseudo-Log scheduling rule in Section IV. The three-part main result of the paper is summarized in Section V. Some preliminary discussion and relevant large deviation principles follow in Section VI. The proofs for the lower and the upper bounds stated in the main result of the paper are discussed in Section VII and VIII respectively. After defining local fluid sample paths and developing essential results in Section IX, the optimality of the p-Log rule, i.e., the last part of main result is proved in Section X.

II. System model

Consider the following problem of dynamically allocating a time-varying server to two queues. For each $i \in I = \{1,2\}$, queue i is fed by an independent arrival process $(\boldsymbol{A}_i(t),\ t = 0,1,\cdots)$ that is i.i.d over t, where $\boldsymbol{A}_i(t) \in \mathbb{Z}_+$ denotes the number of packets arriving in (the beginning of) time slot $[t, t+1)$. We assume that the arrivals are bounded, i.e., $\boldsymbol{A}_i(\cdot) \le C$ for some finite $C > 0$. Let $\boldsymbol{A}(t) = (\boldsymbol{A}_i(t),\ i \in I)$, and $\overline{\lambda} = E[\boldsymbol{A}(1)]$. We use bold face, e.g. $(\boldsymbol{A}(t),\ t = 0,1,\cdots)$, to mean the random process and plain font, e.g. $(A(t),\ t = 0,1,\cdots)$, to mean a realization of the process.

The time-varying state of the server is given by an i.i.d random process $(\boldsymbol{m}(t),\ t = 0,1,\cdots)$, where $\boldsymbol{m}(\cdot) \in \mathcal{M} = \{1,2,\cdots,M\}$ for some finite $M > 0$ denotes the state of the server over $[t, t+1)$, and is drawn from distribution $\pi = (\pi_1,\cdots,\pi_M) > 0$. Associated with each server state $m \in \mathcal{M}$ is an offered service vector $\mu^m \in \mathbb{Z}_+^2$. When in state m over a time slot, the server can either serve at most μ_1^m number of packets from queue 1, or it can serve at most μ_2^m number of packets from queue 2. The scheduling problem then is to allocate the server to a queue $i^*(t) \in I$ for each time slot $[t, t+1)$ such that a given optimality criterion is met.

At any integer t, let the (random) vector $\boldsymbol{Q}(t) = \big(\boldsymbol{Q}_i(t),\ i \in I\big) \in \mathbb{Z}_+^2$, where $\boldsymbol{Q}_i(t)$ denotes the number of packets in the i^{th} queue at the end of time slot $[t-1, t)$. Then,

$$\boldsymbol{Q}_i(t+1) = \big(\boldsymbol{Q}_i(t) + \boldsymbol{A}_i(t) - \mu_i^{\boldsymbol{m}(t)} \mathbb{1}_{\{i^*(t)=i\}}\big)^+ .$$

To be precise, the server state sample path $(m(\tau),\ \tau \le t)$, the queue process sample path $(Q(\tau),\ \tau \le t)$, and the arrival process sample path $(A(\tau),\ \tau < t)$ are available to the scheduler prior to choosing $i^*(t)$. Revealing $A(t)$ too has no bearing on the results presented here. The function i^*, viewed as a mapping from the set of system sample paths to the set of users, is called a *scheduler* or *scheduling policy*. It is easy to see that under a static-state feedback scheduler, i.e. one where $i^*(t) \equiv i^*\big(Q(t), m(t)\big)$, process $(\boldsymbol{Q}(t),\ t = 0,1,\cdots)$ forms a discrete time Markov chain on $\mathbb{Z}_+^2$.

We extend the domain of all discrete time processes and functions to continuous time: a function (originally defined on integer times) has the same value at any real t that it takes at $\lfloor t \rfloor$. Then all processes and functions defined above dwell on the space of *real-valued right continuous functions with left limits*, here denoted by $\mathcal{D}$. We assume that $\mathcal{D}$ is endowed with the topology of uniform convergence over compact sets (u.o.c), and the k-times product space $\mathcal{D}^k$ with the product topology. Lastly, let $(\Omega, \mathcal{F}, P)$ be the probability space that is large enough to define all the random processes in this paper.

Capacity region

For each server state $m \in \mathcal{M}$, let V^m denote the closed *triangle* having vertices $(0,0)$, $(0,\mu_2^m)$, and $(\mu_1^m, 0)$. Then the expected service rates jointly offered to the two queues under any scheduling rule (such that the expectation exists), conditional on the server being in state m, lies in the triangle V^m. Define the capacity region V_π as the set of expected service rate vectors offered to the two queues under all possible scheduling rules, then V_π is a convex polyhedron given by weighted Minkowski sum of regions V^m, i.e.,

$$\begin{aligned} V_\pi &= \pi_1 V^1 \oplus \cdots \oplus \pi_M V^M , \\ &= \left\{ \sum_{m \in \mathcal{M}} \pi_m v(m) :\ v(m) \in V^m,\ m \in \mathcal{M} \right\}. \end{aligned} \quad (1)$$

(See Fig. 1 for a graphical illustration of capacity region: the server has $M = 5$ states with some

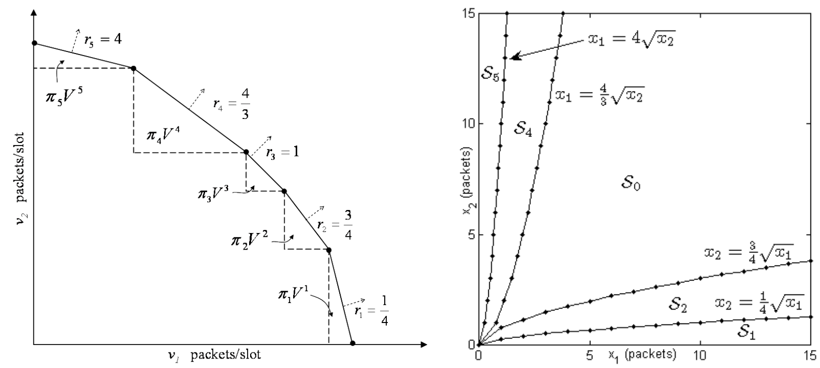

Fig. 1. (Left) Capacity region for $\mu^m \in \{(1,4),(3,4),(1,1),\ (4,3),(4,1)\}$, depicting Minkowski addition and outer-normal vectors; (right) resulting partitions under p-Log.

distribution π and offered service vectors μ^m in set $\{(1,4),(3,4),(1,1),(4,3),(4,1)\}$.)

We assume there exists a $v \in V_\pi$ such that $\overline{\lambda} < v$, which is a sufficient condition for stabilizability of queues [1].

Note that regions V^m are triangles because the server can only be allocated to a single queue over a time slot (or equivalently, the server can only be time-shared.) However, we can relax this to allow that regions V^m be arbitrary convex polyhedrons. For example, V^m can be information theoretic polymatroids as in [5]. The optimality results presented in this paper fall through in spite of this relaxation.

III. Throughput-optimal schedulers and radial sum-rate monotonicity

The throughput-optimal schedulers MaxWeight, Exp rule, and Log rule, mentioned in the introduction, are all static state-feedback. These schedulers can be defined as follows: when the system is in state $(\boldsymbol{Q}(t),\boldsymbol{m}(t)) = (Q,m)$, schedulers MaxWeight, Exp rule, and Log rule serve a user i^*_{MW}, i^*_{EXP}, and i^*_{LOG} respectively that is given by,

$$i^*_{MW}(Q,m) \in \arg\max_{i\in I}\ \{\ b_i Q_i^\alpha \times \mu_i^m\ \}\ ,$$

$$i^*_{EXP}(Q,m) \in \arg\max_{i\in I}\{b_i \exp\left(\frac{a_i Q_i}{c+(\sum_{j\in I} Q_j)^\eta}\right)\times \mu_i^m\},$$

$$i^*_{LOG}(Q,m) \in \arg\max_{i\in I}\ \{\ b_i \log\left(1+a_i Q_i\right) \times \mu_i^m\ \}\ ,$$

for any fixed positive b_i's, a_i's, α, c, and $0<\eta<1$, and augmented with any fixed tie-breaking rule.

Indeed numerous throughput-optimal schedulers can be engineered that *react* differently to the disparity among the users' queue lengths (i.e., make different tradeoffs between *service rate maximization* and *queue balancing*), see, e.g., [9] or Theorem 1 of [3]. A scheduler is called *radial sum-rate monotone* if, as the queues scale up linearly, the scheduling rule allocates server in a manner that de-emphasizes queue-balancing in favor of greedily maximizing the current service rate. More formally, let $v(Q) \in V_\pi$ be the vector of expected service offered to the queues under a static state-feedback scheduler i^*, conditional on queue state being Q, i.e.,

$$v(Q) = \left(\ E[\mu_i^{\boldsymbol{m}} \mathbb{1}_{\{i^*(Q,\boldsymbol{m})=i\}}]:\ i\in I\ \right)\ ,$$

where expectation is w.r.t (random) $\boldsymbol{m}$ drawn from distribution π. Also, for any vector x, let $|x|$ denote the sum of its components, (and not the usual L_1 norm). We say the scheduling policy i^* is *radially sum-rate monotone* if for any Q and scalar θ such that $\theta Q \in \mathbb{Z}_+^n$, the total expected offered service, $|v(\theta Q)|$, is an increasing function of θ, and $\lim_{\theta\to\infty}|v(\theta Q)| = \max(|y| : y\in V_\pi,\ y_i = 0 \text{ if } Q_i = 0)$.

The Log rule described above and the pseudo-Log rule introduced next are both radial sum-rate monotone, whereas, the MaxWeight and the Exp rule are not (see [3]). In fact, Exp rule is the opposite of radial sum-rate monotone, in that, as the queues grow linearly, the Exp rule emphasizes queue-balancing while compromising the current service rate.

IV. The pseudo-Log scheduling rule

We'll define pseudo-Log (p-Log) scheduling rule using a vector field $h=(h_1,h_2)$ on $\mathbb{Z}_+^2$, the state space of queue process: when queue is in state $Q\in\mathbb{Z}_+^2$ and the server in state $m\in\mathcal{M}$, then the server is allocated to queue $i^*_{pLog}(Q,m)$ given by,

$$i^*_{pLog}(Q,m) \in \arg\max_{i\in N} h_i(Q)\mu_i^m\ , \qquad (2)$$

where, in case of tie, if $Q_1 \geq Q_2$ the channel is allocated to queue 1, otherwise to queue 2. Since it is only the slope, $\frac{h_2(Q)}{h_1(Q)}$, that determines the scheduling decision, we'll specify vector field using slope.

Consider the region V^m, let r_m denote the slope of the outer normal to the line segment joining $(0,\mu_2^m)$, and $(\mu_1^m,0)$, i.e., $r_m = \frac{\mu_1^m}{\mu_2^m}$, (where $r_m=\infty$ if $\mu_2^m=0$.) Let's assume that channel states are sorted in ascending order of outer normal slopes r_m, i.e., $r_1\leq r_2\leq\ldots\leq r_M$. Also, let $r_0=0$ and $r_{M+1}=\infty$. Let r_k be the largest slope strictly less than 1 and r_l the smallest slope strictly greater than 1.

Next, we partition $\mathbb{Z}_+^2$ into at most $M+1$ regions and assign a constant slope to vector field h in each region. For each $m\in\{1,\cdots,k\}$, define partition,

$$\mathcal{S}_m = \{(x_1,x_2)\in\mathbb{Z}_+^2 |\ r_{m-1}\sqrt{x_1}\leq x_2 < r_m\sqrt{x_1}\}\ , \quad (3)$$

and assign $h(x)$ a slope of $\frac{r_{m-1}+r_m}{2}$ in partition $\mathcal{S}_m$ (see Fig. 1). It will be useful later to note that assigning $h(x)$ any slope strictly between r_{m-1} and r_m in partition $\mathcal{S}_m$ will yield the same scheduling decision as assigning $h(x)$

a slope of $\frac{r_{m-1}+r_m}{2}$. The above assignment of partitions and vector field slopes has following interpretation (see (2)): when queue lies in any partition $\mathcal{S}_m$ defined so far, the server in states $\{m, \cdots, M\}$ is assigned to queue 1 and in all remaining states to queue 2. Also, for each $m \in \{l, \cdots, M\}$, define a partition $\mathcal{S}_m$,

$$\mathcal{S}_m = \{(x_1, x_2) \in \mathbb{Z}_+^2 |\ r_{m+1}^{-1}\sqrt{x_2} \le x_1 < r_m^{-1}\sqrt{x_2}\} \ ,$$

and assign $h(x)$ a slope of $\left(\frac{r_{m+1}^{-1}+r_m^{-1}}{2}\right)^{-1}$ in partition $\mathcal{S}_m$. Then by (2), when queue lies in any partition $\mathcal{S}_m$ for $m \ge l$, the server in states $\{1, \cdots, m\}$ is assigned to queue 2 and in all remaining states to queue 1. There are points in $\mathbb{Z}_+^2$ which have still not been assigned to any partition; we associate them all with partition $\mathcal{S}_0$, i.e.,

$$\mathcal{S}_0 = \{(x_1, x_2) \in \mathbb{Z}_+^2 |\ r_k\sqrt{x_1} \le x_2 \le (r_l x_1)^2\} \ ,$$

and assign $h(x)$ a slope of 1 in partition $\mathcal{S}_0$. We'll refer to the horn-shaped partition $\mathcal{S}_0$ as max-sum rate partition.

Remark 1: Perhaps at the cost of clarity and expressiveness, a more compact definition of vector field h yielding the same scheduling decision in every system state as the definition above can be given as follows. Fix any $x \in \mathbb{Z}_+^2$, let

$$i = \begin{cases} 1, & \text{if } x_1 \ge x_2; \\ 2, & \text{otherwise,} \end{cases}$$

and let j denote the component not denoted by i. Define the two components $h_1(x)$ and $h_2(x)$ as follows,

$$\begin{aligned} h_i(x) &= \sqrt{x_i} \ , \\ h_j(x) &= \min(x_j, \sqrt{x_i}) \ . \end{aligned}$$

Now p-Log rule can be defined using this definition of vector field, *augmented with an appropriate tie-breaking rule*. We find the original definition of vector field h more useful for exposition of arguments in the this paper.

Remark 2: Consider any non-empty partition $\mathcal{S}_m$, and pick an $x \in \mathcal{S}_m$. Let $v(x) \in V_\pi$ be the vector of expected service offered to the two queues under p-Log rule, conditional on queue state being x, i.e.,

$$v(x) = \Big(\ E[\mu_i^{\boldsymbol{m}} \mathbb{1}_{\{i^*_{pLog}(x,\boldsymbol{m})=i\}}] : \ i \in I \ \Big) \ ,$$

where expectation is w.r.t (random) $\boldsymbol{m}$ drawn from distribution π. Then, by (2) and Lemma 1 of [3],

$$v(x) \in \arg\max_{v \in V_\pi} \langle v, h(x) \rangle \ ,$$

where the argmax is necessarily unique for all $x \notin \mathcal{S}_0$. Hence, in each non-empty partition $\mathcal{S}_m$ for $1 \le m \le M$, vector $v(x)$ is equal to a unique vertex of V_π that has $h(x)$ as an outer normal vector. Moreover, in partition $\mathcal{S}_0$, if the argmax above is not unique, then $v(x)$ depends on the sign of $x_1 - x_2$: if $x_1 - x_2 \ge 0$, then $v(x)$ is such that $v_1(x)$ is the largest among all vertices achieving argmax in above; otherwise, $v(x)$ is such that $v_2(x)$ is the largest.

Now one can see that the similarity between the pseudo-Log rule defined above and the Log rule defined in [3]. In p-Log rule, switching curves are a constant times $\sqrt{\max(x_1, x_2)}$, where the constant varies from one switching curve to the next; whereas, in the Log rule, switching curves are $(1 + \max(x_1, x_2))^\eta - 1$ for some $\eta \in [0, 1]$, where η varies from one switching curve to the next. For comparison see (3) and [3].

V. Main result

The following three-part theorem– stating a lower bound on the tail of overflow probability, an upper bound on the same under p-Log rule, and the optimailty of p-Log rule– summarizes the main results of this paper. The first part is discussed in Section VII, the second in Section VIII (and Appendix), while the last in Section X (after developing the necessary results in Section IX).

Theorem 1: There exists finite $T_0 > 0$ such that for any scheduling rule and any $t > T_0$, we have the following lower bound:

$$\liminf_{n\to\infty} \frac{1}{n} \log P\big(|\boldsymbol{Q}(nt)| > n\big) \ge -J_* \ .$$

where J_* is defined in Section VII.

(ii) Now, consider the system under p-Log rule, then process $(\boldsymbol{Q}(t),\ t = 0, 1, \cdots)$ forms an ergodic Markov chain, and we have the following upper bound under the stationary distribution of Markov chain $\boldsymbol{Q}$,

$$\limsup_{n\to\infty} \frac{1}{n} \log P\big(|\boldsymbol{Q}(0)| > n\big) \le -J_{**} \ .$$

where J_{**} is defined in Section VIII.

(iii) The p-Log rule minimizes the tail of the probability of sum-queue overflow, i.e.,

$$J_{**} = J_* \ .$$

VI. Fluid-scaled processes and large deviation principles

In this section, we define sequences of fluid-scaled processes and functions, and give the Large Deviation Principle [10] on those sequences, as needed for proving Theorem 1. Define arrival flow $\boldsymbol{F} = \big(\boldsymbol{F}(t) = (\boldsymbol{F}_i(t),\ i \in I), t \ge 0\big)$ constructed from process $(\boldsymbol{A}(t),\ t \ge 0)$ like this,

$$\boldsymbol{F}_i(t) = \sum_{k=0}^{\lfloor t-1 \rfloor} \boldsymbol{A}_i(k) \ ,$$

and flow $\boldsymbol{G} = \big(\boldsymbol{G}(t) = (\boldsymbol{G}_m(t),\ m \in \mathcal{M}),\ t \ge 0\big)$ constructed from process $(\boldsymbol{m}(t),\ t \ge 0)$ like this,

$$\boldsymbol{G}_m(t) = \sum_{k=0}^{\lfloor t-1 \rfloor} \mathbb{1}_{\{\boldsymbol{m}(k)=m\}} \ .$$

The set (Q, F, G), where $Q = (Q(t),\ t \ge 0)$ is the queue sample path (under a fixed scheduling rule) corresponding to the sample paths (F, G) and initial state $Q(0)$, denotes a realization of the system $(\boldsymbol{Q}, \boldsymbol{F}, \boldsymbol{G})$. For each $n = 0, 1, \cdots$ consider an independent and stochastically equivalent system $(\boldsymbol{Q}^{(n)}, \boldsymbol{F}^{(n)}, \boldsymbol{G}^{(n)})$. The corresponding sequence of fluid-scaled processes, denoted by $(\boldsymbol{q}^{(n)}, \boldsymbol{f}^{(n)}, \boldsymbol{g}^{(n)})$, is defined as,

$$\boldsymbol{q}^{(n)} = (\boldsymbol{q}^{(n)}(t),\ t \geq 0) = \left(\frac{1}{n}\boldsymbol{Q}^{(n)}(nt),\ t \geq 0\right) \ ,$$

with $\boldsymbol{f}^{(n)}$ and $\boldsymbol{g}^{(n)}$ defined similarly.

For each $i \in I$, define for any $\lambda \geq 0$ (the rate function of sequence $\boldsymbol{f}_i^{(n)}(1)$),

$$L_i(\lambda) = \sup_{\theta \geq 0} \Big(\ \theta\lambda - \log E\left[\ e^{\theta A_i(1)}\ \right]\ \Big) \ ,$$

with $L_i(\cdot) = \infty$ over $(-\infty, 0)$. Also, for any vector $\lambda \in \mathbb{R}_+^2$, let,

$$L_{(f)} = \sum_{i \in I} L_i(\lambda_i) \ .$$

Define for any probability distribution $\gamma = (\gamma_m,\ m \in \mathcal{M})$ (the relative entropy of γ w.r.t distribution π),

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

with $L_{(g)}(\cdot) = \infty$ everywhere outside the simplex on $\mathbb{R}_+^M$. Consider any functions $(f, g) \in \mathcal{D}^{2+M}$. For any $t > 0$, if $(f(0), g(0)) = 0$ and (f, g) are absolutely continuous in interval $[0, t]$, then let,

$$J_t(f, g) = \int_0^t L_{(f)}\big(f'(s)\big) + L_{(g)}\big(g'(s)\big)\ ds \ ,$$

otherwise $J_t(f, g) = \infty$. Function $J_t(f, g)$ is called the cost of trajectories (f, g) over time interval $[0, t]$. The following is a form of Mogulsky theorem [10].

Proposition 1: For any fixed $T > 0$, consider a sequence in n of the fluid-scaled processes $(\boldsymbol{f}^{(n)}, \boldsymbol{g}^{(n)}) = \big((\boldsymbol{f}^{(n)}(t), \boldsymbol{g}^{(n)}(t)),\ t \in [0, T]\big)$. Then for any measurable $B \subseteq \mathcal{D}^{2+M}$, we have,

$$\begin{aligned} &-\inf\{J_T(f,g) | (f,g) \in B^\circ\} \\ &\quad \leq \liminf_{n\to\infty} \frac{1}{n} \log P\big((\boldsymbol{f}^{(n)}, \boldsymbol{g}^{(n)}) \in B\big) \\ &\quad \leq \limsup_{n\to\infty} \frac{1}{n} \log P\big((\boldsymbol{f}^{(n)}, \boldsymbol{g}^{(n)}) \in B\big) \\ &\quad \leq -\inf\{J_T(f,g) | (f,g) \in \overline{B}\} \ , \end{aligned}$$

where, B° and $\overline{B}$ denote the interior and closure of set B respectively.

Let $u(n) = \lceil n^\alpha\ \rceil$ for some fixed $\alpha \in (0, 0.5)$. For any function $h \in \mathcal{D}^{2+M}$, let $U^n h$ denote the piece-wise linear function obtained by linear interpolation over samples $\big(h(\frac{ku(n)}{n}),\ k = 0, 1, \cdots\big)$. The following upper bound was introduced in [8] as refined Mogulsky theorem.

Proposition 2: For any fixed $T > 0$, consider a sequence in n of the fluid-scaled processes $(\boldsymbol{f}^{(n)}, \boldsymbol{g}^{(n)}) = \big((\boldsymbol{f}^{(n)}(t), \boldsymbol{g}^{(n)}(t)),\ t \in [0, T]\big)$. Suppose, for each n there is a fixed measurable $B^{(n)} \subseteq \mathcal{D}^{2+M}$, that is a subset of the set of feasible realizations of $(\boldsymbol{f}^{(n)}, \boldsymbol{g}^{(n)})$ in $[0, T]$. Then,

$$\begin{aligned} &\limsup_{n\to\infty} \frac{1}{n} \log P\big((\boldsymbol{f}^{(n)}, \boldsymbol{g}^{(n)}) \in B^{(n)}\big) \\ &\quad \leq -\liminf_{n\to\infty} \inf\{J_{T^{(n)}} U^n(f,g) | (f,g) \in B^{(n)}\} \ , \end{aligned}$$

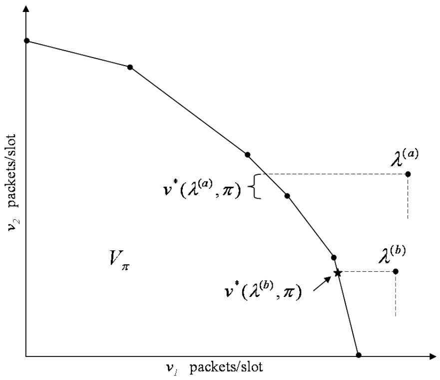

Fig. 2. The *optimal* capacity vector $v^*(\lambda^{(b)}, \pi)$ is unique, whereas, vector $v^*(\lambda^{(a)}, \pi)$ is any point on the annotated part of max-sum-rate face of region V_π.

where $T^{(n)} = \frac{u(n)}{n}\lfloor\frac{nT}{u(n)}\rfloor$.

Also introduced in [8] is a notion of generalized fluid sample path (GFSP) which we describe next. Consider a sequence in n of realization $(q^{(n)}, f^{(n)}, g^{(n)})$ such that along some subsequence (still denoted by $\{n\}$), we have the u.o.c convergence

$$(q^{(n)}, f^{(n)}, g^{(n)}) \rightarrow (q, f, g)$$

for some Lipschitz continuous functions (q, f, g), and the u.o.c convergence

$$\begin{aligned} \overline{J}^{(n)} &= \Big(\overline{J}_t^{(n)},\ t \geq 0\Big) \ , \\ &= \Big(J_{t^{(n)}} U^n(f^{(n)}, g^{(n)}),\ t \geq 0\Big) \rightarrow \overline{J} = \big(\overline{J}_t, t \geq 0\big) \end{aligned}$$

for some non-negative increasing Lipschitz continuous function $\overline{J}$. Then following construction is called a GFSP,

$$\psi = \Big[(q^{(n)}, f^{(n)}, g^{(n)}),\ \overline{J}^{(n)},\ n = 0, 1, \cdots;\ (q, f, g);\ \overline{J}\Big] .$$

The function $\overline{J}$ is called the refined cost function of GFSP.

VII. Lower Bound on Overflow Probability under any Scheduling Rule

For any distribution γ on the set $\mathcal{M}$ of server states, let V_γ be the corresponding capacity region (see (1)). For any vector $\lambda \in \mathbb{R}_+^2$, define a capacity vector $v^*(\lambda, \gamma)$ given by,

$$v^*(\lambda, \gamma) \in \arg\max\{|v| :\ v \leq \lambda,\ v \in V_\gamma\} \ . \tag{4}$$

For an example, see Fig. 2 depicting $v^*(\lambda, \pi)$ for two hypothetical vectors $\lambda^{(a)}$ and $\lambda^{(b)}$ lying outside capacity region V_γ. The interpretation is, if the arrival process were to have an empirical mean λ and the server state to have a "twisted" distribution γ, then serving queues according to the capacity vector $v^*(\lambda, \gamma)$ minimizes the rate of sum-queue build-up, $|\lambda - v^*(\lambda, \gamma)|$.

Finally, define the minimum cost (per unit increase in sum-queue) J_*,

$$J_* = \inf_{\gamma, \lambda} \frac{L_{(g)}(\gamma) + L_{(f)}(\lambda)}{|\lambda - v^*(\lambda, \gamma)|} \ . \tag{5}$$

Sketch of Proof of Theorem 1-(i): Let (λ^*, γ^*) be a point that achieves the minimum in (5), and let $T_0 = (|\lambda^* - v^*(\lambda^*, \gamma^*)|)^{-1}$. We show that regardless of the scheduling rule, all realizations (f, g) sufficiently *close* to $(\lambda^* t, \gamma^* t)$ over interval $[0, T_0]$– thus having cost $J_{T_0}(f, g)$ close to J_*– lead to overflow at time T_0. See [11] for detailed proof.

VIII. Upper bound on overflow probability under the p-Log rule

Let J_{**} denote the lowest refined cost of a GFSP that raises $|q(t)|$ to 1 from the initial state $q(0) = 0$, i.e.,

$$J_{**} = \inf_{t \geq 0} J_{**,t} \ , \tag{6}$$

where,

$$J_{**,t} = \inf\{\overline{J}_t|\psi : \ q(0) = 0, \ |q(t)| \geq 1\} \ .$$

The following is a restatement of Theorem 1-(ii) in terms of a sequence of fluid scaled queues.

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

$$\limsup_{n \to \infty} \frac{1}{n} \log P\big(|\boldsymbol{q}^{(n)}(0)| > 1\big) \leq -J_{**} \ .$$

Remark 3: This is the *equivalent* of Theorem 8.4 of [8], and its proof too follows the same framework and uses classical Wentzel-Freidlin constructions [Fre98]. The theorem establishes two things: firstly, that the upper bound on the probability of overflow when starting with empty queues, given by Stolyar's refinement to Mogulsky's upper bound, indeed reduces to inf over the cost of GFSPs of interest; and secondly, that a GFSP with the cheapest limiting trajectories (f, g) that can raise the sum queue $|q|$ to 1, *starting with empty queues*, indeed has a cost arbitrarily close to the cost of the cheapest trajectory *starting in the stationary regime*. See [11] for proof.

It is clear that $J_{**} \leq J_*$ (since $-J_*$ was lower bound under any scheduling rule.) To prove the optimality of p-Log rule, we need show $J_{**} = J_*$. In the following section, we gradually develop the results needed for this proof.

IX. Local fluid sample path

Let us first motivate the need for a LFSP (local fluid sample path.) For each n, define the set $\mathcal{S}_m^{(n)}$ as *fluid* scaling of partition $\mathcal{S}_m$, i.e.,

$$\mathcal{S}_m^{(n)} = \{x \in \mathbb{R}_+^2 : \ nx \in \mathcal{S}_m\} \ .$$

Then for the n^{th} system, the scheduling decision at time t depends on which set $\mathcal{S}_m^{(n)}$ the queue $q^{(n)}(t)$ lies in. As $n \to \infty$, (the characteristic function of) set $\mathcal{S}_0^{(n)}$ converges pointwise to (the characteristic function of) $\mathcal{S}_0^{(\infty)} = \{x \in \mathbb{R}_+^2 : \ x > 0\}$; while all other scaled partitions except $\mathcal{S}_0^{(n)}$ collapse to one of the axes. Note that this is true for any radial sum-rate monotone scheduling rule w.r.t. weight vector $(1, 1)$. For the limiting trajectory $q(t)$, while we can still show that,

$$\text{if} \ \ q(t) \in \mathcal{S}_0^{(\infty)}, \quad \text{then} \quad \frac{d}{dt}|q(t)| = \frac{d}{dt}|f(t)| - \max_{v \in V_{\gamma(t)}} |v|,$$

we lose information about service rate when $q(t)$ hits an axis. Hence, we define a local FSP, that has a *finer* than fluid scaling, and under which partitions do not collapse.

Consider a GFSP over some interval $[0, T]$ and fix any $\tau \in (0, T)$ such that $q(\tau) \neq 0$. Any sequence $\tau^{(n)} \to \tau$ has a subsequence along which $q^{(n)}(\tau^{(n)}) \to q(\tau)$. Set $\sigma_n = \sqrt{q_*^{(n)}(\tau^{(n)})}/\sqrt{n}$, where, $q_*^{(n)}(\cdot) = \max_{i \in I} q_i^{(n)}(\cdot)$. Following re-scaled functions over interval $[\tau^{(n)}, \tau^{(n)} + \sigma_n S]$, parameterized by $s \in [0, S]$, are called the local fluid sample paths:

$$\begin{aligned}
{}_\diamond q_i^{(n)}(s) &= \frac{1}{\sigma_n}\Big(q_i^{(n)}\big(\tau^{(n)} + \sigma_n s\big) - q_i^{(n)}\big(\tau^{(n)}\big)\Big) \ , \\
{}_\diamond \hat{q}_i^{(n)}(s) &= \frac{1}{\sigma_n} \ q_i^{(n)}\big(\tau^{(n)} + \sigma_n s\big) \ , \\
{}_\diamond d^{(n)}(s) &= {}_\diamond \hat{q}_1^{(n)}(s) - {}_\diamond \hat{q}_2^{(n)}(s) \ , \\
{}_\diamond f_i^{(n)}(s) &= \frac{1}{\sigma_n}\Big(f_i^{(n)}\big(\tau^{(n)} + \sigma_n s\big) - f_i^{(n)}\big(\tau^{(n)}\big)\Big) \ , \\
{}_\diamond g_m^{(n)}(s) &= \frac{1}{\sigma_n}\Big(g_m^{(n)}\big(\tau^{(n)} + \sigma_n s\big) - g_m^{(n)}\big(\tau^{(n)}\big)\Big) \ .
\end{aligned}$$

Then along some subsequence in n, functions $\big({}_\diamond q^{(n)}, \ {}_\diamond f^{(n)}, \ {}_\diamond g^{(n)}\big)$ converge uniformly over $[0, S]$ to Lipschitz continuous functions $\big({}_\diamond q, \ {}_\diamond f, \ {}_\diamond g\big)$, whereas, for each $i \in I$, ${}_\diamond \hat{q}_i^{(n)}$ either converges uniformly to a finite Lipschitz continuous function ${}_\diamond \hat{q}_i$, or is identically equal to ∞, and function ${}_\diamond d^{(n)}(s)$ converges uniformly to a finite Lipschitz continuous function ${}_\diamond d$, or is identically equal to $+\infty$ or $-\infty$. Since $q(\tau) \neq 0$, it must be that ${}_\diamond \hat{q}_i = \infty$ for at least one $i \in I$. Note that the local fluid queue, ${}_\diamond q$, is merely a *re-centered* version of ${}_\diamond \hat{q}$, i.e., ${}_\diamond \hat{q}(s) = {}_\diamond q(s) + {}_\diamond \hat{q}(0)$, and is always finite by virtue of this re-centering. Moreover, the trajectory ${}_\diamond q$ dwells in space $\{x \in \mathbb{R}^2 | x \geq -{}_\diamond \hat{q}(0)\}$, which is at least a half-plane. Lastly, we have the following relation between the *cost* of LFSP and the refined cost sequence of GFSP over $[\tau^{(n)}, \tau^{(n)} + \sigma_n S]$ (see (9.3) of [8],)

$$J_S({}_\diamond f, {}_\diamond g) - J_0({}_\diamond f, {}_\diamond g) \leq \liminf_{n \to \infty} \frac{1}{\sigma_n} \Big(\overline{J}^{(n)}_{\tau^{(n)} + \sigma_n} - \overline{J}^{(n)}_{\tau^{(n)}}\Big) \tag{7}$$

A. Scheduling over time scales of LFSP

Without loss of generality, suppose $q_1(\tau) \geq q_2(\tau)$. Then by (3), for n large enough, the scheduling decision in interval $[\tau^{(n)}, \tau^{(n)} + \sigma_n S]$ depends on the sign of ${}_\diamond d^{(n)}(s)$, (recall the tie-breaking rule mentioned in the description of p-Log rule in Section IV,) and then on the ratio,

$$\frac{Q_2\Big(n\tau^{(n)} + n\sigma_n s\Big)}{\sqrt{Q_1\Big(n\tau^{(n)} + n\sigma_n s\Big)}} = \frac{{}_\diamond \hat{q}_2^{(n)}(s)\sqrt{q_1^{(n)}(\tau^{(n)})}}{\sqrt{q_1^{(n)}(\tau^{(n)} + \sigma_n s)}} \ ,$$

which converges to ${}_\diamond \hat{q}_2(s) = {}_\diamond q_2(s) + {}_\diamond \hat{q}_2(0)$ uniformly in $[0, S]$. Hence, switching curves on the space of ${}_\diamond q$ are now

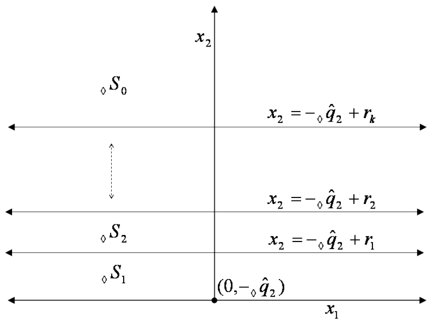

Fig. 3. Partitions and switching curves on space of local fluid queue $_\diamond q$.

parallel to $_\diamond q_1$ axis. More precisely, for each $1 \leq m \leq k$, define set $_\diamond\mathcal{S}_m$ by appropriately re-scaling the partition $\mathcal{S}_m$, i.e.,

$$_\diamond\mathcal{S}_m = \{(x_1, x_2) \in \mathbb{R}^2 |\ r_{m-1} \leq x_2 + {_\diamond\hat{q}_2}(0) < r_m\} \ .$$

Lastly, define,

$$_\diamond\mathcal{S}_0 = \{(x_1, x_2) \in \mathbb{R}^2 |\ r_k \leq x_2 + {_\diamond\hat{q}_2}(0)\} \ .$$

(See Fig. 3 for graphical illustration.) Then $\mu(s)$, service rate seen by the local fluid queue $_\diamond q(s)$ at time s, depends on which re-scaled partition $_\diamond\mathcal{S}_m$ contains $_\diamond q(s)$. We also define a vector field $_\diamond h$ over the state space of $_\diamond q$ like this: for each $m \leq k$, assign $_\diamond h$ the same slope in set $_\diamond\mathcal{S}_m$ as h has in $\mathcal{S}_m$. The following Lemma– essentially relating the service rate $\mu(s)$ to the vector field $_\diamond h({_\diamond q}(s))$– can be derived without much effort from the results on MaxWeight scheduler in [1], (these results relate the service rate under MaxWeight scheduler to the vector field defining the scheduler.)

Lemma 1: (See Lemma 9.1 of [8]) For any LFSP, for almost all $s \in [0, S]$, the following derivatives exist and are finite.

$$\begin{aligned} \lambda(s) &= \frac{d}{ds}\, {_\diamond f}(s) \ , \\ \gamma(s) &= \frac{d}{ds}\, {_\diamond g}(s) \ , \\ \frac{d}{ds}\, {_\diamond q}(s) &= \lambda(s) - \mu(s) \ , \end{aligned}$$

for some,

$$\mu(s) \in \arg\max_{v \in V_{\gamma(s)}} \langle {_\diamond h}({_\diamond q}(s)), v\rangle \ . \quad (8)$$

Remark 4: Under the definition of $_\diamond h$ (or originally h), argmax in (8) is necessarily unique unless $_\diamond q(s) \in {_\diamond\mathcal{S}_0}$. Even in case $_\diamond q(s) \in {_\diamond\mathcal{S}_0}$, the service rate $\mu(s)$ can still be uniquely identified by further considering the sign of $_\diamond d(s)$: if $_\diamond d(s) \geq 0$, $\mu(s)$ is such $\mu_1(s)$ is largest among all points achieving max in (8); similarly, if $_\diamond d(s) < 0$, $\mu(s)$ is such that $\mu_2(s)$ is the largest.

Points in $[0, S]$ where the above derivatives exist– and they exist a.e.– are called regular. The following lemma shows that any interesting GFSP contains a *similar* cost LFSP. It serves the same purpose for the problem at hand as the result in Section 11 of [8] does for the max-queue overflow problem. See [11] for proof.

Lemma 2: Suppose a GFSP ψ is given that satisfies $q(0) = 0$, $|q(T)| = 1$ for some $T > 0$ and has a cost $\overline{J}_T < \infty$. Then, for an arbitrarily small $\epsilon > 0$, an LFSP $\left({_\diamond q},\ {_\diamond\hat{q}},\ {_\diamond d},\ {_\diamond f},\ {_\diamond g}\right)$ over an arbitrarily large interval $[0, S]$ can be constructed from the elements of ψ, such that,

$$|{_\diamond q}(S)| - |{_\diamond q}(0)| \geq \theta S \ , \quad (9)$$

for some $\theta > 0$ (independent of ϵ), and cost (per unit increase in sum queue) of this LFSP is bounded above by $\overline{J}_T + \epsilon$, i.e.,

$$\frac{J_S({_\diamond f},\ {_\diamond g}) - J_0({_\diamond f},\ {_\diamond g})}{|{_\diamond q}(S)| - |{_\diamond q}(0)|} \leq \overline{J}_T + \epsilon \ . \quad (10)$$

X. Proof of Theorem 1-(iii): Optimality of the p-Log rule

Recall that to prove optimality of p-Log rule (Theorem 1-(iii)), we need show $J_{**} \geq J_*$. We'll do this by showing that assuming $J_{**} < J_*$ leads to a contradiction (with the definition of J_*.)

Suppose $J_{**} < J_*$, then by definition of J_{**} in (6), there exists a GFSP ψ satisfying $q(0) = 0$, $|q(T)| = 1$ for some finite $T > 0$, and having a cost $\overline{J}_T < J_*$. Then by Lemma 2, we can construct an LFSP $\left({_\diamond q},\ {_\diamond\hat{q}},\ {_\diamond d},\ {_\diamond f},\ {_\diamond g}\right)$ satisfying (9) and (10) for an ϵ small enough so that $J_{***} \equiv \overline{J}_T + \epsilon < J_*$.

Without loss of generality, suppose $_\diamond\hat{q}_1 = \infty$, (and all switching curves on the space of $_\diamond q$ are parallel to $_\diamond q_1$ axis.) Let S_1 and S_2 respectively be the first and the last time in $[0, S]$ such that the trajectory $_\diamond q_2(s) \leq -{_\diamond\hat{q}_2}(0) + r_k$, with $S_1 = S_2 = S$ if $_\diamond q_2(s)$ never hits $[-{_\diamond\hat{q}_2}(0), -{_\diamond\hat{q}_2}(0) + r_k]$. Note that trajectory $_\diamond q(s)$ lies in $_\diamond\mathcal{S}_0$ in $(0, S_1)$ and (S_2, S). Then one of the following must be true over interval $[0, S_1]$ (similarly $[S_2, S]$):

(a) $|{_\diamond q}(0)| < |{_\diamond q}(S_1)|$ and the cost per unit increase in sum-queue over the interval $[0, S_1]$ is less than J_{***}, i.e.
$$\frac{J_{S_1}({_\diamond f},\ {_\diamond g}) - J_0({_\diamond f},\ {_\diamond g})}{|{_\diamond q}(S_1)| - |{_\diamond q}(0)|} \leq J_{***} \ .$$

(b) $|{_\diamond q}(0)| < |{_\diamond q}(S_1)|$ and the cost per unit increase in sum-queue over the interval $[0, S_1]$ is strictly greater than J_{***}, i.e.
$$\frac{J_{S_1}({_\diamond f},\ {_\diamond g}) - J_0({_\diamond f},\ {_\diamond g})}{|{_\diamond q}(S_1)| - |{_\diamond q}(0)|} > J_{***} \ .$$

(c) $|q(0)| \geq |q(S_1)|$.

If (a) is true for either one of the intervals (suppose its true for $[0, S_1]$,) proceed like this: define vectors $\hat{\gamma},\ \hat{\lambda}$, and $\hat{\mu}$ as the average channel distribution, arrival rate, and service rate respectively over $[0, S_1]$, i.e.,

$$\left(\hat{\gamma}, \hat{\lambda}, \hat{\mu}\right) = \frac{1}{S_1}\int_0^{S_1} \Big(g'(s), f'(s), \mu(s)\Big) ds,$$

By Lemma 1 and the fact that $({_\diamond q}(s) : 0 < s < S_1)$ lies in $_\diamond\mathcal{S}_0$, we have $\mu(s) \in \arg\max_{v \in V_{\gamma(s)}} |v|$. This and the *linearity* of $\mu(s)$ in $\gamma(s)$ (see (1)) implies $\hat{\mu} \in \arg\max_{v \in V_{\hat{\gamma}}} |v|$. Then,

$$|{_\diamond q}(S_1)| - |q(0)| \ = \ \left(|\hat{\lambda}| - |\hat{\mu}|\right) S_1 \ ,$$

$$\leq \quad |\hat{\lambda} - v^*(\hat{\lambda},\hat{\gamma})| S_1 \ , \qquad (11)$$

where $v^*(\hat{\lambda},\hat{\gamma})$ is as defined in (4). Finally,

$$\begin{aligned} J_* > J_{***} &\geq \frac{J_{S_1}(\diamond f, \diamond g) - J_0(\diamond f, \diamond g)}{|\diamond q(S_1)| - |\diamond q(0)|} \ , \\ &\geq \frac{\left(L_{(f)}(\hat{\lambda}) + L_{(g)}(\hat{\gamma})\right) S_1}{|\diamond q(S_1)| - |\diamond q(0)|} \ , \\ &\geq \frac{L_{(f)}(\hat{\lambda}) + L_{(g)}(\hat{\gamma})}{|\hat{\lambda} - v^*(\hat{\lambda},\hat{\gamma})|} \ , \end{aligned}$$

where the first inequality follows from the assumption that (a) is true, the second from convexity of rate functions, and the last one from (11). However, by the definition of J_* in (5), the RHS of last inequality cannot be less than J_*, hence giving the contradiction we needed.

Now, if (a) is not true for both intervals $[0, S_1]$ and $[S_2, S]$, then we must have, for arbitrarily small $\epsilon_1 > \epsilon$,

$$\frac{J_{S_2}(\diamond f, \diamond g) - J_{S_1}(\diamond f, \diamond g)}{|\diamond q(S_2)| - |\diamond q(S_1)|} \leq J_{***} \ ,$$

and, for some fixed $\theta_1 > 0$,

$$\begin{aligned} J_{S_2}(\diamond f, \diamond g) - J_{S_1}(\diamond f, \diamond g) &\geq \theta_1 S \ , \\ |\diamond q(S_2)| - |\diamond q(S_1)| &\geq \theta_1 S \ , \\ S_2 - S_1 &\geq \theta_1 S. \end{aligned}$$

Both the terminal values, ${}_\diamond q_2(S_1)$ and ${}_\diamond q_2(S_2)$, lie within the bounded interval $[-{}_\diamond\hat{q}_2(0), -{}_\diamond\hat{q}_2(0) + r_k]$. Any such LFSP can be extended from the interval $[S_1, S_2]$ to a longer interval $[S_1, S_2']$, so that ${}_\diamond q_2(S_2') = {}_\diamond q_2(S_1)$. Moreover, it can be done in such a way that the increments $J_{S_2'}(\diamond f, \diamond g) - J_{S_2}(\diamond f, \diamond g)$ and $|\diamond q(S_2')| - |\diamond q(S_2)|$ are uniformly bounded on the terminal values ${}_\diamond q_2(S_1)$ and ${}_\diamond q_2(S_2)$. Finally, constant S can be chosen large enough such that,

$${}_\diamond q_2(S_1) = {}_\diamond q_2(S_2'), \quad |\diamond q(S_2')| > |\diamond q(S_1)| \ ,$$

and for some $\epsilon_4 > 0$,

$$\frac{J_{S_2'}(\diamond f, \diamond g) - J_{S_1}(\diamond f, \diamond g)}{|\diamond q(S_2')| - |\diamond q(S_1)|} < J_* - \epsilon_4 \ . \qquad (12)$$

Set $b_0 = -{}_\diamond\hat{q}_2(s)$; for each $1 \leq m \leq k$, choose a $b_m \in [-{}_\diamond\hat{q}_2(s) + r_{m-1}, -{}_\diamond\hat{q}_2(s) + r_m]$ such that the counting measure of set $\{s \in [0, S_2'] : \ {}_\diamond q(s) = b_m\}$ is finite; lastly, choose a $b_{k+1} < \infty$ large enough that ${}_\diamond q_2(s) < b_{m+1}$, (such $\{b_m\}$ exist since ${}_\diamond q$ is Lipschitz.) For each $1 \leq m \leq k+1$, define set $B_m = \{s \in [S_1, S_2'] : \ b_{m-1} \leq {}_\diamond q_2(s) \leq b_m\}$. Then for all $m \in \{1, 2, ..., k+1\}$,

$$\int_{B_m} {}_\diamond q_2'(s) \ ds = 0 \ . \qquad (13)$$

Moreover, there exists a set B_m such that,

$$\int_{B_m} |\diamond q'(s)| \ ds > 0 \ ,$$

and,

$$\frac{\int_{B_m} \left(L_{(f)}\left(\diamond f'(s)\right) + L_{(g)}\left(\diamond g'(s)\right)\right) ds}{\int_{B_m} |\diamond q'(s)| \ ds} \leq J_* - \epsilon_4 \ ,$$

otherwise (12) will not hold. Define $\hat{\lambda}$, $\hat{\gamma}$, and $\hat{\mu}$ as,

$$\left(\hat{\lambda}, \hat{\gamma}, \hat{\mu}\right) = \frac{1}{\nu(B_m)} \int_{B_m} \left(\diamond f'(s), \diamond g'(s), \mu(s)\right) ds \ .$$

Since set B_m intersects with at most two adjacent partitions ${}_\diamond\mathcal{S}_{(\cdot)}$, the service vector $\hat{\mu}$ is a convex combination of at most two adjacent vertices of capacity region $V_{\hat{\gamma}}$ (see (8).) Therefore, $\hat{\mu}$ is necessarily a maximal element of $V_{\hat{\gamma}}$. This together with the fact $\hat{\lambda}_2 = \hat{\mu}_2$, (which follows from (13),) we get $\hat{\mu} = \arg\min_{v \in V_{\hat{\gamma}}} |(\hat{\lambda} - v)^+| = v^*(\hat{\lambda}, \hat{\gamma})$. Then,

$$\int_{B_m} |\diamond q'(s)| \ ds = |\hat{\lambda} - v^*(\hat{\lambda},\hat{\gamma})| \ \nu(B_m) \ .$$

Finally,

$$\begin{aligned} J_* - \epsilon_4 &\geq \frac{\int_{B_m} \left(L_{(f)}\left(\diamond f'(s)\right) + L_{(g)}\left(\diamond g'(s)\right)\right) ds}{\int_{B_m} |\diamond q'(s)| \ ds} \ , \\ &\geq \frac{L_{(f)}(\hat{\lambda}) + L_{(g)}(\hat{\gamma})}{|\hat{\lambda} - v^*(\hat{\lambda},\hat{\gamma})|} \ . \end{aligned}$$

By the definition of J_*, the RHS of last inequality cannot be less than J_*, hence giving the contradiction we needed. ∎

REFERENCES

[1] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar, and P. Whiting, "Scheduling in a queuing system with asynchronously varying service rates," *Probab. Eng. Inf. Sci.*, vol. 18, no. 2, pp. 191–217, 2004.

[2] S. Shakkottai and A. Stolyar, "Scheduling for multiple flows sharing a time-varying channel: The Exponential Rule," *American Mathematical Society Translations, Series 2*, vol. 207, 2002.

[3] B. Sadiq, S. J. Baek, and G. de Veciana, "Delay-optimal opportunistic scheduling and approximations: the Log rule," *Submitted*.

[4] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Trans. Inf. Theory*, vol. 39, no. 2, pp. 466–478, Mar 1993.

[5] E. M. Yeh and A. S. Cohen, "Delay optimal rate allocation in multiaccess fading communications," in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, 2004, pp. 140–149.

[6] B. Hajek, "Optimal control of two interacting service stations," *Automatic Control, IEEE Transactions on*, vol. 29, no. 6, pp. 491–499, Jun 1984.

[7] D. Bertsimas, I. Paschalidis, and J. Tsitsiklis, "Asymptotic buffer overflow probabilities in multiclass multiplexers: an optimal control approach," *Automatic Control, IEEE Transactions on*, vol. 43, no. 3, pp. 315–335, Mar 1998.

[8] A. L. Stolyar, "Large deviations of queues sharing a randomly time-varying server," *Queueing Syst. Theory Appl.*, vol. 59, no. 1, pp. 1–35, 2008.

[9] C. Zhou and G. Wunder, "General stability conditions in wireless broadcast channels," in *Forty-Sixth Annual Allerton Conference on Communication, Control, and Computing*, Sep 2008.

[10] A. Dembo and O. Zeitouni, *Large deviations techniques and applications*, 2nd ed. Springer, 1998.

[11] B. Sadiq and G. de Veciana, "Optimality and large deviations of queues under the psuedo-Log rule scheduling," *Technical Report*, 2008. [Online]. Available: www.ece.utexas.edu/~gustavo/logruleoptimality.pdf