

Large deviation sum-queue optimality of a radial sum-rate monotone opportunistic scheduler *

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

June 24, 2009

Abstract

A centralized wireless system is considered that is serving a fixed set of users with time varying channel capacities. An opportunistic scheduling rule in this context selects a user (or users) to serve based on the current channel state and user queues. Unless the user traffic is symmetric and/or the underlying capacity region a polymatroid, little is known concerning how performance optimal schedulers should tradeoff *maximizing current service rate* (being opportunistic) versus *balancing unequal queues* (enhancing user-diversity to enable future high service rate opportunities). By contrast with currently proposed opportunistic schedulers, e.g., MaxWeight, Exp Rule, a radial sum-rate monotonic (RSM) scheduler de-emphasizes queue-balancing in favor of greedily maximizing the system service rate as the queue-lengths are scaled up linearly. In this paper it is shown that an RSM opportunistic scheduler, p-Log Rule, is not only throughput-optimal, but also maximizes the asymptotic exponential decay rate of the sum-queue distribution for a two-queue system. The model and result complement existing optimality results for opportunistic scheduling and point to RSM schedulers as a good design choice given the need for robustness in wireless systems with both heterogeneity and high degree of uncertainty.

Key words and phrases: Opportunistic scheduling, wireless communication, large deviations, queues sharing time-varying server.

1 Introduction

We consider a wireless node shared by two users, where each user's data arrives to a queue at the node as a random process and each user's channel, in terms of data rate that it can support, varies randomly over time. If the channel state is available, a policy can schedule users so as to exploit favorable channels, e.g., schedule the user which currently has the highest rate – this is referred to as opportunistic scheduling (see, e.g., [1, 2]). More generally, a scheduler can be permitted to allocate rates to the users from a compact polytope, where the polytope depends on the channel state.

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

An opportunistic or channel-aware scheduler, however, may not be stable, i.e., keep the user queues bounded, unless it is chosen carefully, e.g., using prior knowledge of the arrival and channel process. 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 the current service rate* (scheduling the queue seeing the highest rate) with *balancing unequal queues* (scheduling the longest queue). Note that by balancing queues, one can increase the likelihood of more non-empty queues in the future and, therefore, can potentially achieve higher service rates in the future. Moreover, the *performance* of any scheduler, in terms of Quality-of-Service (QoS) measures like mean delay or probability of queue overflow, depends on how this tradeoff is made.

Much work to date has focused on engineering queue-and-channel-aware schedulers that are *throughput-optimal*, i.e., schedulers that can achieve stability without any knowledge of arrival or channel statistics, if stability is at all feasible. Examples are MaxWeight [3], Exponential (Exp) rule [4], and Log rule [5]. In fact, necessary and sufficient conditions for MaxWeight-type schedulers to be throughput-optimal have been shown in [6, 7].

Stability, however, is a weak form of optimality. In view of various QoS goals, 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 users, or ones which minimize the probability that either the sum-queue or the longest queue overflows a large buffer. These schedulers are harder to characterize for channels/servers¹ with time-varying capacity, but some results are available that we briefly discuss next.

In [8] and [9] the Longest-Connected-Queue (LCQ) and Longest-Queue-Highest-Possible-Rate (LQHPR) scheduling policies are introduced. Strong results are shown for these policies; they stochastically minimize the max- and sum- queue processes, and thus also the tails of max- and sum- queue distributions and mean delay. However, in addition to assuming certain symmetry conditions on arrival and channel statistics, [8] is limited to on-off channel capacities where only a single queue can be scheduled per time slot, and [9] assumes that the scheduler can allocate service rates from the information theoretic polymatroid multiuser capacity region associated with the current channel state. In both cases, the above-mentioned tradeoff between queue balancing and service rate maximization is absent. Indeed in [8], all policies that pick a connected queue result in the same service rate, whereas, in the case of [9], 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 total service rate. Thus one can achieve the *queue balancing* goal, without compromising the total *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 [10]. The paper considers minimizing the mean delay in a two queue system where each queue has a dedicated server and a third server can be dynamically shared between them. As a result, the two queues can be allocated service rates from a polymatroid capacity region, thus the objective of queue balancing can again be achieved without

¹The terms channel and server, and sometimes the terms user and queue, are used interchangeably through this section.

compromising the total service rate. However, without the underlying symmetry assumptions of [8] and [9], and using a dynamic programming approach, only the existence of a monotone increasing switching curve on the state space of queue process is shown; note that the switching curve under LCQ and LQHPR policies lies along the line where both queues are equal. For a system with a general compact, convex, and coordinate convex capacity region and any finite number of queues, [11] gives a large deviations principle (LDP) for transient queue process under MaxWeight scheduler. This LDP can be used to compute, e.g, the asymptotic probability of sum-queue or max-queue overflow, as well as the corresponding likely modes of overflow. Although the capacity region is not changing over time, the region is such that a scheduler must tradeoff maximizing total service rate with balancing unequal queues. Therefore this result is insightful in relating the modes overflow to the tradeoff made by the MaxWeight scheduler. A more recent work [12] gives many sources large deviations for MaxWeight scheduler for a similar capacity region.

Finally, relaxing the symmetry assumptions of [8] and [9], the works in [13–15] consider the asymptotic probability of max-queue overflow. The server capacity in [13], though time-varying, is identical for all users at any given time, thus the need to tradeoff queue-balancing versus service rate maximization is again absent. In fact, the sum-queue process in [13] is identical for all work conserving schedulers. However [14] and [15] consider a server with *asynchronously* time-varying capacity across users. [14] studies the asymptotic probability of max-queue overflow under MaxWeight scheduler and shows that as the exponent of queue length in the MaxWeight scheduler, α , becomes large, the asymptotic probability of max-queue overflow under MaxWeight approaches the minimum achievable under any other scheduler. A stronger result is shown in [15], that is, the Exp rule scheduler in fact minimizes the steady state asymptotic probability of max-queue overflow. Indeed the models in [14] and [15] accurately capture a wireless channel shared by heterogenous users, and exhibit the tradeoff between queue-balancing and service rate maximization. Existence of this tradeoff also implies that, unlike the LCQ and LQHPR policies, the asymptotic optimality of Exp rule does not translate to minimizing the asymptotic probability of sum-queue overflow or the mean delay. In order to minimize the asymptotic probability of max-queue overflow, the desired² mode of overflow is one where all queues (more precisely, the set of overflowing queues which then exclusively share the server,) grow at the same rate and overflow at the same time. This constrains the system throughput, while, of course, keeping the queue lengths *equal* across users.

Contributions: In this paper we address the problem of minimizing the asymptotic probability of sum-queue overflow. First, we give a tight lower bound on this probability under *any* scheduler. By contrast with the desired mode overflow to minimize the asymptotic probability of *max*-queue overflow, we will see that in order to minimize the asymptotic probability of sum-queue overflow, the desired mode of overflow is one where the system throughput is the highest possible and queues may build up at different rates. Second, we show that a radial sum-rate monotone scheduler (see [5] or Section 3 for definition of radial sum-rate monotonicity), called the pseudo-Log rule, minimizes the asymptotic probability of sum-queue overflow. Although our focus is on overflows of the sum-queue instead of overflows of the max-queue as in [15], the general technique of proof in [15] lends itself well to our problem and we rely heavily on the results developed therein. Other desirable

²By the “desired” mode of max-queue overflow we mean the mode which gives the lower bound on the asymptotic probability of max-queue overflow under *any* scheduler, as given in [15]; the likely mode of max-queue overflow under the Exp rule is indeed the desired one.

features of radial sum-rate monotone schedulers have been explained in [5], which include,

- (a) reducing mean delay;
- (b) graceful degradation of service in terms of fraction of users that can meet their Quality-of-Service requirements during transient overloads;
- (c) robustness to uncertainty in the traffic and channel statistics.

The rest of the paper is organized as follows. The system model is described in Section 2. Queue-and-channel aware schedulers of interest, namely, MaxWeight, Exp rule, and Log rule, and the property of radial sum-rate monotonicity are reviewed in Section 3, followed by the introduction of pseudo-Log scheduling rule in Section 4. The main result of the paper is summarized in Section 5. Some preliminary discussion and relevant large deviation principles follow in Section 6. The proofs for the lower and the upper bounds stated in the main result of the paper are given in Section 7 and 8 respectively. After defining local fluid sample paths and developing essential results in Section 9, the optimality of the p-Log rule, i.e., the last part of main result is proved in Section 10. Immediate extensions of the main result to some other interesting system models are presented in the concluding Section 11.

2 System model

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

A server with randomly varying service rates is available to the two queues and modeled as follows. The server has a time-varying state that is modeled by an i.i.d. random process $(\mathbf{m}(t), t = 0, 1, \dots)$, where $\mathbf{m}(t) \in \mathcal{M} = \{1, 2, \dots, M\}$ for some finite $M > 0$ denotes the state of the server over $[t, t+1)$, and is drawn from distribution $\pi = (\pi_1, \dots, \pi_M) > 0$. Associated with each server state $m \in \mathcal{M}$ is a 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 packets from Queue 1, or at most μ_2^m packets from Queue 2. The scheduling problem is thus to allocate the server to one of the queues for each time slot such that a given optimality criterion is met. This will be formally described later.

At any integer time t , $\mathbf{Q}(t) = (\mathbf{Q}_i(t), i \in I) \in \mathbb{Z}_+^2$ is a random vector, where $\mathbf{Q}_i(t)$ denotes the number of packets in the i^{th} queue at the end of time slot $[t-1, t)$. Let S_t denote a system sample path up to time t , i.e., $S_t \equiv \left((m(\tau), Q(\tau), A(\tau-1)), \tau \leq t \right)$ and \mathcal{S}_t denote the space of all feasible realizations S_t . Let $i_t^* : \mathcal{S}_t \rightarrow I$ denote the queue scheduled to receive service during time slot $[t, t+1)$, i.e., we assume that the system sample path S_t is available for making the scheduling decision for time slot $[t, t+1)$. The evolution of the queue process under the scheduling decision $i_t^*(S_t)$ is given by,

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

The sequence of functions $(i_t^*(\cdot), t = 0, 1, \dots)$ is called a *scheduler* or *scheduling policy*. It is easy to see that under a static state-feedback scheduler, i.e., one where $i_t^*(S_t) = i^*(Q(t), m(t))$, the process $(Q(t), t = 0, 1, \dots)$ forms a discrete time Markov chain on \mathbb{Z}_+^2 .

In the sequel, we will 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 $\lfloor t \rfloor$. All such processes and functions lie in the space of *real-valued right continuous functions with left limits*, 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 (Ω, \mathcal{F}, P) be a probability space that is large enough to define all the random processes in this paper.

For a given weight vector $b = (b_1, b_2) > 0$, we are interested in finding a scheduler which, informally speaking, minimizes the tail of the distribution of weighted sum-queue $\sum_{i \in I} b_i Q_i(\cdot)$. We will show that under a static state-feedback scheduler, namely the p-Log rule described in Section 4, the asymptotic probability of weighted sum-queue overflow in the steady state, i.e.,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log P \left(\sum_{i \in I} b_i Q_i(0) \geq n \right) ,$$

is *minimized* (see Theorem 1 for a formal statement.) The p-Log rule depends only on weight vector b and does not require any knowledge of arrival or server-state distributions.

Capacity region

For each server state $m \in \mathcal{M}$, let $V^m \in \mathbb{R}_+^2$ denote the convex hull of vertices $(0, 0)$, $(0, \mu_2^m)$, and $(\mu_1^m, 0)$. Then, conditional on the server being in state m , the average service jointly offered to the two queues under any scheduling rule (such that the average exists) lies in the *triangle* V^m . Define the capacity region V_π as the set of average service vectors offered to the two queues under all possible scheduling rules, then V_π is a convex polyhedron given by the weighted Minkowski sum of regions V^m , i.e.,

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

See Fig. 1-a for a graphical illustration of capacity region: the server has $M = 6$ states with some distribution π and vectors μ^m in set $\{(1, 4), (3, 4), (1, 1), (4, 3), (4, 1), (1, 0)\}$.

Let $\{r_1, \dots, r_{M'}\}$ for some $M' \leq M$ be the set of strictly positive and finite slopes of the outer normal vectors to the facets of capacity region V_π . The slopes are indexed such that $0 < r_1 < r_2 < \dots < r_{M'} < \infty$. Also, let $r_0 = 0$ and $r_{M'+1} = \infty$. For example, see Fig. 1-a for a depiction of a capacity region with $M' = 5$ facets with outer normal slopes in $(0, \infty)$. Finally, let $\hat{V}_\pi = \{\hat{v}^1, \dots, \hat{v}^{M'+1}\}$ denote the set of maximal vertices of the capacity region V_π . The vertices in set \hat{V}_π are indexed such that $\hat{v}_1^1 > \hat{v}_1^2 > \dots > \hat{v}_1^{M'+1}$, i.e., the vertex \hat{v}^m lies at the intersection of the facets with outer normal slopes r_{m-1} and r_m ; e.g., see vertex \hat{v}^2 in Fig. 1-a.

We assume that there exists a $v \in V_\pi$ such that $\bar{\lambda} < v$, which is a sufficient condition for stabilizability of the queues [3]. If the above condition is met, then there exists at least one static state feedback scheduler under which the Markov chain $(\mathbf{Q}(t), t = 0, 1, \dots)$ is ergodic.

Remark 1 Since we assumed that the server can be allocated to at most one queue per time slot, the region V^m is obtained by taking the convex hull of service vectors in the set $\{(0, 0), (\mu_1^m, 0), (0, \mu_2^m)\}$. However, we can relax this assumption and allow the server to be shared between the two queues during a time slot. That is, we can associate with each server state m a set of k_m service vectors $\{(\mu_1^m(1), \mu_2^m(1)), \dots, (\mu_1^m(k_m), \mu_2^m(k_m))\}$, and allow the server to operate at any one of these service vectors. In this more general case, each region V^m will be an arbitrary convex polyhedron obtained by taking the convex hull of all feasible service vectors associated with server state m . For example, V^m can be information theoretic polymatroids as in [9]. The optimality results presented in this paper will still hold with such a relaxation; see Section 11 for some details.

3 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. We formally define them next. Let the vector fields $h^{mw}(\cdot)$, $h^{exp}(\cdot)$ and $h^{log}(\cdot)$ on \mathbb{R}_+^2 be given as follows: for all $x \in \mathbb{R}_+^2$,

$$\begin{aligned} h^{mw}(x) &= (b_i x_i^\alpha, i \in I) , \\ h^{exp}(x) &= \left(b_i \exp\left(\frac{a_i x_i}{c + (0.5(a_1 x_1 + a_2 x_2))^\eta}\right), i \in I \right) , \\ h^{log}(x) &= (b_i \log(1 + a_i x_i), i \in I) , \end{aligned}$$

for any fixed positive b_i 's, a_i 's, α , c , and $0 < \eta < 1$. Then, when the system is in state $(\mathbf{Q}(t), \mathbf{m}(t)) = (Q, m) \in \mathbb{Z}_+^2 \times \mathcal{M}$, the MaxWeight scheduler serves a queue i_{mw}^* given by,

$$i_{mw}^*(Q, m) \in \arg \max_{i \in I} h_i^{mw}(Q) \mu_i^m , \quad (2)$$

augmented with any fixed tie-breaking rule. The Exp rule i_{exp}^* and the Log rule i_{log}^* are defined similarly by substituting h^{exp} and h^{log} respectively in place of h^{mw} . Indeed numerous vector field based throughput-optimal schedulers can be engineered so as to respond differently to the disparity among the users' queue lengths, i.e., make different tradeoffs between *service rate maximization* and *queue balancing*. For reference see, e.g., [7], which gives necessary and sufficient conditions for a vector field based scheduler to be throughput-optimal.

We refer to a scheduler as *radial sum-rate monotone* if, as the queues scale up linearly, the scheduling rule allocates the 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 average service offered to the queues under a static state-feedback scheduler i^* , conditional on queue state being Q , i.e.,

$$v(Q) = (E[\mu_i^{m^*} \mathbb{1}_{\{i^*(Q, m) = i\}}], i \in I) , \quad (3)$$

where expectation is with respect to (random) \mathbf{m} drawn from distribution π .

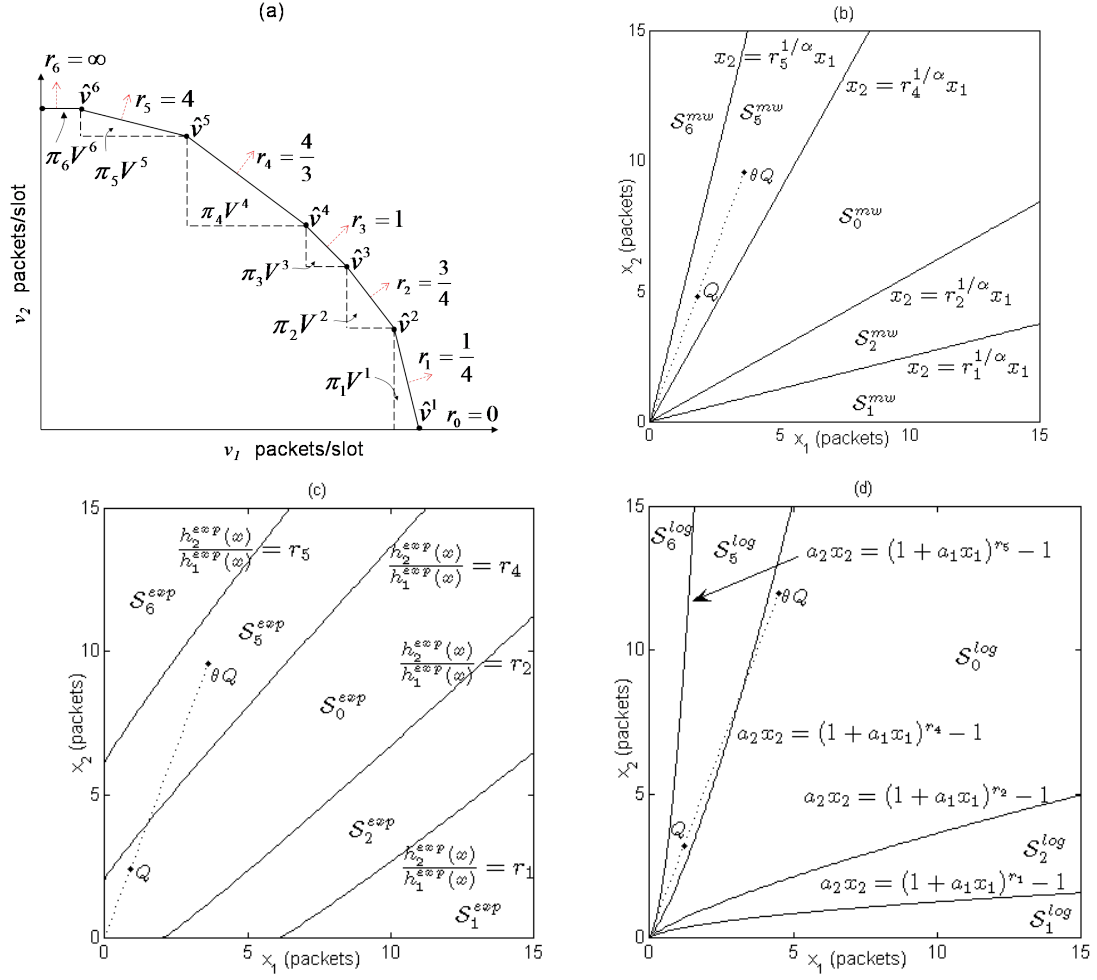


Figure 1: (a) The capacity region for $\mu^m \in \{(1,4), (3,4), (1,1), (4,3), (4,1), (1,0)\}$, depicting Minkowski addition, outer-normal vectors and maximal vertices; the resulting switching curves under (b) MaxWeight, (c) Exp rule, (d) Log rule, for weight vector $b = (1,1)$.

Definition 1 A scheduling policy i^* is *radial sum-rate monotone* with respect to weight vector $b > 0$ if for any Q and scalar θ such that $\theta Q \in \mathbb{Z}_+^n$, the weighted sum of expected offered service, $\langle b, v(\theta Q) \rangle$, is an increasing function of θ , and $\lim_{\theta \rightarrow \infty} \langle b, v(\theta Q) \rangle = \max(\langle b, y \rangle : y \in V_\pi \text{ and } y_i = 0 \text{ if } Q_i = 0)$.

Let $v^{mw}(Q)$ denote the expected service vector under the MaxWeight scheduler, i.e., the vector obtained by substituting i_{mw}^* for i^* in (3); similarly, let $v^{exp}(Q)$ and $v^{log}(Q)$ denote the expected service vectors under the Exp and Log rule respectively. Also, fix a weight vector $b > 0$. Next, under schedulers Maxweight, Exp rule, and Log rule respectively, we will identify the sets \mathcal{S}_0^{mw} , \mathcal{S}_0^{exp} , and \mathcal{S}_0^{log} given *approximately*³ by,

$$\mathcal{S}_0^{(\cdot)} \approx \left\{ Q \in \mathbb{Z}_+^2 : v^{(\cdot)}(Q) \in \arg \max_{y \in V_\pi} \langle y, b \rangle \right\} ,$$

in words, the set of queue states such that the expected service vector $v^{(\cdot)}(Q)$ has the maximum weighted sum with respect to weight vector b . For example, for the capacity region shown in Fig. 1-*a* and weight vector $b = (1, 1)$, Fig. 1-*b-d* illustrates the sets \mathcal{S}_0^{mw} , \mathcal{S}_0^{exp} , and \mathcal{S}_0^{log} . The set \mathcal{S}_0^{mw} is a cone, the set \mathcal{S}_0^{exp} a cylinder with gradually increasing diameter, and the set \mathcal{S}_0^{log} resembles a French horn. The set \mathcal{S}_0^{log} is such that for any $Q > 0$, we have $\theta Q \in \mathcal{S}_0^{log}$ for all θ large enough, indicating that the Log rule is radial sum-rate monotone. The MaxWeight and Exp rule are not radial sum-rate monotone. In fact, Exp rule is the ‘opposite’ of radial sum-rate monotone in that, for any Q such that $a_1 Q_1 \neq a_2 Q_2$, we have $\theta Q \notin \mathcal{S}_0^{exp}$ for all θ large enough. A formal description of the sets $\mathcal{S}_0^{(\cdot)}$ is as follows.

The vector $h^{(\cdot)}(Q)$ can be shown to be an outer normal vector to the capacity region V_π at point $v^{(\cdot)}(Q)$ for each scheduler $(\cdot) \in \{mw, exp, log\}$, i.e.,

$$v^{(\cdot)}(Q) \in \arg \max_{y \in V_\pi} \langle y, h^{(\cdot)}(Q) \rangle , \quad (4)$$

See, e.g., Lemma 2.1 of [16]. Recall the set of outer normal slopes $\{r_0, r_1, \dots, r_{M'}, r_{M'+1}\}$, and let r_k be the largest slope strictly less than b_2/b_1 (i.e. the slope of vector b) and r_l the smallest slope strictly greater than b_2/b_1 ; thus if $r_{k+1} \neq b_2/b_1$ then $l = k + 1$, otherwise $l = k + 2$. For example, in Fig. 1-*a*, for weight vector $b = (1, 1)$, we have $k = 2$ and $l = 4$. For each scheduler $(\cdot) \in \{mw, exp, log\}$ we define the set,

$$\mathcal{S}_0^{(\cdot)} = \left\{ x \in \mathbb{R}_+^2 : r_k < \frac{h_2^{(\cdot)}(x)}{h_1^{(\cdot)}(x)} < r_l \right\} .$$

Then, by (4), for any $Q \in \mathcal{S}_0^{(\cdot)} \cap \mathbb{Z}_+^2$, we must have $v^{(\cdot)}(Q) \in \arg \max_{y \in V_\pi} \langle y, b \rangle$. Each region $\mathcal{S}_0^{(\cdot)}$ is bounded by *switching curves* given by,

$$\left\{ x \in \mathbb{R}_+^2 : \frac{h_2^{(\cdot)}(x)}{h_1^{(\cdot)}(x)} = r_k \right\} \quad \text{and} \quad \left\{ x \in \mathbb{R}_+^2 : \frac{h_2^{(\cdot)}(x)}{h_1^{(\cdot)}(x)} = r_l \right\} .$$

³Under the formal definition of set $\mathcal{S}_0^{(\cdot)}$ given later, the queue states on the boundary of set $\{Q \in \mathbb{Z}_+^2 : v^{(\cdot)}(Q) \in \arg \max_{y \in V_\pi} \langle y, b \rangle\}$ may or may not lie in $\mathcal{S}_0^{(\cdot)}$.

For any queue state Q lying on these switching curves, the argmax in (4) is not unique and whether $v^{(\cdot)}(Q)$ lies in the set $\arg \max_{y \in V_\pi} \langle y, b \rangle$ depends on the tie breaking rule associated with (2). For each $m \in \{1, \dots, k, l+1, \dots, M'+1\}$, we can also define a set,

$$\mathcal{S}_m^{(\cdot)} = \left\{ x \in \mathbb{R}_+^2 : r_{m-1} < \frac{h_2^{(\cdot)}(x)}{h_1^{(\cdot)}(x)} < r_m \right\} .$$

Then, again by (4), for any $Q \in \mathcal{S}_m^{(\cdot)} \cap \mathbb{Z}_+^2$, we must have $v^{(\cdot)}(Q) = \hat{v}^m$. That is, each set $\mathcal{S}_m^{(\cdot)}$ for $m \in \{1, \dots, k, l+1, \dots, M'+1\}$ is associated with a unique vertex. The switching curve between any two regions $\mathcal{S}_m^{(\cdot)}$ and $\mathcal{S}_{m+1}^{(\cdot)}$ for $m \in \{1, \dots, k-1, l+1, \dots, M'\}$ is given by,

$$\left\{ x \in \mathbb{R}_+^2 : \frac{h_2^{(\cdot)}(x)}{h_1^{(\cdot)}(x)} = r_m \right\} ,$$

and as before, for any queue state Q lying on this switching curve, the argmax in (4) is not unique and whether it lies in set $\mathcal{S}_m^{(\cdot)}$ or $\mathcal{S}_{m+1}^{(\cdot)}$ depends on the tie breaking rule associated with (2). Fig. 1-*b-d* also show the switching curves and sets $\mathcal{S}_m^{(\cdot)}$ under MaxWeight, Exp rule, and Log rule. In fact, for MaxWeight and Log rule, the switching curves can be given in closed form, e.g., for Log rule, the switching curves are given by,

$$a_2 x_2 = (1 + a_1 x_1)^{\frac{b_1}{b_2} r} - 1 ,$$

for $r \in \{r_1, \dots, r_{M'}\}$.

4 The pseudo-Log scheduling rule

In this section we introduce a static state-feedback and radial sum-rate monotone scheduler, denoted the pseudo-Log (p-Log) rule. We subsequently show in Theorem 1 that the p-Log rule minimizes the asymptotic probability of weighted sum-queue overflow for any weight vector b . The p-Log rule takes weight vector b as a parameter but does not require any knowledge of arrival or server-state distributions.

We will define the p-Log scheduling rule through a vector field $h = (h_1, h_2)$ on \mathbb{R}_+^2 . For any $x \in [0, 1]^2$, let $h(x) = 0$. For any $x \in \mathbb{R}_+^2 \setminus [0, 1]^2$,

$$\begin{aligned} \text{if } x_1 \geq x_2, & \quad \text{then let } h_1(x) = b_1 \sqrt{x_1} , & \quad h_2(x) = b_2 \min(x_2, \sqrt{x_1}) ; \\ \text{if } x_1 < x_2, & \quad \text{then let } h_1(x) = b_1 \min(x_1, \sqrt{x_2}) , & \quad h_2(x) = b_2 \sqrt{x_2} . \end{aligned} \quad (5)$$

The p-Log rule is given as follows: when the queues are 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,

$$i_{pLog}^*(Q, m) \in \arg \max_{i \in I} h_i(Q) \mu_i^m , \quad (6)$$

where, in the case of a tie, if $Q_1 \geq Q_2$ the server is allocated to Queue 1, otherwise to Queue 2.

Note that it is only the slope, $\frac{h_2(Q)}{h_1(Q)}$, of the vector $h(Q)$ that determines the scheduling decision, and so, for example, when $Q_1 \geq Q_2$, $Q_1 \neq 0$, the slope is given by,

$$\frac{h_2(Q)}{h_1(Q)} = \frac{b_2}{b_1} \min \left(1, \frac{Q_2}{\sqrt{Q_1}} \right) .$$

Switching curves under the p-Log rule

Let $v^{pLog}(Q)$ denote the vector of average service offered to the queues under p-Log rule, conditional on queue state being Q , i.e., the vector obtain by substituting i_{pLog}^* for i^* in (3). Recall the set of outer normal slopes $\{r_0, r_1, \dots, r_{M'}, r_{M'+1}\}$ and, in particular, the slopes r_k and r_l from this set as defined in Section 3. Similar to the sets $\mathcal{S}_0^{(\cdot)}$ for $(\cdot) \in \{mw, exp, log\}$, we define the set \mathcal{S}_0 for p-Log rule as,

$$\begin{aligned} \mathcal{S}_0 &= \left\{ x \in \mathbb{R}_+^2 : r_k < \frac{h_2(x)}{h_1(x)} < r_l \right\} , \\ &= \left\{ x \in \mathbb{R}_+^2 : \frac{b_1}{b_2} r_k \sqrt{x_1} < x_2 < \left(\frac{b_1}{b_2} r_l x_1 \right)^2 \right\} . \end{aligned}$$

The region $\mathcal{S}_0^{(\cdot)}$ is shaped like a French horn, see Fig. 2 for an illustration. For any $Q \in \mathcal{S}_0 \cap \mathbb{Z}_+^2$, we have $v^{pLog}(Q) \in \arg \max_{y \in V_\pi} \langle y, b \rangle$. Therefore, we will refer to the region \mathcal{S}_0 as the weighted-max-sum rate region (with respect to weight vector b). Moreover, for any $Q > 0$, we have $\theta Q \in \mathcal{S}_0$ for all θ large enough, thus indicating that p-Log rule is radial sum-rate monotone. Also, for each $m \in \{1, \dots, k, l+1, \dots, M'+1\}$, we define a set \mathcal{S}_m as,

$$\mathcal{S}_m = \left\{ x \in \mathbb{R}_+^2 : r_{m-1} < \frac{h_2(x)}{h_1(x)} < r_m \right\} .$$

Then for any $Q \in \mathcal{S}_m \cap \mathbb{Z}_+^2$, we have $v^{pLog}(Q) = \hat{v}^m$. All switching curves in the half plane $\{x_1 \geq x_2\}$ are given by,

$$x_2 = \frac{b_1}{b_2} r \sqrt{x_1} ,$$

for $r \in \{r_0, \dots, r_k\}$ and $x_1 \geq 1$. Similarly, all switching curves in the half plane $x_2 > x_1$ are given by,

$$x_1 = \frac{b_2}{b_1} r^{-1} \sqrt{x_2} ,$$

for $r \in \{r_l, \dots, r_{M'+1}\}$ and $x_2 \geq 1$. We will refer to the collection of switching curves and regions \mathcal{S}_m as a *partition* of \mathbb{R}_+^2 (or the queue state space \mathbb{Z}_+^2) under the p-Log rule. It will be useful to note that this partition, as well as the partitions under MaxWeight, Exp rule, and Log rule, depend only on the set of vectors $\{\mu^1, \dots, \mu^M\}$ associated with the M server states and not on the distribution $\pi = (\pi_1, \dots, \pi_M) > 0$ over these states.

5 Main result

The following three-part theorem summarizes the main results of this paper. It includes a lower bound on the tail of the weighted sum-queue overflow probability, an upper bound on the same,




Figure 2: The switching curves under the p-Log rule for the capacity region depicted in Fig. 1-a and weight vector $b = (1, 1)$.

and the optimality of the p-Log rule. The first part is proved in Section 7, the second in Section 8 (and Appendix), while the last in Sections 9 and 10.

Theorem 1 *For the system model detailed in Section 2, the following hold.*

(i) *Given a weight vector $b = (b_i : b_i > 0, i \in I)$, there exists finite $T_0 > 0$ such that for any $t > T_0$ and under any scheduling rule starting in any initial state $Q(0)$, we have the following lower bound,*

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log P \left(\sum_{i \in I} b_i Q_i(nt) \geq n \right) \geq -J_* ,$$

where J_* is defined in Section 7.

(ii) *For a stabilizable system, under p-Log scheduling rule the process $(Q(t), t = 0, 1, \dots)$ forms an ergodic Markov chain, and we have the following upper bound for a random vector Q drawn from the stationary distribution of the Markov chain,*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log P \left(\sum_{i \in I} b_i Q_i \geq n \right) \leq -J_{**} ,$$

where J_{**} is defined in Section 8.

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

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

Remark 2 Since a scheduler can, in principle, be non-stationary, the lower bound in (i) is expressed in a more general form than the upper bound in (ii) which is specific to a static state-feedback scheduler, namely the p-Log rule. For stationary schedulers under which the process $(Q(t), t = 0, 1, \dots)$ forms an ergodic Markov chain, the lower bound in (i) also implies the same lower bound under the steady state distribution of the Markov chain. This is because the lower bound will hold if the initial state $Q(0)$ were random and drawn from the steady state distribution.

6 Fluid-scaled processes and large deviation principles

In this section, we define sequences of fluid-scaled processes and functions, and a Large Deviation Principle [17] on those sequences, which is used in proving Theorem 1. Define the cumulative arrivals process $\mathbf{F} = (\mathbf{F}(t) = (\mathbf{F}_i(t), i \in I), t \geq 0)$ obtained from the process $(\mathbf{A}(t), t \geq 0)$ as,

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

and cumulative time the channel is in each state $\mathbf{G} = (\mathbf{G}(t) = (\mathbf{G}_m(t), m \in \mathcal{M}), t \geq 0)$ obtained from the process $(\mathbf{m}(t), t \geq 0)$ as,

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

The triplet (Q, F, G) , where $Q = (Q(t), t \geq 0)$ 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 $(\mathbf{Q}, \mathbf{F}, \mathbf{G})$. For each $n = 0, 1, \dots$, let $(\mathbf{Q}^{(n)}, \mathbf{F}^{(n)}, \mathbf{G}^{(n)})$ denote an independent and identically distributed system. We define a corresponding sequence of fluid-scaled processes, denoted by $(\mathbf{q}^{(n)}, \mathbf{f}^{(n)}, \mathbf{g}^{(n)})$, as,

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

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

The arrival and service processes are i.i.d. and bounded and, therefore, satisfy large deviation principles [17]. In particular, for each $i \in I$, define for any scalar $\lambda_i \geq 0$ the rate function $L_i(\cdot)$ for the sequence $\mathbf{f}_i^{(n)}(1)$ as,

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

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

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

For any probability distribution $\gamma = (\gamma_m, m \in \mathcal{M})$, we define the relative entropy $L_{(g)}(\cdot)$ of γ with respect to distribution π as,

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

where $L_{(g)}(\cdot) = \infty$ everywhere outside the simplex in \mathbb{R}_+^M . Now 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 on the interval $[0, t]$, then let,

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

otherwise $J_t(f, g) = \infty$. The functional $J_t(f, g)$ is referred to as the cost of the trajectories (f, g) over the time interval $[0, t]$. The following is a form of Mogulsky theorem [17].

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

$$\begin{aligned} - \inf_{(f,g)} \{J_T(f,g) | (f,g) \in B^\circ\} &\leq \liminf_{n \rightarrow \infty} \frac{1}{n} \log P((\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B) \\ &\leq \limsup_{n \rightarrow \infty} \frac{1}{n} \log P((\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B) \leq - \inf_{(f,g)} \{J_T(f,g) | (f,g) \in \bar{B}\}, \end{aligned}$$

where, B° and \bar{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 $d \in \mathcal{D}^{2+M}$, let $U^n d$ denote the piece-wise linear function obtained by linear interpolation over samples $(d(\frac{ku(n)}{n}), k = 0, 1, \dots)$. The following upper bound is a refinement of Mogulsky theorem and was first introduced in [15].

Proposition 2 For any fixed $T > 0$, consider a sequence in n of the fluid-scaled processes $(\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) = ((\mathbf{f}^{(n)}(t), \mathbf{g}^{(n)}(t)), t \in [0, T])$. 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 $(\mathbf{f}^{(n)}, \mathbf{g}^{(n)})$ in $[0, T]$. Then,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log P((\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B^{(n)}) \leq - \liminf_{n \rightarrow \infty} \inf_{(f,g)} \{J_{T^{(n)}} U^n(f,g) | (f,g) \in B^{(n)}\}, \quad (7)$$

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

Note that by contrast with Proposition 1, $\{B^{(n)}\}$ in Proposition 2 corresponds to a sequence of sets of scaled feasible trajectories, and the bound on the right side is an infimum over the cost of sampled and linearly interpolated scaled trajectories. Thus this proposition provides a refinement allowing, for example, the consideration of large deviations for sets $\{B^{(n)}\}$ which can distinguish among the trajectories in the set $\{f^{(k)} \in \mathcal{D}^2 : f^{(k)} \text{ converges u.o.c. to a Lipschitz } f\}$. That is to say, even if all the fluid scaled trajectories in a set converge to the same limiting trajectory, the events $B^{(n)}$ can be defined to include only a subset of these trajectories.

Also introduced in [15] is a notion of generalized fluid sample path (GFSP) which naturally appears when applying the bound in Proposition 2; we describe this next. Consider a sequence of realizations $(q^{(n)}, f^{(n)}, g^{(n)})$ $n = 1, 2, \dots$ such that along some subsequence (still denoted by $\{n\}$), we have u.o.c convergence

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

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

$$\bar{J}^{(n)} = (\bar{J}_t^{(n)}, t \geq 0) = (J_{t^{(n)}} U^n(f^{(n)}, g^{(n)}), t \geq 0) \rightarrow \bar{J} = (\bar{J}_t, t \geq 0)$$

to some non-negative increasing Lipschitz continuous function \bar{J} . Then the following construct is called a GFSP,

$$\psi = \left[(q^{(n)}, f^{(n)}, g^{(n)}), \bar{J}^{(n)}, n = 0, 1, \dots; (q, f, g); \bar{J} \right],$$

and the function \bar{J} is referred to as the refined cost function of the GFSP. Since this construct contains not only the limiting trajectory (q, f, g) but also the sequence converging to it, the construct is useful when one needs to *zoom in* and study the limiting trajectories (q, f, g) at a finer-than-fluid scaling. Moreover, we will see that for the events $B^{(n)}$ of interest in this paper, the bound in Proposition 2 reduces to an infimum over the refined cost of a well-defined set of GFSPs.




Figure 3: For weights $b_i = 1$, the *optimal* capacity vector $v^*(\lambda^{(2)}, \pi)$ is unique, whereas, vector $v^*(\lambda^{(1)}, \pi)$ is any point on the annotated part of max-sum-rate face of region V_π .

7 Lower bound on overflow probability under any scheduling rule

For any distribution γ on the set \mathcal{M} of server states, let V_γ denote the corresponding capacity region, see (1). Let $b = (b_1, b_2) > 0$ be the given weight vector. For any vector $\lambda \in \mathbb{R}_+^2$, let $v^*(\lambda, \gamma)$ denote a service vector such that,

$$v^*(\lambda, \gamma) \in \arg \max_v \{ \langle b, v \rangle : v \leq \lambda, v \in V_\gamma \} , \quad (8)$$

where, if the argmax is not unique, then $v^*(\lambda, \gamma)$ can be set to any maximizer. For example, see Fig. 3 depicting $v^*(\lambda, \pi)$ for weight vector $b = (1, 1)$ and two hypothetical vectors $\lambda^{(1)}$ and $\lambda^{(2)}$ lying outside capacity region V_γ . The interpretation is as follows: if the arrival process were to exhibit an empirical mean of λ and the server state were to exhibit an empirical distribution γ , then serving the queues according to the service vector $v^*(\lambda, \gamma)$ minimizes the rate of weighted-sum-queue build-up, i.e., $\langle b, \lambda - v^*(\lambda, \gamma) \rangle$.

Finally, we define the minimum cost (per unit increase of the weighted-sum-queue) J_* as,

$$J_* = \min_{\gamma, \lambda} \frac{L_{(g)}(\gamma) + L_{(f)}(\lambda)}{\langle b, \lambda - v^*(\lambda, \gamma) \rangle} , \quad (9)$$

7.1 Proof of Theorem 1-(i)

Let (λ^*, γ^*) be a point that achieves the minimum in (9), and let $T_0 = \langle b, \lambda^* - v^*(\lambda^*, \gamma^*) \rangle^{-1}$. We will show that regardless of the scheduling rule, all realizations (f, g) sufficiently *close* to (λ^*t, γ^*t) over the interval $[0, T_0]$ — and thus having cost $J_{T_0}(f, g)$ close to J_* — lead to an overflow at time T_0 .

Let $\|\cdot\|$ denote the L_∞ norm. For any $\epsilon > 0$, define a set of trajectories over the interval $[0, T_0]$,

$$B_\epsilon = \left\{ (f, g) \in \mathcal{D}^{2+M} : (f, g) \text{ is absolutely continuous; } \forall t \in [0, T_0], \right.$$

$$\|L_{(f)}(f'(t)) - L_{(f)}(\lambda^*)\| \leq \frac{\epsilon}{2T_0}; \quad (10)$$

$$\|L_{(g)}(g'(t)) - L_{(g)}(\gamma^*)\| \leq \frac{\epsilon}{2T_0}; \quad (11)$$

$$\forall u \in V_{g'(t)}, \inf_{v \in V_{\gamma^*}} \|u - v\| \leq \frac{\epsilon}{5T_0}; \text{ and} \quad (12)$$

$$\left. \|\lambda^* - \frac{1}{T_0} f(T_0)\| \leq \frac{\epsilon}{5T_0} \right\} \quad (13)$$

The set B_ϵ is measurable, compact, and the cost for any $(f, g) \in B_\epsilon$ satisfies $\|J_{T_0}(f, g) - J_*\| \leq \epsilon$. Next, we will show that any sequence of realizations $(q^{(n)}, f^{(n)}, g^{(n)})$, such that $q^{(n)}(0) = 0$ and $(f^{(n)}, g^{(n)})$ converge uniformly in $[0, T_0]$ to some $(f, g) \in B_\epsilon$, must have $\sum_{i \in I} b_i q_i^{(n)}(T_0) > 1 - 2\epsilon$ for all sufficiently large n . Therefore,

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log P \left(\sum_{i \in I} b_i q_i^{(n)}(T_0) > 1 - 2\epsilon \right) \geq \liminf_{n \rightarrow \infty} \frac{1}{n} \log P((\mathbf{f}^{(n)}, \mathbf{g}^{(n)}) \in B_\epsilon) \geq -J_*,$$

where the rightmost inequality follows from the Mogulsky theorem (Proposition 1).

Let $\lambda^{(n)} = \frac{1}{T_0} f^{(n)}(T_0)$, and $u^{(n)}$ denote the average service vector seen by the queues; that is to say, $u_i^{(n)} T_0$ is equal to the number of packets served from the i^{th} queue over the interval $[0, T_0]$. By (13) in the definition of B_ϵ , for all n sufficiently large,

$$\lambda^* - \left(\frac{\epsilon}{4T_0}, \frac{\epsilon}{4T_0} \right) \stackrel{(i)}{\leq} \lambda^{(n)} \stackrel{(ii)}{\leq} \lambda^* + \left(\frac{\epsilon}{4T_0}, \frac{\epsilon}{4T_0} \right), \quad (14)$$

and by (12) there exists a $v \in V_{\gamma^*}$ such that

$$v - \left(\frac{\epsilon}{4T_0}, \frac{\epsilon}{4T_0} \right) \stackrel{(i)}{\leq} u^{(n)} \stackrel{(ii)}{\leq} v + \left(\frac{\epsilon}{4T_0}, \frac{\epsilon}{4T_0} \right). \quad (15)$$

We must also have $u^{(n)} \leq \lambda^{(n)}$, since the total service cannot exceed the total arrivals. This, along with inequalities (14-ii) and (15-i), implies $v \leq \lambda^* + \left(\frac{\epsilon}{2T_0}, \frac{\epsilon}{2T_0} \right)$. Without loss of generality, we assume that the weight vector b is normalized to have $\sum_{i \in I} b_i = 2$. Then using (15-ii), we get,

$$\begin{aligned} \langle b, u^{(n)} \rangle &\leq \langle b, v \rangle + \frac{\epsilon}{2T_0} \leq \max \left(\langle b, y \rangle : y \in V_{\gamma^*}, y \leq \lambda^* + \left(\frac{\epsilon}{2T_0}, \frac{\epsilon}{2T_0} \right) \right) + \frac{\epsilon}{2T_0}, \\ &\leq \langle b, v^*(\lambda^*, \gamma^*) \rangle + \frac{\epsilon}{T_0}. \end{aligned} \quad (16)$$

Finally, using (14-i) and (16), we have that,

$$\begin{aligned} \langle b, q^{(n)}(T_0) \rangle &= \left(\langle b, \lambda^{(n)} \rangle - \langle b, u^{(n)} \rangle \right) T_0 > \left(\langle b, \lambda^* \rangle - \langle b, v^*(\lambda^*, \gamma^*) \rangle \right) T_0 - 2\epsilon, \\ &= 1 - 2\epsilon. \end{aligned}$$

In order to obtain an overflow at any time after T_0 while still incurring a cost close to J_* , the trajectory $(\lambda^* t, \gamma^* t)$ can be prepended with the zero-cost trajectory for an appropriate amount of time. ■

8 Upper bound on overflow probability under the p-Log rule

Recall the notion of a GFSP and its refined cost function from Section 6. Let J_{**} denote the lowest refined cost of a GFSP that, under p-Log rule, raises $\sum_{i \in I} b_i q_i(t)$ to 1 from the initial state $q(0) = 0$, i.e.,

$$J_{**} = \inf_{t \geq 0} J_{**,t} , \quad (17)$$

where,

$$J_{**,t} = \inf_{\psi} \left\{ \bar{J}_t | \psi : q(0) = 0, \sum_{i \in I} b_i q_i(t) \geq 1 \right\} .$$

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

Theorem 2 *For each $n = 1, 2, \dots$, consider the system under the p-Log scheduling rule in a stationary regime, then, the corresponding sequence of fluid-scaled processes is such that,*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log P \left(\sum_{i \in I} b_i q_i^{(n)}(0) \geq 1 \right) \leq -J_{**} .$$

Remark 3 This is the *equivalent* of Theorem 8.4 of [15], and its proof (given in the Appendix) follows the same framework and uses classical Wentzel-Freidlin constructions [18]. 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 of 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 $\sum_{i \in I} b_i q_i$ 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 Appendix for a 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 develop the results needed to show this.

9 Local fluid sample path

Let us first motivate the need for defining LFSP (local fluid sample path.) For each n , define the set $\mathcal{S}_m^{(n)}$ as the "fluid" scaled version of set \mathcal{S}_m of the state space partition associated with the p-Log rule, 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 $\mathcal{S}_m^{(n)}$, $m \in \{0, \dots, k, l+1, \dots, M'+1\}$ (or the corresponding switching curve) the fluid scaled queue $q^{(n)}(t)$ lies in. As $n \rightarrow \infty$, 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 sets from the partition collapse to one of the axes. Note that this is true for the partition under any radial sum-rate monotone scheduling

rule. Now consider a Lipschitz continuous limiting trajectory (q, f, g) for the fluid scaled process $(q^{(n)}, f^{(n)}, g^{(n)})$. One can show that,

$$\text{if } q(t) \in \mathcal{S}_0^{(\infty)}, \quad \text{then } \frac{d}{dt} \langle b, q(t) \rangle = \langle b, f'(t) \rangle - \max_{v \in V_{g'(t)}} \langle b, v \rangle,$$

but if $q(t)$ hits an axis, we lose information about service rates of the queues. Hence, we define a LFSP using a *finer-than-fluid* scaling such that the sets of the partition do not collapse and we are able to state the proper derivative of the limiting queue trajectory.

Consider a GFSP over some interval $[0, T]$ and fix any $\tau \in (0, T)$ such that $q(\tau) \neq 0$. Any sequence $\tau^{(n)} \rightarrow \tau$ has a subsequence along which $q^{(n)}(\tau^{(n)}) \rightarrow q(\tau)$. Set $\sigma_n = \sqrt{q_*^{(n)}(\tau^{(n)})}/\sqrt{n}$, where, $q_*^{(n)}(\cdot) = \max_{i \in I} q_i^{(n)}(\cdot)$. To obtain a local fluid sample path, we will *magnify* in both space and time the fluid scaled trajectories $(q^{(n)}, f^{(n)}, g^{(n)})$ by a factor of σ_n^{-1} , i.e., an order $O(\sqrt{n})$ term. More formally, for any fixed $S > 0$, the following re-scaled functions over the interval $[\tau^{(n)}, \tau^{(n)} + \sigma_n S]$, parameterized by $s \in [0, S]$, are called the local fluid sample paths⁴: for all $i \in I$ and $m \in \mathcal{M}$,

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

Then along some subsequence in n ,

- the functions $(\diamond q^{(n)}, \diamond f^{(n)}, \diamond g^{(n)})$ converge uniformly over $[0, S]$ to Lipschitz continuous functions $(\diamond q, \diamond f, \diamond g)$;
- for each $i \in I$, the function $\diamond \hat{q}_i^{(n)}$ either converges uniformly over $[0, S]$ to a finite Lipschitz continuous function $\diamond \hat{q}_i$, or is identically equal to ∞ ;
- the function $\diamond d^{(n)}$ converges uniformly over $[0, S]$ 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 $\diamond \hat{q}_i(\cdot) = \infty$ 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 the set $\{x \in \mathbb{R}^2 | x \geq -\diamond \hat{q}(0)\}$, which is at least a half-plane. Lastly, we have the following relation between the *cost* of LFSP over $[0, S]$ and the refined cost sequence of GFSP over $[\tau^{(n)}, \tau^{(n)} + \sigma_n S]$ (see (9.4) of [15]),

$$J_S(\diamond f, \diamond g) - J_0(\diamond f, \diamond g) \leq \liminf_{n \rightarrow \infty} \frac{1}{\sigma_n} (\bar{J}_{\tau^{(n)} + \sigma_n S}^{(n)} - \bar{J}_{\tau^{(n)}}^{(n)}). \quad (18)$$

⁴The definitions of $\diamond f_i^{(n)}(\cdot)$ and $\diamond g_m^{(n)}(\cdot)$ are the same as in [15], whereas, $\diamond q_i^{(n)}(\cdot)$ are *scaled* as in [15] but *centered* differently.

9.1 Scheduling over time scales of LFSP

By (6) the scheduling decision in the interval $[\tau^{(n)}, \tau^{(n)} + \sigma_n S]$ depends on the slope,

$$\frac{h_2(Q(n\tau^{(n)} + n\sigma_n s))}{h_1(Q(n\tau^{(n)} + n\sigma_n s))} ,$$

and the sign of $\diamond d^{(n)}(s)$, recall the tie-breaking rule mentioned in the description of p-Log rule in Section 4. Without loss of generality suppose $q_1(\tau) \geq q_2(\tau)$. Then by (5), for n large enough, we have that,

$$\begin{aligned} \frac{h_2(Q(n\tau^{(n)} + n\sigma_n s))}{h_1(Q(n\tau^{(n)} + n\sigma_n s))} &= \frac{b_2}{b_1} \min \left(1, \frac{Q_2(n\tau^{(n)} + n\sigma_n s)}{\sqrt{Q_1(n\tau^{(n)} + n\sigma_n s)}} \right) , \\ &= \frac{b_2}{b_1} \min \left(1, \frac{\diamond \hat{q}_2^{(n)}(s) \sqrt{q_1^{(n)}(\tau^{(n)})}}{\sqrt{q_1^{(n)}(\tau^{(n)} + \sigma_n s)}} \right) . \end{aligned}$$

Then, as $n \rightarrow \infty$, the above converges to,

$$\frac{b_2}{b_1} \min (1, \diamond q_2(s) + \diamond \hat{q}_2(0)) ,$$

where the convergence is uniform on $[0, S]$. Let us define a vector field $\diamond h$ over the state space of $\diamond q$, i.e., $\{x \in \mathbb{R}^2 | x \geq -\diamond \hat{q}(0)\}$, as follows:

$$\begin{aligned} \diamond h_1(x) &= 1 , \\ \diamond h_2(x) &= \frac{b_2}{b_1} \min (1, x_2 + \diamond \hat{q}_2(0)) . \end{aligned}$$

That is, we can restate the above convergence as,

$$\frac{h_2(Q(n\tau^{(n)} + n\sigma_n s))}{h_1(Q(n\tau^{(n)} + n\sigma_n s))} \rightarrow \frac{\diamond h_2(\diamond q(s))}{\diamond h_1(\diamond q(s))}$$

uniformly on $[0, S]$. Hence, the switching curves on the space of $\diamond q$ are given by,

$$\left\{ (x_1, x_2) \in \mathbb{R}^2 : \frac{\diamond h_2(x)}{\diamond h_1(x)} = r_m \right\} = \left\{ (x_1, x_2) \in \mathbb{R}^2 : x_2 = -\diamond \hat{q}_2(0) + \frac{b_1}{b_2} r_m \right\}$$

for $m \in \{0, \dots, k\}$, and are now parallel to x_1 axis (i.e., the axis of $\diamond q_1(s)$, see Fig. 4). For each $m \in \{1, \dots, k\}$, define the set $\diamond \mathcal{S}_m$ by appropriately re-scaling the set \mathcal{S}_m , i.e.,

$$\begin{aligned} \diamond \mathcal{S}_m &= \left\{ (x_1, x_2) \in \mathbb{R}^2 : r_{m-1} < \frac{\diamond h_2(x)}{\diamond h_1(x)} < r_m \right\} , \\ &= \left\{ (x_1, x_2) \in \mathbb{R}^2 : -\diamond \hat{q}_2(0) + \frac{b_1}{b_2} r_{m-1} < x_2 < -\diamond \hat{q}_2(0) + \frac{b_1}{b_2} r_m \right\} , \end{aligned} \tag{19}$$




Figure 4: Partitions and switching curves on space of local fluid queue $\diamond q$.

and define,

$$\begin{aligned} \diamond \mathcal{S}_0 &= \left\{ (x_1, x_2) \in \mathbb{R}^2 : r_k < \frac{\diamond h_2(x)}{\diamond h_1(x)} \right\}, \\ &= \left\{ (x_1, x_2) \in \mathbb{R}^2 : -\diamond \hat{q}_2(0) + \frac{b_1}{b_2} r_k < x_2 \right\}. \end{aligned} \quad (20)$$

See Fig. 4 for a graphical illustration of sets $\diamond \mathcal{S}_0, \dots, \diamond \mathcal{S}_k$. The service rate $\mu(s)$ allocated to the local fluid queue $\diamond q(s)$ at time s , depends on which re-scaled set $\diamond \mathcal{S}_m$ (or the associated switching curve) the local fluid queue $\diamond q(s)$ lies in. More formally, the following lemma relates the service rate $\mu(s)$ to the vector field $\diamond h(\diamond q(s))$, and can be derived without much effort from the results shown for the Exponential rule scheduler in [4]⁵.

Lemma 1 *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,

$$\mu(s) \in \arg \max_{v \in V_{\gamma(s)}} \langle \diamond h(\diamond q(s)), v \rangle. \quad (21)$$

Remark 4 Note that if $\diamond q_2(s) \geq -\diamond \hat{q}_2(0) + 1$, then $\diamond h(\diamond q(s)) \propto b$, and the argmax in (21) may not be unique. However, $\mu(s)$ can still be uniquely identified as long as $\diamond d(s) \neq 0$. That is, if $\diamond d(s) > 0$, then $\mu(s)$ is such that $\mu_1(s)$ is the largest possible among all points achieving max in (21); similarly, if $\diamond d(s) < 0$, then $\mu(s)$ is such that $\mu_2(s)$ is the largest. The argmax in (21) may again be non-unique if $\diamond q(s)$ lies on a switching curve, i.e., $\diamond q_2(s) = -\diamond \hat{q}_2(0) + \frac{b_1}{b_2} r_m$ for some $m \in \{0, 1, \dots, r_k\}$, however, in this case $\mu(s)$ can not be uniquely identified using only the components of LFSP, nor will we need to uniquely identify $\mu(s)$.

⁵Specifically, Proposition 1 and results leading from (34) to (35) of [4] relate the service rate seen by the local fluid queue under Exponential rule scheduler to the vector field defining the scheduler.

We will also need the following two crucial lemmas that show for the p-Log rule what Lemmas 9.2 and 9.3 of [15] show for the Exponential rule. The following two lemmas also implicitly prove the throughput-optimality of p-Log rule, see Appendix for details.

Lemma 2 *There exist fixed constants $\epsilon_1 > 0$ and $\delta_1 > 0$, and a Lipschitz continuous Lyapunov function H (constructed in the proof below) such that for any regular point $s \in [0, S]$, if*

$$\frac{d}{ds} J_s(\diamond f, \diamond g) \leq \epsilon_1, \quad \text{then} \quad \frac{d}{ds} H(\diamond q(s)) \leq -\delta_1.$$

Proof If $\frac{d}{ds} J_s(\diamond f, \diamond g)$ is small, then $\lambda(s)$ must be close to the mean arrival rate $\bar{\lambda}$, and $\gamma(s)$ close to the server state distribution π . Since the capacity region V_γ is continuous in the distribution γ (see (1)), there exist vectors $\lambda^* < v^*$ such that uniformly on all sufficiently small values of ϵ_1 , we have $\lambda(s) < \lambda^*$, and v^* lies in the interior of $V_{\gamma(s)}$. Let $\delta_1 = \min_{i \in I} (v_i^* - \lambda_i^*)$. By Lemma 1, we have,

$$\begin{aligned} \langle \diamond h(\diamond q(s)), \mu(s) \rangle &= \arg \max_{v \in V_{\gamma(s)}} \langle \diamond h(\diamond q(s)), v \rangle, \\ &> \langle \diamond h(\diamond q(s)), v^* \rangle. \end{aligned} \quad (22)$$

Note that continuous and bounded vector field $\diamond h$ is in fact a gradient field associated with a convex, increasing in each direction, and Lipschitz continuous function H , i.e., $\nabla H = \diamond \hat{h}$ on $\{x \in \mathbb{R}^2 | x \geq -\diamond \hat{q}(0)\}$. Then we have,

$$\begin{aligned} \frac{d}{ds} H(\diamond q(s)) &= \langle \diamond h(\diamond q(s)), \diamond q'(s) \rangle, \\ &= \langle \diamond h(\diamond q(s)), \lambda(s) - \mu(s) \rangle, \\ &\leq \langle \diamond h(\diamond q(s)), \lambda^* - v^* \rangle, \\ &\leq -\delta_1, \end{aligned}$$

where the first inequality follows from (22) and the second from the definition of δ_1 . ■

Lemma 3 *There exist fixed constants $\epsilon_2 > 0$ and $\hat{\delta}_2 > \delta_2 > 0$ such that, if*

$$J_S(\diamond f, \diamond g) - J_0(\diamond f, \diamond g) \leq \epsilon_2 S, \quad \text{then} \quad H(\diamond q(S)) - H(\diamond q(0)) \leq -\hat{\delta}_2 S,$$

which further implies that uniformly for all large S , we have the following bound on the change in the weighted sum queue,

$$\langle b, \diamond q(S) - \diamond q(0) \rangle \leq -\delta_2 S.$$

Proof Noting that $H(\diamond q(s))$ (as a function of s) is Lipschitz continuous with some Lipschitz constant denoted by c , the proof of first statement is identical to that of Lemma 9.3 of [15]: pick a positive $\epsilon_2 < \epsilon_1$, let $B_1 = \{s \in [0, S] : \frac{d}{ds} J_s(\diamond f, \diamond g) \geq \epsilon_1\}$ and $B_2 = S \setminus B_1$. Then the Lebesgue measures of B_1 and B_2 satisfy, $\nu(B_1) \leq \frac{\epsilon_2}{\epsilon_1} S$ and $\nu(B_2) \geq (1 - \frac{\epsilon_2}{\epsilon_1}) S$. Finally,

$$\begin{aligned} H(\diamond q(S)) - H(\diamond q(0)) &= \int_{B_1} \frac{d}{ds} H(\diamond q(s)) + \int_{B_2} \frac{d}{ds} H(\diamond q(s)) \leq c \frac{\epsilon_2}{\epsilon_1} S - \delta_1 (1 - \frac{\epsilon_2}{\epsilon_1}) S, \\ &= -S (\delta_1 - \frac{\epsilon_2}{\epsilon_1} (c + \delta_1)). \end{aligned}$$

Fix any positive $\hat{\delta}_2 < \delta_1$, then a sufficiently small ϵ_2 can be chosen to prove the first statement. To prove the last statement, note that, for any $x \geq y$,

$$\left\langle \left(1, \frac{b_2}{b_1}\right), x - y \right\rangle - \frac{b_2}{b_1} \leq H(x) - H(y) \leq \left\langle \left(1, \frac{b_2}{b_1}\right), x - y \right\rangle ,$$

Using this we get,

$$\begin{aligned} -\hat{\delta}_2 S &\geq H(\diamond q(S)) - H(\diamond q(0)) , \\ &\geq \left\langle \left(1, \frac{b_2}{b_1}\right), \diamond q(S) - \diamond q(0) \right\rangle - b_2 , \\ \Rightarrow -b_1 S \left(\hat{\delta}_2 - \frac{b_2}{S} \right) &\geq \langle b, \diamond q(S) - \diamond q(0) \rangle . \end{aligned}$$

Now one can take S large enough to obtain a δ_2 . ■

In the following lemma, identical to the result in Section 11 of [15], we show that any GFSP that raises the weighted sum queue to unity contains a LFSP with cost *close* to that of the GFSP. Subsequently, we will use this lemma on a GFSP of cost close to J_{**} (see (17)) in order to obtain a LFSP with cost close to J_{**} .

Lemma 4 *Suppose a GFSP ψ is given that satisfies $q(0) = 0$, $\langle b, q(T) \rangle = 1$ for some $T > 0$ and has a cost $\bar{J}_T < \infty$. Then, for an arbitrarily small $\epsilon > 0$, an LFSP $(\diamond q, \diamond \hat{q}, \diamond d, \diamond f, \diamond g)$ over an arbitrarily large interval $[0, S]$ can be constructed from the elements of ψ , such that,*

$$\langle b, \diamond q(S) - \diamond q(0) \rangle \geq \theta S , \quad (23)$$

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

$$\frac{J_S(\diamond f, \diamond g) - J_0(\diamond f, \diamond g)}{\langle b, \diamond q(S) - \diamond q(0) \rangle} \leq \bar{J}_T + \epsilon . \quad (24)$$

Proof Recall that $(\bar{J}_T - \bar{J}_0) / \langle b, q(T) - q(0) \rangle = \bar{J}_T$. For any $0 < \xi_1 < \xi_2 < 1$, define time $t_1 = \max(t : \langle b, q(t) \rangle = \xi_1)$ and time $t_2 = \min(t > t_1 : \langle b, q(t) \rangle = \xi_2)$. Then for any $\epsilon > 0$, there must exist $0 < \xi_1 < \xi_2 < 1$ such that,

$$\frac{\bar{J}_{t_2} - \bar{J}_{t_1}}{\langle b, q(t_2) - q(t_1) \rangle} < \bar{J}_T + \frac{\epsilon}{2} . \quad (25)$$

Fix an $S > 0$ large enough as required by Lemma 3. Then there exists a sequence $\{\tau^{(n)}\}$ within $[t_1, t_2]$ such that for all n large enough,

$$\begin{aligned} \left\langle b, q^{(n)}(\tau^{(n)} + \sigma_n S) - q^{(n)}(\tau^{(n)}) \right\rangle &> 0 , \\ \frac{\bar{J}_{(\tau^{(n)} + \sigma_n S)}^{(n)} - \bar{J}_{\tau^{(n)}}^{(n)}}{\left\langle b, q^{(n)}(t_2^{(n)}) - q^{(n)}(t_1^{(n)}) \right\rangle} &< \bar{J}_T + \epsilon , \end{aligned}$$

where, as before, $\sigma_n = \sqrt{q_*^{(n)}(\tau^{(n)})}/\sqrt{n}$; such a sequence must exist, otherwise (25) cannot hold. The above two inequalities can be re-written as,

$$\langle b, \diamond q^{(n)}(S) - \diamond q^{(n)}(0) \rangle > 0, \quad (26)$$

$$\frac{\sigma_n^{-1}(\bar{J}_{(\tau^{(n)} + \sigma_n S)}^{(n)} - \bar{J}_{\tau^{(n)}}^{(n)})}{\langle b, \diamond q^{(n)}(S) - \diamond q^{(n)}(0) \rangle} < \bar{J}_T + \epsilon. \quad (27)$$

Now we can pick a subsequence in n along which $\tau^{(n)} \rightarrow \tau$ and $(\diamond q^{(n)}, \diamond \hat{q}^{(n)}, \diamond d^{(n)}, \diamond f^{(n)}, \diamond g^{(n)})$ converge to $(\diamond q, \diamond \hat{q}, \diamond d, \diamond f, \diamond g)$, as described in Section 9. From (26) we also have that $\langle b, \diamond q(S) - \diamond q(0) \rangle \geq 0$. Then by Lemma 3, we get,

$$J_S(\diamond f, \diamond g) - J_0(\diamond f, \diamond g) > \epsilon_2 S,$$

which, alongside (27) and (18), further implies that for some $\theta > 0$,

$$\langle b, \diamond q(S) - \diamond q(0) \rangle \geq \theta S,$$

and finally,

$$\frac{J_S(\diamond f, \diamond g) - J_0(\diamond f, \diamond g)}{\langle b, \diamond q(S) - \langle b, \diamond q(0) \rangle} \leq \bar{J}_T + \epsilon.$$

■

10 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 will 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 (17), there exists a GFSP ψ satisfying $q(0) = 0$, $\langle b, q(T) \rangle = 1$ for some finite $T > 0$, and having a cost $\bar{J}_T < J_*$. Then by Lemma 4, we can construct an LFSP $(\diamond q, \diamond \hat{q}, \diamond d, \diamond f, \diamond g)$ satisfying (23) and (24) for an ϵ small enough so that $J_{***} \equiv \bar{J}_T + \epsilon < J_*$.

Since $\diamond \hat{q}_i(\cdot) = \infty$ for at least one $i \in I$, without loss of generality suppose $\diamond \hat{q}_1(\cdot) = \infty$, thus all switching curves on the space of $\diamond q$ are parallel to $\diamond q_1$ axis. Recall that the lower boundary of the set $\diamond \mathcal{S}_0$ is given by the switching curve $x_2 = -\diamond \hat{q}_2(0) + \frac{b_1}{b_2} r_k$ (see Fig. 4). Let S_1 and S_2 respectively be the first and the last time in $[0, S]$ such that the trajectory $\diamond q_2(s) \leq -\diamond \hat{q}_2(0) + \frac{b_1}{b_2} r_k$, with $S_1 = S_2 = S$ if $\diamond q_2(s)$ never hits $[-\diamond \hat{q}_2(0), -\diamond \hat{q}_2(0) + \frac{b_1}{b_2} r_k]$. Note that the trajectory of $\diamond q(s)$ lies in $\diamond \mathcal{S}_0$ in $(0, S_1)$ and (S_2, S) . Then one of the following must be true over the interval $[0, S_1]$ (similarly $[S_2, S]$):

- (a) $\langle b, \diamond q(0) \rangle < \langle b, \diamond q(S_1) \rangle$ and the cost per unit increase in weighted-sum-queue over the interval $[0, S_1]$ is less than J_{***} , i.e.

$$\frac{J_{S_1}(\diamond f, \diamond g) - J_0(\diamond f, \diamond g)}{\langle b, \diamond q(S_1) - \diamond q(0) \rangle} \leq J_{***}.$$

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

$$\frac{J_{S_1}(\diamond f, \diamond g) - J_0(\diamond f, \diamond g)}{\langle b, \diamond q(S_1) - \diamond q(0) \rangle} > J_{***} .$$

(c) $\langle b, q(0) \rangle \geq \langle b, q(S_1) \rangle$.

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

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

By Lemma 1 and the fact that $(\diamond q(s) : 0 < s < S_1)$ lies in $\diamond S_0$, we have $\mu(s) \in \arg \max_{v \in V_{\gamma(s)}} \langle b, v \rangle$. This and the *linearity* of $\mu(s)$ in $\gamma(s)$ (see (1)) implies $\hat{\mu} \in \arg \max_{v \in V_{\hat{\gamma}}} \langle b, v \rangle$. Then,

$$\begin{aligned} \langle b, \diamond q(S_1) - q(0) \rangle &= \langle b, \hat{\lambda} - \hat{\mu} \rangle S_1 , \\ &\leq \langle b, \hat{\lambda} - v^*(\hat{\lambda}, \hat{\gamma}) \rangle S_1 , \end{aligned} \tag{28}$$

where $v^*(\hat{\lambda}, \hat{\gamma})$ is as defined in (8). Finally,

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

where the first inequality follows from the assumption that (a) is true, the second from convexity of rate functions, and the last one from (28). However, by the definition of J_* in (9), the right side of last inequality cannot be less than J_* , giving the contradiction we needed; therefore, we must have $J_* = J_{**}$.

Now, if (a) is not true for both intervals $[0, S_1]$ and $[S_2, S]$, then for (23) and (24) to hold, we must have,

$$\frac{J_{S_2}(\diamond f, \diamond g) - J_{S_1}(\diamond f, \diamond g)}{\langle b, \diamond q(S_2) - \diamond q(S_1) \rangle} \leq J_{***} < J_* , \tag{29}$$

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

$$\begin{aligned} \langle b, \diamond q(S_2) - \diamond q(S_1) \rangle &\geq \theta_1 \langle b, \diamond q(S) - \diamond q(0) \rangle , \\ &\geq \theta_1 \theta S , \end{aligned} \tag{30}$$

where the second inequality follows by applying (23) of Lemma 4. Moreover, since $\langle b, \diamond q(\cdot) \rangle$ is Lipschitz, for some fixed $\theta_2 > 0$,

$$S_2 - S_1 \geq \theta_2 \langle b, \diamond q(S_2) - \diamond q(S_1) \rangle \geq \theta_2 \theta_1 \theta S .$$




Figure 5: Illustration of original trajectory $(\diamond q(s), s \in [S_1, S_2])$ and the extension $(\diamond q(s), s \in [S_2, S'_2])$ in order to obtain $\diamond q_2(S'_2) = \diamond q_2(S_1)$.

This, along with (30) and Lemma 3, imply,

$$\begin{aligned} J_{S_2}(\diamond f, \diamond g) - J_{S_1}(\diamond f, \diamond g) &\geq \epsilon_2(S_2 - S_1) , \\ &\geq \epsilon_2 \theta_2 \theta_1 \theta S . \end{aligned} \quad (31)$$

Subsequently, we will use (30) and (32) to make the quantities on the left side of these inequalities as large as needed by choosing a large S .

Consider the trajectory $(\diamond q(s), s \in [S_1, S_2])$ which is an element of the LFSP $(\diamond q, \diamond \hat{q}, \diamond d, \diamond f, \diamond g)$ over $[S_1, S_2]$. For some $S'_2 \geq S_2$, we can append to this LFSP an *extension* over the time interval $[S_2, S'_2]$ so as to obtain $\diamond q_2(S'_2) = \diamond q_2(S_1)$ (see Fig. 5). Moreover, since the terminal values $\diamond q_2(S_1)$ and $\diamond q_2(S_2)$, lie within the bounded interval $[-\diamond \hat{q}_2(0), -\diamond \hat{q}_2(0) + \frac{b_1}{b_2} r_k]$, therefore the extension LFSP can be constructed such that it has a bounded cost, i.e.,

$$J_{S'_2}(\diamond f, \diamond g) - J_{S_2}(\diamond f, \diamond g) \leq \Delta , \quad (32)$$

and a bounded increase in weighted-sum-queue, i.e.,

$$|\langle b, \diamond q(S'_2) - \diamond q(S_2) \rangle| \leq \Delta , \quad (33)$$

for some large fixed $\Delta < \infty$ which is independent of any components of the LFSP over $[0, S]$, (e.g. the value of S or the terminal values $\diamond q_2(S_1)$ and $\diamond q_2(S_2)$.) Finally, the constant S can be chosen large enough such that, by (30) and (33), we have,

$$\langle b, \diamond q(S'_2) - \diamond q(S_1) \rangle > 0 ,$$

and by (29)–(33), for some small $\epsilon_4 > 0$ that satisfies $J_{***} + \epsilon_4 < J_*$, we have,

$$\frac{J_{S'_2}(\diamond f, \diamond g) - J_{S_1}(\diamond f, \diamond g)}{\langle b, \diamond q(S'_2) \rangle - \langle b, \diamond q(S_1) \rangle} \leq J_{***} + \epsilon_4 < J_* . \quad (34)$$

Set $c_0 = -\diamond \hat{q}_2(0)$ and for each $m \in \{1, \dots, k\}$, choose a $c_m \in \left(-\diamond \hat{q}_2(s) + \frac{b_1}{b_2} r_{m-1}, -\diamond \hat{q}_2(s) + \frac{b_1}{b_2} r_m\right)$ such that the counting measure of set $\{s \in [0, S'_2] : \diamond q(s) = c_m\}$ is finite; such $\{c_m\}$ exist since $\diamond q$ is

Lipschitz. Lastly, choose a $c_{k+1} < \infty$ large enough such that $\max_{s \in [S_1, S'_2]} \diamond q_2(s) < c_{k+1}$. For each $m \in \{1, \dots, k+1\}$, define a set $C_m = \{s \in [S_1, S'_2] : \diamond q_2(s) \in [c_{m-1}, c_m]\}$. We make the following three observations which will be used in the sequel: for each $m \in \{1, \dots, k+1\}$,

- (i) the trajectory $(\diamond q_2(s), s \in [S_1, S'_2])$ intersects with end points of interval $[c_{m-1}, c_m]$ only finitely many times, therefore, the corresponding set C_m can be written as a union of finitely many intervals;
- (ii) if set C_m is non-empty, then $\diamond q_2(\min_{s \in C_m} s) = \diamond q_2(\max_{s \in C_m} s)$;
- (iii) the trajectory $(\diamond q(s), s \in C_m)$ can intersect with at most one switching curve $x_2 = -\diamond \hat{q}_2(s) + \frac{b_1}{b_2} r_{m-1}$, or equivalently, with at most two adjacent regions $\diamond \mathcal{S}_{(\cdot)}$, (e.g., for $m = k+1$ the above trajectory can intersect with adjacent regions $\diamond \mathcal{S}_k$ and $\diamond \mathcal{S}_0$).

Then using (i) and (ii) above, for all sets C_m , we must have,

$$\int_{C_m} \diamond q'_2(s) ds = 0. \quad (35)$$

Moreover, there exists a set C_m such that,

$$\int_{C_m} \langle b, \diamond q'(s) \rangle ds > 0,$$

and,

$$\frac{\int_{C_m} \left(L_{(f)}(\diamond f'(s)) + L_{(g)}(\diamond g'(s)) \right) ds}{\int_{C_m} \langle b, \diamond q'(s) \rangle ds} < J_*,$$

otherwise (34) will not hold. With a set C_m for which the above two hold, let,

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

By (iii) above, (19–21), and the fact that $\diamond d(s) > 0$ in $[S_1, S'_2]$, the service vector $\hat{\mu}$ is a convex combination of at most two adjacent vertices of capacity region $V_{\hat{\gamma}}$ and lies on the *facet* with outer normal slope r_{m-1} . That is, $\hat{\mu}$ is necessarily a maximal element of $V_{\hat{\gamma}}$. This, together with the fact that $\hat{\lambda}_2 = \hat{\mu}_2$ which follows from (35), gives $\hat{\mu} = \arg \min_{v \in V_{\hat{\gamma}}} \langle b, (\hat{\lambda} - v)^+ \rangle = v^*(\hat{\lambda}, \hat{\gamma})$, and then,

$$\int_{C_m} \langle b, \diamond q'(s) \rangle ds = \langle b, \hat{\lambda} - v^*(\hat{\lambda}, \hat{\gamma}) \rangle \nu(C_m).$$

Finally,

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

By the definition of J_* , the right side of last inequality cannot be less than J_* , giving the required contradiction; hence we conclude that $J_* = J_{**}$. \blacksquare

11 Conclusion and extensions

In order to minimize the asymptotic probability of weighted-sum-queue overflow, the desirable mode of overflow is one where queues may build up at different rates, however, the total weighted service rate seen by the queues is the highest possible, (service rate subject to being not more than the arrival rate.) The p-Log scheduling rule minimizes the asymptotic probability of weighted-sum-queue overflow and exhibits such mode of overflow. This property of the p-Log rule is related to the collapse under fluid scaling of all but one set of the state space partition to either of the axes; the set that does not collapse is the horn-shaped weighted-max-sum-rate set. This collapse under fluid scaling is typical of the partition under any radial sum-rate monotone scheduler. However, the scaling or magnifying factor required to obtain a useful LFSP, and the shape of the sets of partition on the local fluid state space will vary for different radial sum-rate monotone schedulers. In this regard, p-Log rule yields easily tractable partition of the local fluid space where all switching curves are parallel to one of the axes.

The following extensions of the main results of the paper as well as the system model are possible without much effort.

First, the lower bound (i.e., Theorem 1-i and its proof) goes through without any changes for any fixed number of queues (instead of only two) sharing the time-varying server.

Second, the main result (Theorem 1) is also applicable to the following different and *simpler* system model. Instead of a single server with time-varying state $\mathbf{m}(t) \in \{1, 2, \dots, M\}$ having distribution π , suppose there are M distinct servers with fixed but asymmetric capacities across the two queues. More specifically, the m^{th} server, if allocated to the $i^{\text{th}} \in I$ queue, can serve $\pi_i \mu_i^m \in \mathbb{Z}_+$ packets from the queue. Then the scheduling problem is to dynamically assign the servers to the queues based on the queue state, and the total service offered to the queue is taken to be the sum of service offered by each server assigned to the queue. When queue is in state Q , the p-Log scheduler in this context allocates the m^{th} server to the queue $i_{pLog}^*(Q, m)$ as given by (6). This system model is simpler in that the only random process now driving the system are the arrivals, but is also different from the original model in that there are multiple servers with each server having asymmetric capacities across the two queues.

Third, the main result also goes through if the capacity regions V^m are permitted to be arbitrary convex polyhedra instead of just triangles obtained as a convex hull of service vectors $(0, 0)$, $(\mu_1^m, 0)$, and $(0, \mu_2^m)$. In a more general model, in any state m , the server can be permitted to operate at any one of the k_m service vectors from the set,

$$R_m = \{(\mu_1^m(1), \mu_2^m(1)), \dots, (\mu_1^m(k_m), \mu_2^m(k_m))\} .$$

The only change needed is to generalize the definition of pLog rule as follows. When the system is in state (Q, m) , operate the server at a service vector $v_{pLog}(Q, m) \in V^m$ given by,

$$\mu_{pLog}^*(Q, m) \in \arg \max_{y \in R^m} \langle y, h(Q) \rangle ,$$

where, in the case of a tie, if $Q_1 \geq Q_2$, then $v_{pLog}(Q, m)$ maps to the maximizer with the largest capacity for Queue 1, otherwise $v_{pLog}(Q, m)$ maps to the maximizer with the largest capacity for

Queue 2. This generalization affects the fluid and local fluid sample paths through (21) in Lemma 1, which can be shown to hold exactly as in [19, 20].

Besides the above extensions, we will conclude by stating one more interesting application of the p-Log scheduler. A throughput-optimal scheduler can also be used to offer minimum and maximum average service rate guarantees to *infinitely backlogged* queues, referred to as tasks, sharing a time-varying server/wireless channel [21, 22]. This is done by using virtual token queues that are fed by deterministic arrivals at a constant rate λ_i , and making scheduling decisions to serve tasks based on the virtual token queues (augmented with a scheduling rule to use when all token queues are empty). If rates λ_i are feasible (i.e., vector λ lies in the interior of capacity region V_π associated with the time-varying server), then under any throughput-optimal scheduler, each task i will be offered an average service rate $v_i \geq \lambda_i$ (such that $v \in V_\pi$). However, if rates λ_i are not feasible, then main result of this paper implies that the average service rate vector v has the following interesting and desirable property under pLog rule: $\langle b, v \rangle$ is maximized subject to $v_i \leq \lambda_i$. That is, pLog rule splits the tasks in two sets, for one set of tasks $v_i = \lambda_i$, whereas for the other $v_i < \lambda_i$, and the sets are chosen such that the total weighted service rate $\langle b, v \rangle$ is maximized.

References

- [1] R. Knopp and P. Humblet, "Information capacity and power control in single-cell multiuser communications," in *Proc. IEEE Int. Conf. on Comm.*, vol. 1, Jun 1995, pp. 331–335.
- [2] X. Liu, E. K. P. Chong, and N. B. Shroff, "A framework for opportunistic scheduling in wireless networks," *Comput. Networks*, vol. 41, no. 4, pp. 451–474, 2003.
- [3] 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.
- [4] 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.
- [5] B. Sadiq, S. Baek, and G. de Veciana, "Delay-optimal opportunistic scheduling and approximations: the Log rule," in *Proc. IEEE INFOCOM*, Rio de Janeiro, Brazil, 2009.
- [6] A. Eryilmaz, R. Srikant, and J. Perkins, "Stable scheduling policies for fading wireless channels," *IEEE/ACM Trans. Networking.*, vol. 13, no. 2, pp. 411–424, April 2005.
- [7] C. Zhou and G. Wunder, "General stability conditions in wireless broadcast channels," in *Proc. of Allerton Conf. on Comm. Control and Comp.*, Sep 2008.
- [8] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Trans. Information Theory*, vol. 39, no. 2, pp. 466–478, Mar 1993.
- [9] E. M. Yeh and A. S. Cohen, "Delay optimal rate allocation in multiaccess fading communications," in *Proc. of Allerton Conf. on Comm. Control and Comp.*, 2004, pp. 140–149.
- [10] B. Hajek, "Optimal control of two interacting service stations," *Automatic Control, IEEE Transactions on*, vol. 29, no. 6, pp. 491–499, Jun 1984.

- [11] V. Subramanian, “Large deviations of max-weight scheduling policies on convex rate regions,” in *Proc. of IEEE Information Theory Workshop*, Feb. 2008, pp. 535–544.
- [12] V. Subramanian, T. Javidi, and S. Kittipiyakul, “Many sources large deviations of max-weight scheduling.” *submitted*.
- [13] 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.
- [14] V. J. Venkataramanan and X. Lin, “On wireless scheduling algorithms for minimizing the queue-overflow probability,” *submitted*.
- [15] 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.
- [16] S. Borst, “User-level performance of channel-aware scheduling algorithms in wireless data networks,” *IEEE/ACM Trans. Networking.*, vol. 13, no. 3, pp. 636–647, 2005.
- [17] A. Dembo and O. Zeitouni, *Large deviations techniques and applications*, 2nd ed. Springer, 1998.
- [18] M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems*, 2nd ed. Springer, 1998.
- [19] A. L. Stolyar, “Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic,” *The Annals of Applied Probability*, vol. 14, no. 1, pp. 1–53, 2004.
- [20] M. Airy, “Multiple antenna wireless systems: Capacity and user performance limits,” Ph.D. dissertation, The University of Texas at Austin, May 2006.
- [21] S. Shakkottai and A. Stolyar, “Scheduling algorithms for a mixture of real-time and non-real-time data in HDR,” in *Proceedings of the 17th International Teletraffic Congress (ITC-17)*, 2001.
- [22] M. Andrews, L. Qian, and A. Stolyar, “Optimal utility based multi-user throughput allocation subject to throughput constraints,” in *Proc. IEEE INFOCOM*, vol. 4, no. 13-17, Mar 2005, pp. 2415–2424.
- [23] V. A. Malyshev and M. V. Menshikov, “Ergodicity, continuity, and analyticity of countable markov chains,” *Transactions of Moscow Mathematical Society*, 1979.

12 Appendix: proof of Theorem 2

We begin by developing the necessary results need for the proof of Theorem 2. The corresponding proof in [15] has some parts– lemmas and theorems– that are not specific to the Exp rule; these mostly go through by interpreting $q_*(t)$ (which is the notation for max-queue in [15]) as weighted-sum-queue $\sum_{i \in I} b_i q_i(t)$, while other lemmas and theorems require more work specific to the p-Log rule. We proceed by stating the following two theorems whose proofs are short and identical to those of Theorem 8.5 and Theorem 8.6 in [15] by interpreting $q_*(t)$ as $\sum_{i \in I} b_i q_i(t)$.

Theorem 3 (See Theorem 8.5 of [15]) For any fixed $T \geq 0$ and $0 \leq c < 1$, let us denote

$$J_{**, \leq T, c} = \inf \{ \bar{J}_t | \psi : \sum_{i \in I} b_i q_i(0) \leq c, \sum_{i \in I} b_i q_i(t) \geq 1 \text{ for some } t \leq T \} .$$

Then, we have,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \sup_{\sum_{i \in I} b_i q_i^{(n)}(0) \leq c} P \left(\sup_{t \in [0, T]} \sum_{i \in I} b_i q_i^{(n)}(t) > 1 \right) \leq -J_{**, \leq T, c} ,$$

and as $c \rightarrow 0$,

$$J_{**, \leq T, c} \nearrow J_{**, \leq T, 0} = \inf_{t \leq T} J_{**, t} . \quad (36)$$

Theorem 4 (See Theorem 8.6 of [15]) For any fixed $C^* > \delta > 0$, and $T > 0$, let us denote

$$K(C^*, \delta, T) = \inf \{ \bar{J}_T | \psi : \sum_{i \in I} b_i q_i(0) \leq C^*, \sum_{i \in I} b_i q_i(t) \geq \delta \text{ for all } t \in [0, T] \} .$$

Then, we have,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \sup_{\sum_{i \in I} b_i q_i^{(n)}(0) \leq C^*} P \left(\inf_{t \in [0, T]} \sum_{i \in I} b_i q_i^{(n)}(t) \geq \delta \right) \leq -K(C^*, \delta, T) . \quad (37)$$

Theorem 5 (See Theorem 8.7 of [15]) For any $C^* > 0$, there exists $\Delta_1 > 0$ such that for all sufficiently large T and all $\delta \in (0, C^*)$, we have $K(C^*, \delta, T) \geq \Delta_1 T$.

Proof It is clear that $K(C^*, \delta, T)$ is increasing in δ . We will show that there exists an $\epsilon_5 > 0$ and $\delta_5 > 0$ such that for any GFSP satisfying $\langle b, q(\cdot) \rangle > 0$ over $[0, T]$, if

$$\bar{J}_T - \bar{J}_0 \leq \epsilon_5 T, \quad \text{then} \quad \langle b, q(T) - q(0) \rangle \leq -\delta_5 T , \quad (38)$$

hence for all $T > C^*/\delta_5$ and all $\delta \in (0, C^*)$, we must have $K(C^*, \delta, T) > \epsilon_5 T$, thus the desired result.

Fix S large enough as required by Lemma 3 and recall ϵ_2 and δ_2 therein. For each n , define sequence $\{\tau_l^{(n)}, l = 0, 1, \dots, l^*\}$ in interval $[0, T]$ such that $\tau_0^{(n)} = 0$, and,

$$\tau_{l+1}^{(n)} = \tau_l^{(n)} + \sigma_n(\tau_l^{(n)})S ,$$

where $\sigma_n(\cdot) = \sqrt{q_*^{(n)}(\cdot)}/\sqrt{n}$, and l^* is the largest integer (that depends on n) such that $\tau_{l^*}^{(n)} \leq T$. Let $B_2^{(n)}$ be the union of intervals $[\tau_l^{(n)}, \tau_{l+1}^{(n)})$ over which the refined cost is strictly less than $\epsilon_2 \sigma_n(\tau_l^{(n)})S$, i.e.,

$$B_2^{(n)} = \{ \cup [\tau_l^{(n)}, \tau_{l+1}^{(n)}) : \bar{J}_{\tau_{l+1}^{(n)}}^{(n)} - \bar{J}_{\tau_l^{(n)}}^{(n)} < \epsilon_2 \sigma_n(\tau_l^{(n)})S \} ,$$

and $B_1^{(n)} = [0, T] \setminus B_2^{(n)}$. Now pick a positive $\epsilon_5 < \epsilon_2$. From this point on, we assume n large enough. The Lebesgue measures of $B_1^{(n)}$ and $B_2^{(n)}$ satisfy $\nu(B_1^{(n)}) \leq \frac{2\epsilon_5}{\epsilon_2} T$ and $\nu(B_2^{(n)}) \geq (1 - \frac{2\epsilon_5}{\epsilon_2})T$. Moreover, over any interval $[\tau_l^{(n)}, \tau_{l+1}^{(n)}) \subseteq B_2^{(n)}$, we must have,

$$\langle b, q^{(n)}(\tau_{l+1}^{(n)}) - q^{(n)}(\tau_l^{(n)}) \rangle \leq -\frac{\delta_2}{2} \sigma_n(\tau_l^{(n)})S .$$

otherwise we could construct a LFSP contradicting Lemma 3 by choosing a subsequence in n along which the above inequality does not hold and $\tau_l^{(n)}$ converges to some $\tau \in [0, T]$. Then we have the following bound on the total increment of $\langle b, q^{(n)}(\cdot) \rangle$ over $B_2^{(n)}$,

$$\sum_{[\tau_l^{(n)}, \tau_{l+1}^{(n)}] \subseteq B_2^{(n)}} \langle b, q^{(n)}(\tau_{l+1}^{(n)}) - q^{(n)}(\tau_l^{(n)}) \rangle \leq -\frac{\delta_2}{2} \left(1 - \frac{2\epsilon_5}{\epsilon_2}\right) T .$$

Also, because we have assumed bounded arrivals, there is a finite $c_1 > 0$ such that the total increment of $\langle b, q^{(n)}(\cdot) \rangle$ over $[0, T] \setminus B_2^{(n)}$ satisfies,

$$\langle b, q^{(n)}(T) - q^{(n)}(\tau_{l^*}) \rangle + \sum_{[\tau_l^{(n)}, \tau_{l+1}^{(n)}] \subseteq B_1^{(n)}} \langle b, q^{(n)}(\tau_{l+1}^{(n)}) - q^{(n)}(\tau_l^{(n)}) \rangle \leq c_1 \frac{2\epsilon_5}{\epsilon_2} T .$$

Finally, summing the above two inequalities and taking limit, we get

$$\langle b, q(T) - q(0) \rangle \leq \limsup_{n \rightarrow \infty} \langle b, q^{(n)}(T) - q^{(n)}(0) \rangle \leq -T \left(\frac{\delta_2}{2} - \frac{\epsilon_5}{\epsilon_2} (\delta_2 + 2c_1) \right) .$$

Now we can fix a positive $\delta_5 < \delta_2/2$ and choose $\epsilon_5 > 0$ small enough to satisfy (38), completing the proof. \blacksquare

Remark 5 While the reader could have concluded the throughput-optimality of the p-Log rule from Lemma 2 alone, nevertheless, the throughput-optimality explicitly bears out as a corollary of the above theorem. Specifically, the theorem shows that for the zero cost limiting trajectories (f, g, q) — i.e. $(\mathbf{f}^{(n)}, \mathbf{g}^{(n)}, \mathbf{q}^{(n)})$ converge u.o.c to (f, g, q) with probability 1— with initial condition $\langle b, q(t) \rangle = 1$, we have $\langle b, q(T) \rangle = 0$ for all $T > \frac{1}{\delta_5}$. Consequently,

$$\limsup_{n \rightarrow \infty} \sup_{\sum_{i \in I} b_i q_i^{(n)}(0) \leq 1} E \left[\sum_{i \in I} b_i q_i^{(n)}(T) \right] = 0. \quad (39)$$

By [23] (or see Theorem 4.1 of [4],) this proves the throughput-optimality of p-Log rule.

Equation (39) along with Dynkin's formula implies the following result.

Lemma 5 (See Lemma 8.8 of [15]) *Let constants $C^* > \delta > 0$ be fixed. Consider the stopping time,*

$$\beta^{(n)} = \inf \left\{ t \geq 0 : \langle b, \mathbf{q}^{(n)}(t) \rangle \leq \delta \right\} .$$

Then, for all sufficiently large n , uniformly on the initial states with $\langle b, \mathbf{q}^{(n)}(0) \rangle \leq C^$, we have,*

$$E[\beta^{(n)}] \leq \Delta_2 C^* ,$$

for some finite $\Delta_2 > 0$.

Proof of Theorem 2: Now we can proceed with the proof of Theorem 2. Consider the scaled (random) process $\mathbf{q}^{(n)}$. For any fixed constants $C^* > 1 > \epsilon^* > \epsilon > \delta > 0$, define the following stopping times,

$$\begin{aligned}\alpha^{(n)} &= \inf \left\{ t > 0 : \langle b, \mathbf{q}^{(n)}(t) \rangle \geq 1 \right\} , \\ \beta^{(n)} &= \inf \left\{ t > 0 : \langle b, \mathbf{q}^{(n)}(t) \rangle \leq \delta \right\} , \\ \eta^{(n)} &= \inf \left\{ t > \beta^{(n)} : \langle b, \mathbf{q}^{(n)}(t) \rangle \geq \epsilon \right\} .\end{aligned}$$

Let $p^{(n)}$ denote the stationary distribution of process $\mathbf{q}^{(n)}(t)$, $p_x^{(n)}$ its distribution conditional on $\mathbf{q}^{(n)}(0) = x$, and E_x expectation under $p_x^{(n)}$. Then, for any arbitrary $T > 0$, it is easy to show (see (8.14) [15]) the following upper bound on the probability of overflow,

$$p^{(n)} \left(\sum_{i \in I} b_i \mathbf{q}_i^{(n)} > 1 \right) \leq \frac{\sup_{x: \langle b, x \rangle \leq \epsilon^*} \left(p_x^{(n)}(\beta^{(n)} \geq T) + p_x^{(n)}(\alpha^{(n)} \leq T) \right) \cdot \sup_{y: \langle b, y \rangle \leq C^*} E_y[\beta^{(n)}]}{\inf_{z: \langle b, z \rangle \geq \epsilon} E_z[\eta^{(n)}]} . \quad (40)$$

Now, by uniform upper bound in Lemma 5, we have $\limsup_{n \rightarrow \infty} \sup_{y: \langle b, y \rangle \leq C^*} E_y[\beta^{(n)}] \leq \Delta_2 C^*$; and by bounded arrivals, we have the lower bound $\liminf_{n \rightarrow \infty} \inf_{z: \langle b, z \rangle \geq \epsilon} E_z[\eta^{(n)}] > 0$. Therefore, the terms that are going to decide the limit,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log p^{(n)} \left(\sum_{i \in I} b_i \mathbf{q}_i^{(n)} > 1 \right) ,$$

are $p_x^{(n)}(\beta^{(n)} \geq T)$ and $p_x^{(n)}(\alpha^{(n)} \leq T)$.

By Theorem 5, we can choose T large enough such that $K(\epsilon^*, \delta, T) \geq K(C^*, \delta, T) \geq J_{**}$. Then by Theorem 4,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \sup_{x: \langle b, x \rangle \leq \epsilon^*} p_x^{(n)}(\beta^{(n)} \geq T) \leq -K(C^*, \delta, T) \leq -J_{**} .$$

By (36) and the definition of J_{**} in (17), for any $\epsilon_6 > 0$, we can choose an even larger a T (if required) and an $\epsilon^* > 0$ small enough such that $J_{**, \leq T, \epsilon^*} > J_{**} - \epsilon_6$. Then by Theorem 3,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \sup_{x: \langle b, x \rangle \leq \epsilon^*} p_x^{(n)}(\alpha^{(n)} \leq T) \leq -J_{**, \leq T, \epsilon^*} < -(J_{**} - \epsilon_6) .$$

Since we can choose ϵ_6 arbitrarily small (and subsequently fix constants $1 > \epsilon^* > \epsilon > \delta > 0$), substituting the above two bounds in (40) completes the proof. \blacksquare