

Queue Length Asymptotics for Generalized Max-Weight Scheduling in the presence of Heavy-Tailed Traffic

Krishna Jagannathan, Mihalis Markakis, Eytan Modiano, and John N. Tsitsiklis
Massachusetts Institute of Technology
Cambridge, MA 02139

Email: {krishnaj, mihalis, modiano, jnt}@mit.edu

Abstract—We investigate the asymptotic behavior of the steady-state queue length distribution under generalized max-weight scheduling in the presence of heavy-tailed traffic. We consider a system consisting of two parallel queues, served by a single server. One of the queues receives heavy-tailed traffic, and the other receives light-tailed traffic. We study the class of throughput optimal max-weight- α scheduling policies, and derive an exact asymptotic characterization of the steady-state queue length distributions. In particular, we show that the tail of the light queue distribution is heavier than a power-law curve, whose tail coefficient we obtain explicitly. Our asymptotic characterization also shows that the celebrated max-weight scheduling policy leads to the *worst possible* tail of the light queue distribution, among all non-idling policies.

Motivated by the above ‘negative’ result regarding the max-weight- α policy, we analyze a log-max-weight (LMW) scheduling policy. We show that the LMW policy guarantees an exponentially decaying light queue tail, while still being throughput optimal.

I. INTRODUCTION

Traditionally, traffic in telecommunication networks has been modeled using Poisson and Markov-modulated processes. These simple traffic models exhibit ‘local randomness’, in the sense that much of the variability occurs in short time scales, and only an average behavior is perceived at longer time scales. With the spectacular growth of packet-switched networks such as the internet during the last couple of decades, these traditional traffic models have been shown to be inadequate. This is because the traffic in packetized data networks is intrinsically more ‘bursty’, and exhibits correlations over longer time scales than can be modeled by any Markovian point process. Empirical evidence, such as the famous Bellcore study on self-similarity and long-range dependence in ethernet traffic [14] lead to increased interest in traffic models with high variability.

Heavy-tailed distributions, which have long been used to model high variability and risk in finance and insurance, were considered as viable candidates to model traffic in data networks. Further, theoretical work such as [12], linking heavy-tails to long-range dependence (LRD) lent weight to the belief that extreme variability in the internet file sizes is ultimately

responsible for the LRD traffic patterns reported in [14] and elsewhere.

Many of the early queueing theoretic results for heavy-tailed traffic were obtained for the single server queue; see [4], [5], [19] for surveys of these results. It turns out that the service discipline plays an important role in the delay experienced in a queue, when the traffic is heavy-tailed. For example, it was shown in [1] that any non-preemptive service discipline leads to infinite expected delay, when the traffic is sufficiently heavy-tailed. Further, the asymptotic behavior of the delay under various service disciplines such as first-come-first-served (FCFS) and processor sharing (PS), is markedly different under light-tailed and heavy-tailed scenarios [4], [24]. This is important, for example, in the context of scheduling jobs in server farms [11].

In the context of communication networks, a subset of the traffic flows may be well modeled as heavy-tailed, and the rest better modeled as light-tailed. In such a scenario, there are relatively few studies on the problem of scheduling *between* the different flows, and the ensuing nature of interaction between the heavy-tailed and light-tailed traffic. Perhaps the earliest, and one of the most important studies in this category is [3], where the interaction between light and heavy-tailed traffic flows under generalized processor sharing (GPS) is studied. In that paper, the authors derive the asymptotic workload behavior of the light-tailed flow, when its GPS weight is greater than its traffic intensity.

One of the key considerations in the design of a scheduling policy for a queueing network is *throughput optimality*, which is the ability to support the largest set of traffic rates that is supportable by a given queueing network. Queue length based scheduling policies, such as max-weight scheduling [22], [23] and its many variants, are known to be throughput optimal in a general queueing network. For this reason, the max-weight family of scheduling policies has received much attention in various networking contexts, including switches [16], satellites [17], wireless [18], and optical networks [6].

In spite of a large and varied body of literature related to max-weight scheduling, it is somewhat surprising that the policy has not been adequately studied in the context of heavy-tailed traffic. Specifically, a question arises as to what behavior we can expect due to the interaction of heavy and light-tailed

This work was supported by NSF grants CNS-0626781, CNS-0915988, and CCF-0728554, and by ARO Muri grant number W911NF-08-1-0238.

flows, when a throughput optimal max-weight-like scheduling policy is employed. Our present work is aimed at addressing this basic question.

In a recent paper [15], a special case of the problem considered here is studied. Specifically, it was shown that when the heavy-tailed traffic has an infinite variance, the light-tailed traffic experiences an infinite expected delay under max-weight scheduling. Further, it was shown that the max-weight policy can be tweaked to favor the light-tailed traffic, so as to make the expected delay of the light-tailed traffic finite. In the present paper, we considerably generalize these results by providing a precise asymptotic characterization of the occupancy distributions under the max-weight scheduling family, for a large class of heavy-tailed traffic distributions.

We study a system consisting of two parallel queues, served by a single server. One of the queues is fed by a heavy-tailed arrival process, while the other is fed by light-tailed traffic. We refer to these queues as the ‘heavy’ and ‘light’ queues, respectively. In this setting, we analyze the asymptotic performance of max-weight- α scheduling, which is a generalized version of max-weight scheduling. Specifically, while max-weight scheduling makes scheduling decisions by comparing the queue lengths in the system, the max-weight- α policy uses different powers of the queue lengths to make scheduling decisions. Under this policy, we derive an asymptotic characterization of the light queue occupancy distribution, and specify all the bounded moments of the queue lengths.

An interesting conclusion from our asymptotic characterization is that the ‘plain’ max-weight scheduling policy induces the worst possible asymptotic behavior on the light queue tail. We also show that by a choice of parameters in the max-weight- α policy that increases the preference afforded to the light queue, the tail behavior of the light queue can be improved. Ultimately however, the tail of the light queue distribution is lower bounded by a power-law-like curve, for any scheduling parameters used in the max-weight- α scheduling policy. Intuitively, the reason max-weight- α scheduling induces a power-law-like decay on the light queue distribution is that the light queue has to compete with a typically large heavy queue for service.

The simplest way to guarantee a good asymptotic behavior for the light queue distribution is to give the light queue complete priority over the heavy queue, so that it does not have to compete with the heavy queue for service. We show that under priority for the light queue, the tail distributions of both queues are asymptotically as good as they can possibly be under any policy. Be that as it may, giving priority to the light queue has an important shortcoming – it is not throughput optimal for a general constrained queueing system.

We therefore find ourselves in a situation where on the one hand, the throughput optimal max-weight- α scheduling leads to poor asymptotic performance for the light queue. On the other hand, giving priority to the light queue leads to good asymptotic behavior for both queues, but is not throughput optimal in general. To remedy this situation, we propose a throughput optimal log-max-weight (LMW) scheduling pol-




Fig. 1. A system of two parallel queues, with one of them receiving heavy-tailed traffic.

icy, which gives significantly more importance to the light queue compared to max-weight- α scheduling. We analyze the asymptotic behavior of the LMW policy and show that the light queue occupancy distribution decays exponentially. We also obtain the exact large deviation exponent of the light queue tail under a regularity assumption on the heavy-tailed input. Thus, the LMW policy has both desirable attributes – it is throughput optimal, and ensures an exponentially decaying tail for the light queue distribution.

The remainder of this paper is organized as follows. In Section II, we describe the system model. In Section III, we present the relevant definitions and mathematical preliminaries. Section IV deals with the queue length behavior under priority scheduling. Sections V and VI respectively contain our asymptotic results for max-weight- α scheduling, and the LMW policy. We conclude the paper in Section VII.

We omit several proofs due to space constraints. We refer the reader to [13] for detailed proofs of the results in the paper.

II. SYSTEM MODEL

Our system consists of two parallel queues, H and L , served by a single server, as depicted in Fig. 1. Time is slotted, and stochastic arrivals of packet bursts occur to each queue in each slot. The server is capable of serving one packet per time slot, from only one of the queues according to a scheduling policy. Let $H(t)$ and $L(t)$ denote the number of packets that arrive during slot t to H and L respectively. Although we postpone the precise assumptions on the traffic to Section III-B, let us loosely say that the input $L(t)$ is light-tailed, and $H(t)$ is heavy-tailed. We will refer to the queues H and L as the heavy and light queues, respectively.

Let $q_H(t)$ and $q_L(t)$, respectively, denote the number of packets in H and L during slot t , and let q_H and q_L denote the steady-state queue lengths, when they exist. Our aim is to characterize the behavior of $\mathbb{P}\{q_L > b\}$ and $\mathbb{P}\{q_H > b\}$ as b becomes large, under various scheduling policies.

III. DEFINITIONS AND MATHEMATICAL PRELIMINARIES

A. Heavy-tailed distributions

We define some properties regarding the tail distributions of non-negative random variables.

Definition 1: A random variable X is said to be *light-tailed* if there exists $\theta > 0$ for which $\mathbb{E}[\exp(\theta X)] < \infty$. A random variable is *heavy-tailed* if it is not light-tailed.

In other words, a light-tailed random variable is one that has a well defined moment generating function in a neighborhood of the origin. The complementary distribution function of a

light-tailed random variable decays at least exponentially fast. Heavy-tailed random variables are those which have complementary distribution functions that decay slower than any exponential. This class is often too general to study, so subclasses of heavy-tailed distributions, such as sub-exponentials have been defined and studied in the past [21]. We now define the tail-coefficient of a random variable.

Definition 2: The *tail coefficient* of a random variable X is defined by

$$C_X = \sup\{c \mid \mathbb{E}[X^c] < \infty\}.$$

In words, the tail coefficient is the threshold where the power moment of a random variable starts to blow up. Note that the tail coefficient of a light-tailed random variable is infinite. On the other hand, the tail coefficient of a heavy-tailed random variable may be infinite (e.g., log-normal) or finite (e.g., Pareto).

We restrict our attention to a fairly general class of heavy-tailed distributions, known in the literature as *order-regular varying* distributions; see [2], [7]. We repeat the relevant definitions for ease of reference. Let X denote a non-negative random variable, with complementary distribution function $\bar{F}(x) = \mathbb{P}\{X > x\}$.

Notation: If $f(x)$ and $g(x)$ are positive functions defined on $[0, \infty]$, we write $f(x) \sim g(x)$ to mean $\lim_{x \rightarrow \infty} \frac{f(x)}{g(x)} = 1$. Similarly, $f(x) \gtrsim g(x)$ means $\liminf_{x \rightarrow \infty} \frac{f(x)}{g(x)} \geq 1$.

Definition 3: $\bar{F}(x)$ is said to have a *regularly varying* tail of index ν , notation $\bar{F} \in \mathcal{R}(\nu)$, if $\bar{F}(x) = U(x)x^{-\nu}$, where $U(x)$ is a *slowly varying* function, i.e., a function that satisfies $U(kx) \sim U(x)$, $\forall k > 0$.

Definition 4: $\bar{F}(x)$ is said to be *order-regularly varying*, notation $\bar{F} \in \mathcal{OR}$, if for some $\Lambda > 1$,

$$0 < \liminf_{x \rightarrow \infty} \frac{\bar{F}(\lambda x)}{\bar{F}(x)} \leq \limsup_{x \rightarrow \infty} \frac{\bar{F}(\lambda x)}{\bar{F}(x)} < \infty, \quad \forall \lambda \in [1, \Lambda].$$

It can be shown that $\mathcal{R} \subset \mathcal{OR}$ and that the inclusion is strict. In what follows, a statement such as $X \in \mathcal{OR}$ should be construed to mean $\mathbb{P}\{X > x\} \in \mathcal{OR}$.

It turns out that the any order-regularly varying distribution is asymptotically heavier than some power-law curve (see [20]). Further, any heavy-tailed random variable with an order-regularly varying tail has a finite tail coefficient [13]. We now state the precise assumptions on the arrival processes.

B. Assumptions on the arrival processes

- 1) The arrival processes $H(t)$ and $L(t)$ are independent of each other, and independent of the current state of the system.
- 2) $H(t)$ is independent and identically distributed (i.i.d.) from slot-to-slot.
- 3) $L(t)$ is i.i.d. from slot-to-slot.
- 4) $L(\cdot)$ is light-tailed with $\mathbb{E}[L(t)] = \lambda_L$.
- 5) $H(\cdot) \in \mathcal{OR}$ with tail coefficient $C_H > 1$, and $\mathbb{E}[H(t)] = \lambda_H$.

We also assume that $\lambda_L + \lambda_H < 1$, so that the input rate does not overwhelm the service rate. Then, it can be shown that

the system is stable¹ under any non-idling policy, and that the steady-state queue lengths q_H and q_L exist.

C. Residual and Age distributions

Here, we define the residual and age distributions for the heavy-tailed input process, which will be useful later. First, we note that $H(\cdot)$ necessarily has a non zero probability mass at zero, since $\lambda_H < 1$. Define H_+ as the strictly positive part of $H(\cdot)$. Specifically, $\mathbb{P}\{H_+ = m\} = \mathbb{P}\{H(\cdot) = m \mid H(\cdot) > 0\}$, $m = 1, 2, \dots$. Note that H_+ has tail coefficient equal to C_H .

Now consider a discrete-time renewal process with inter-renewal times distributed as H_+ . Let $H_R \in \{1, 2, \dots\}$ denote the residual random variable, and $H_A \in \{0, 1, \dots\}$ the age of the renewal process [9].² The joint distribution of the residual and the age can be derived using basic renewal theory:

$$\mathbb{P}\{H_R = k, H_A = l\} = \frac{\mathbb{P}\{H_+ = k + l\}}{\mathbb{E}[H_+]}, \quad (1)$$

where $k \in \{1, 2, \dots\}$, $l \in \{0, 1, \dots\}$. The marginal of H_R can be derived from (1):

$$\mathbb{P}\{H_R = k\} = \frac{\mathbb{P}\{H_+ \geq k\}}{\mathbb{E}[H_+]}, \quad k \in \{1, 2, \dots\}. \quad (2)$$

We now state without proof, the important result that the residual and age distributions are *one order heavier* than the original distribution.

Proposition 1: If $H(\cdot) \in \mathcal{OR}$ has tail coefficient equal to C_H , then H_R and H_A have tail coefficient equal to $C_H - 1$.

IV. THE PRIORITY POLICIES

In this section, we study the two ‘extreme’ scheduling policies, namely priority for L and priority for H . Our analysis helps us arrive at the important conclusion that the tail of the heavy queue is asymptotically insensitive to the scheduling policy. In other words, there is not much we can do to improve or hurt the tail distribution of H by the choice of a scheduling policy.

A. Priority for H

Under priority for H , the heavy queue receives service whenever it is non-empty, and L receives service only when H is empty. It should be intuitively clear at the outset that this policy is bound to have an undesirable impact on the light queue. The reason we analyze this policy is that it gives us a best case scenario for the heavy queue. In particular, the following result shows that the heavy queue occupancy is one order heavier than its input distribution.

Proposition 2: Under priority for H , the steady-state queue occupancy distribution of the heavy queue satisfies the following bounds.

¹The notion of stability used here is the positive recurrence of the system occupancy Markov chain.

²We have defined the residual time and age such that if a renewal occurs at a particular time slot, the age at that time slot is zero, and the residual time is equal to the length of the upcoming renewal interval.

1) For every $\epsilon > 0$, there exists a constant $\kappa_H(\epsilon) > 0$ such that

$$\mathbb{P}\{q_H > b\} < \kappa_H(\epsilon)b^{-(C_H-1-\epsilon)}, \forall b. \quad (3)$$

2)

$$\mathbb{P}\{q_H > b\} \geq \lambda_H \mathbb{P}\{H_R > b\}, \forall b. \quad (4)$$

Further, q_H is a heavy-tailed random variable with tail coefficient equal to $C_H - 1$. That is, for each $\epsilon > 0$, we have

$$\mathbb{E}\left[q_H^{C_H-1-\epsilon}\right] < \infty, \quad (5)$$

and

$$\mathbb{E}\left[q_H^{C_H-1+\epsilon}\right] = \infty. \quad (6)$$

Proof: Equation (5) can be shown using a straightforward Lyapunov argument, along the lines of [15, Proposition 6]. Equation (3) follows from (5) and the Markov inequality.

Next, to show (4), we consider a time instant t at steady-state, and write

$$\begin{aligned} \mathbb{P}\{q_H(t) > b\} &= \mathbb{P}\{q_H(t) > b | q_H(t) > 0\} \mathbb{P}\{q_H(t) > 0\} \\ &= \lambda_H \mathbb{P}\{q_H(t) > b | q_H(t) > 0\}. \end{aligned}$$

We have used Little's law at steady-state to write $\mathbb{P}\{q_H(t) > 0\} = \lambda_H$. Let us now lower bound the term $\mathbb{P}\{q_H(t) > b | q_H(t) > 0\}$. Conditioned on H being non-empty, denote by $\tilde{B}(t)$ the number of packets that belong to the burst in service that still remain in the queue at time t . Then, clearly, $q_H(t) \geq \tilde{B}(t)$, from which $\mathbb{P}\{q_H(t) > b | q_H(t) > 0\} \geq \mathbb{P}\{\tilde{B}(t) > b\}$. Now, since the H queue receives service whenever it is non-empty, it is clear that the time spent at the head-of-line by a burst is *equal* to its size. It can therefore be shown that in steady-state, $\tilde{B}(t)$ is distributed according to the residual variable H_R . Thus, $\mathbb{P}\{q_H(t) > b | q_H(t) > 0\} \geq \mathbb{P}\{H_R > b\}$, and (4) is proved. Finally, (6) follows from (4) and Proposition 1. \square

Furthermore, when the distribution of $H(\cdot)$ is regularly varying, the lower bound (4) takes on a power-law form that agrees with the upper bound (3).

Corollary 1: If $H(\cdot) \in \mathcal{R}(C_H)$, then

$$\mathbb{P}\{q_H > b\} > U(b)b^{-(C_H-1)}, \forall b,$$

where $U(\cdot)$ is some slowly varying function.

Since priority for H affords the most favorable treatment to the heavy queue, it follows that the asymptotic behavior of H can be no better than the above under *any* policy.

Proposition 3: Under any scheduling policy, q_H is heavy-tailed with tail coefficient at most $C_H - 1$. That is, Equation (6) holds for all scheduling policies.

Interestingly, under priority for H , the steady-state light queue occupancy q_L is also heavy-tailed with the *same* tail coefficient as q_H . This should not be surprising, since the light queue has to wait for the entire heavy queue to clear, before it receives any service.

Proposition 4: Under priority for H , q_L is heavy-tailed with tail coefficient $C_H - 1$.

B. Priority for L

We now study the policy that serves L whenever it is non-empty, and serves H only if L is empty. This policy affords the best possible treatment to L and the worst possible treatment to H , among all non-idling policies. Under this policy, L is completely oblivious to the presence of H , in the sense that it receives service whenever it has a packet to be served. Therefore, L behaves like a discrete time G/D/1 queue, with light-tailed inputs. Classical large deviation bounds can be derived for such a queue; see [10] for example.

Recall that since $L(\cdot)$ is light-tailed, the log moment generating function

$$\Lambda_L(\theta) = \log \mathbb{E}\left[e^{\theta L(\cdot)}\right]$$

exists for some $\theta > 0$. Define

$$E_L = \sup\{\theta | \Lambda_L(\theta) - \theta < 0\}. \quad (7)$$

Proposition 5: Under priority for L , q_L satisfies the large deviation principle (LDP)

$$\lim_{b \rightarrow \infty} -\frac{1}{b} \log \mathbb{P}\{q_L > b\} = E_L. \quad (8)$$

In words, the above proposition asserts that the tail of q_L is asymptotically exponential, with rate function E_L . We will refer to E_L as the *intrinsic exponent* of the light queue. An equivalent expression for the intrinsic exponent that is often used in the literature is

$$E_L = \inf_{a>0} \frac{1}{a} \Lambda_L^*(1+a), \quad (9)$$

where $\Lambda_L^*(\cdot)$ is the Fenchel-Legendre transform [10] of $\Lambda_L(\theta)$.

It is clear that priority for L gives the best possible asymptotic behavior for the light queue, and the worst possible treatment for the heavy queue. Surprisingly however, it turns out that the heavy queue tail under priority for L is asymptotically as good as it is under priority for H .

Proposition 6: Under priority for L , q_H is heavy-tailed with tail coefficient $C_H - 1$.

Proof: This is a special case of Theorem 2, given in the next section. \square

The above result also implies that the tail coefficient of H cannot be worse than $C_H - 1$ under any other scheduling policy.

Proposition 7: Under any non-idling scheduling policy, q_H has a tail coefficient of at least $C_H - 1$. That is, Equation (5) holds for all non-idling scheduling policies.

Propositions 3 and 7 together imply the insensitivity of the heavy queue's tail distribution to the scheduling policy. We state this important result in the following theorem.

Theorem 1: Under *any* non-idling scheduling policy, q_H is heavy-tailed with tail coefficient equal to $C_H - 1$. Further, $\mathbb{P}\{q_H > b\}$ satisfies bounds of the form (3) and (4) under all non-idling policies.

Therefore, it is not possible to either improve or hurt the heavy queue's asymptotic behavior, by the choice of a scheduling policy.

It is evident that the light queue has the best possible asymptotic behavior under priority for L . Although priority for L is non-idling, and therefore throughput-optimal in this simple setting, we are ultimately interested in studying more sophisticated network models, where priority for L may not be throughput optimal. We therefore analyze the asymptotic behavior of general throughput optimal policies belonging to the max-weight family.

V. QUEUE LENGTH ASYMPTOTICS FOR MAX-WEIGHT- α SCHEDULING

In this section, we analyze the asymptotic tail behavior of the light queue distribution under max-weight- α scheduling. For fixed parameters $\alpha_H > 0$ and $\alpha_L > 0$, the max-weight- α policy operates as follows: During each time slot t , perform the comparison

$$q_L(t)^{\alpha_L} \stackrel{\geq}{\leq} q_H(t)^{\alpha_H},$$

and serve one packet from the queue that wins the comparison. Ties can be broken arbitrarily, but we break them in favor of the light queue for the sake of definiteness. Note that $\alpha_L = \alpha_H$ corresponds to the usual max-weight policy, which serves the longest queue in each slot. $\alpha_L/\alpha_H > 1$ corresponds to emphasizing the light queue over the heavy queue, and vice-versa.

We provide an asymptotic characterization of the light queue occupancy distribution under max-weight- α scheduling by deriving matching upper and lower bounds. Our characterization shows that the light queue occupancy is heavy-tailed under max-weight- α scheduling for all values of the parameters α_H and α_L . Since we obtain distributional bounds on the light queue occupancy, our results also shed further light on the moment results derived in [15] for max-weight- α scheduling.

A. Upper Bound

In this section, we derive two different upper bounds on the tail probability $\mathbb{P}\{q_L > b\}$, that both hold under max-weight- α scheduling. However, depending on the values of α_H and α_L , one of them would be tighter. The first upper bound holds for all non-idling policies, including max weight- α scheduling.

Theorem 2: Under any non-idling policy, and for every $\epsilon > 0$, there exists a constant $\kappa_1(\epsilon) > 0$, such that

$$\mathbb{E}[q_L^{C_H-1-\epsilon}] < \infty \quad (10)$$

and

$$\mathbb{P}\{q_L > b\} < \kappa_1(\epsilon)b^{-(C_H-1-\epsilon)}. \quad (11)$$

Proof: (Sketch) The main idea is to combine the two queues into one composite queue which is fed by the sum input process $L(\cdot) + H(\cdot)$. This sum input process can be shown to have a tail coefficient equal to C_H . Now, under any non-idling policy in the original system, the composite queue behaves like a G/D/1 queue, with heavy-tailed input of tail coefficient C_H . The occupancy of the composite queue is given by $q = q_L + q_H$. Therefore, $\mathbb{E}[(q_H + q_L)^{C_H-1-\epsilon}] = \mathbb{E}[q^{C_H-1-\epsilon}] < \infty$ for any $\epsilon > 0$, by applying Equation (5)

to the composite queue. This implies (10), and (11) follows from (10) using the Markov inequality. \square

The above result asserts that the tail coefficient of q_L is at least $C_H - 1$ under any non-idling policy, and that $\mathbb{P}\{q_L > b\}$ is uniformly upper bounded by a power-law curve. Our second upper bound is specific to max-weight- α scheduling. It hinges on a simple observation regarding the scaling of the α parameters, in addition to a theorem in [15]. We first state the following elementary observation due to its usefulness.

Observation: (Scaling of α parameters) Let α_H and α_L be given parameters of a max-weight- α policy, and let $\beta > 0$ be arbitrary. Then, the max-weight- α policy that uses the parameters $\beta\alpha_H$ and $\beta\alpha_L$ for the queues H and L respectively, is *identical* to the original policy. That is, in each time slot, the two policies make the same scheduling decision.

Next, let us invoke an important result from [15].

Theorem 3: If max-weight- α scheduling is performed with $0 < \alpha_H < C_H - 1$, then, for any $\alpha_L > 0$, we have $\mathbb{E}[q_L^{\alpha_L}] < \infty$.

Thus, by choosing a large enough α_L , any moment of the light queue length can be made finite, as long as $\alpha_H < C_H - 1$. Our second upper bound, which we state next, holds regardless of how the α parameters are chosen.

Theorem 4: Define

$$\gamma = \frac{\alpha_L}{\alpha_H}(C_H - 1). \quad (12)$$

Under max weight α scheduling, for every $\epsilon > 0$, there exists a constant $\kappa_2(\epsilon) > 0$, such that

$$\mathbb{E}[q_L^{\gamma-\epsilon}] < \infty \quad (13)$$

and

$$\mathbb{P}\{q_L > b\} < \kappa_2(\epsilon)b^{-(\gamma-\epsilon)}. \quad (14)$$

Proof: Given $\epsilon > 0$, let us choose $\beta = (C_H - 1)/\alpha_H - \epsilon/\alpha_L$, and perform max-weight- α scheduling with parameters $\beta\alpha_H$ and $\beta\alpha_L$. According to the above observation, this policy is identical to the original max-weight- α policy. Next, since $\beta\alpha_H < C_H - 1$, Theorem 3 applies, and we have $\mathbb{E}[q_L^{\beta\alpha_L}] = \mathbb{E}[q_L^{\gamma-\epsilon}] < \infty$, which proves (13). Finally, (14) can be proved using (13) and the Markov inequality. \square

The above theorem asserts that the tail coefficient of q_L is at least γ under the max weight- α policy. We remark that Theorem 2 and Theorem 4 both hold for max-weight- α scheduling with any parameters. However, one of them yields a stronger bound than the other, depending on the α parameters. Specifically, we have the following two cases:

- (i) $\frac{\alpha_L}{\alpha_H} \leq 1$: This is the regime where the light queue is given lesser priority compared to the heavy queue. In this case, Theorem 2 yields a stronger bound.
- (ii) $\frac{\alpha_L}{\alpha_H} > 1$: This is the regime where the light queue is given more priority compared to the heavy queue. In this case, Theorem 4 gives the stronger bound.

Remark 1: The upper bounds in this section hold whenever $H(\cdot)$ is heavy-tailed with tail coefficient C_H . We need the

assumption $H(\cdot) \in \mathcal{OR}$ only to derive the lower bounds in the next subsection.

B. Lower Bound

In this section, we state our main lower bound result, which asymptotically lower bounds the tail of the light queue distribution in terms of the tail of the residual variable H_R .

Theorem 5: Under max-weight- α scheduling with parameters α_H and α_L , the distribution of the light queue occupancy satisfies the following asymptotic lower bounds:

1) If $\frac{\alpha_L}{\alpha_H} < 1$,

$$\mathbb{P}\{q_L \geq b\} \gtrsim \lambda_H \mathbb{P}\left\{H_R \geq \frac{b}{\lambda_L}\right\} \quad (15)$$

2) If $\frac{\alpha_L}{\alpha_H} = 1$,

$$\mathbb{P}\{q_L \geq b\} \gtrsim \lambda_H \mathbb{P}\left\{H_R \geq b\left(1 + \frac{1}{\lambda_L}\right)\right\} \quad (16)$$

3) If $\frac{\alpha_L}{\alpha_H} > 1$,

$$\mathbb{P}\{q_L \geq b\} \gtrsim \lambda_H \mathbb{P}\left\{H_R \geq b^{\alpha_L/\alpha_H}\right\}. \quad (17)$$

The above theorem holds when the heavy-tailed input distribution is order-regular. As a special case of the above theorem, when $H(\cdot)$ is regularly varying with index C_H , the lower bounds take on a more pleasing power-law form that matches the upper bounds (11) and (14).

Corollary 2: If $H(\cdot) \in \mathcal{R}(C_H)$, then, under max-weight- α scheduling with parameters α_H and α_L , the distribution of the light queue satisfies the following asymptotic lower bounds:

1) If $\frac{\alpha_L}{\alpha_H} \leq 1$,

$$\mathbb{P}\{q_L \geq b\} \gtrsim U(b)b^{-(C_H-1)} \quad (18)$$

2) If $\frac{\alpha_L}{\alpha_H} > 1$,

$$\mathbb{P}\{q_L \geq b\} \gtrsim U(b)b^{-\gamma}, \quad (19)$$

where $U(\cdot)$ is some slowly varying function.

The proof of Theorem 5 is quite involved, and is omitted due to space constraints. We refer the reader to [13] for the complete proof.

C. Tail Coefficient of q_L

In this section, we characterize the exact tail coefficient of the light queue distribution under max-weight- α scheduling.

Theorem 6: The tail coefficient of the steady-state queue length q_L of the light queue is given by

- (i) $C_H - 1$ for $\frac{\alpha_L}{\alpha_H} \leq 1$, and
- (ii) $\gamma = \frac{\alpha_L}{\alpha_H}(C_H - 1)$ for $\frac{\alpha_L}{\alpha_H} > 1$.

Proof: Equations (10) and (13) give lower bounds on the tail coefficient of q_L . Matching upper bounds on the tail coefficient follow from Theorem 5 and Proposition 1. We omit the details in the interest of brevity. \square

In Fig. 2, we show the tail coefficient of q_L as a function of the ratio α_L/α_H . We see that the tail coefficient is constant at the value $C_H - 1$ as α_L/α_H varies from 0 to 1. Recall that




Fig. 2. The tail coefficient of q_L under max-weight- α scheduling, as a function of α_L/α_H , for $C_H = 2.5$.

$\alpha_L/\alpha_H = 1$ corresponds to max-weight scheduling, while $\alpha_L/\alpha_H \downarrow 0$ corresponds to priority for H . Thus, the tail coefficient of q_L under max-weight scheduling is the same as the tail coefficient under priority for H , implying that the max-weight policy leads to the *worst possible* asymptotic behavior for the light queue among all non-idling policies. However, the tail coefficient of q_L begins to improve in proportion to the ratio α_L/α_H in the regime where the light queue is given more importance.

Remark 2: If the heavy-tailed input has infinite variance ($C_H < 2$), then it follows from Theorem 6 that the expected delay in the light queue is infinite under max-weight scheduling. Thus, [15, Proposition 5] is a special case of the above theorem.

VI. LOG-MAX-WEIGHT SCHEDULING

We showed in Theorem 6 that the light queue occupancy distribution is necessarily heavy-tailed with a finite tail coefficient, under max-weight- α scheduling. On the other hand, priority for L ensures the best possible asymptotic behavior for both queues, but suffers from possible instability effects in more general queueing networks.

In this section, we analyze the log-max-weight (LMW) policy. We show that the *light queue distribution is light-tailed* under LMW scheduling, i.e., that $\mathbb{P}\{q_L > b\}$ decays exponentially in b . However, unlike priority for L , the LMW policy can be shown to be throughput optimal even in more general settings, as a consequence of [8, Theorem 1]. For our simple system model, we define the LMW policy as follows.

In each time slot t , the log-max-weight policy compares

$$q_L(t) \gtrless \log(1 + q_H(t)),$$

and serves one packet from the queue that wins the comparison. Ties are broken in favor of the light queue.

The main idea in the LMW policy is to give preference to the light queue to a far greater extent than any max-weight- α policy. Specifically, for $\alpha_L/\alpha_H > 1$, the max-weight- α policy

compares q_L to a power of q_H that is smaller than 1. On the other hand, LMW scheduling compares q_L to a logarithmic function of q_H , leading to a significant preference for the light queue. It turns out that this significant de-emphasis of the heavy queue with respect to the light queue is sufficient to ensure an exponential decay for the distribution of q_L in our setting.

Furthermore, the LMW policy has another useful property when the heavy queue gets overwhelmingly large. Although the LMW policy significantly de-emphasizes the heavy queue, it does not *ignore* it, unlike priority for L . That is, if the H queue occupancy gets overwhelmingly large compared to L , the LMW policy will serve the H queue. In contrast, priority for L will ignore any build-up in H , as long as L is non-empty. This property turns out to be crucial in more complex queueing models, where throughput optimality is non-trivial to obtain. For example, when the queues have time-varying connectivity to the server, the LMW policy will stabilize both queues for all rates within the rate region, whereas priority for L can lead to instability effects in H [23].

Our main result in this section shows that under the LMW policy, $\mathbb{P}\{q_L > b\}$ decays exponentially in b , unlike under max-weight- α scheduling.

Theorem 7: Under log-max-weight scheduling, q_L is light-tailed. Specifically, it holds that

$$\liminf_{b \rightarrow \infty} -\frac{1}{b} \log \mathbb{P}\{q_L \geq b\} \geq \min(E_L, C_H - 1), \quad (20)$$

where E_L is the intrinsic exponent, given by (7), (9).

Proof: See Appendix A. \square

Thus, the light queue tail is upper bounded by an exponential term, whose rate of decay is given by the smaller of the intrinsic exponent E_L , and $C_H - 1$. We remark that Theorem 7 utilizes only the light-tailed nature of $L(\cdot)$, and the tail coefficient of $H(\cdot)$. Specifically, we do not need to assume any regularity property such as $H(\cdot) \in \mathcal{OR}$ for the result to hold. However, if we assume that the tail of $H(\cdot)$ is regularly varying, we can obtain a matching lower bound to the upper bound in Theorem 7.

Theorem 8: Suppose that $H(\cdot) \in \mathcal{R}(C_H)$. Then, under LMW scheduling, the tail distribution of q_L satisfies an LDP with rate function given by

$$\lim_{b \rightarrow \infty} -\frac{1}{b} \log \mathbb{P}\{q_L \geq b\} = \min(E_L, C_H - 1).$$

Fig. 3 shows the large deviation exponent given by Theorem 8 as a function of λ_L , for $C_H = 2.5$, and Poisson bursts feeding the light queue. There are two distinct regimes in the plot, corresponding to two fundamentally different modes of overflow. For relatively large values of λ_L , the exponent for the LMW policy equals E_L , the intrinsic exponent. In this regime, the light queue overflows entirely due to atypical behavior in the input process $L(\cdot)$. In other words, q_L would have grown close to the level b even if the heavy queue was absent. This mode of overflow is more likely for larger values of λ_L , which explains the diminishing exponent in this regime.




Fig. 3. The large deviation exponent for q_L under LMW scheduling, as a function of λ_L . The light queue is fed by Poisson bursts, and $C_H = 2.5$.

The flat portion of the curve in Fig. 3 corresponds to a second overflow mode. In this regime, the overflow of the light queue occurs due to extreme misbehavior on the part of the heavy-tailed input. Specifically, the heavy queue becomes larger than e^b after receiving a very large burst. After this instant, the heavy queue hogs all the service, and the light queue gets starved until it gradually builds up to the level b . In this regime, the light queue input behaves typically, and plays no role in the overflow of L . That is, the exponent is independent of λ_L , being equal to a constant $C_H - 1$. The exponent is decided entirely by the ‘burstiness’ of the heavy-tailed traffic, which is reflected in the tail coefficient.

VII. CONCLUDING REMARKS

We considered a system of parallel queues fed by a mix of heavy-tailed and light-tailed traffic, and served by a single server. We studied the asymptotic behavior of the queue size distributions under various scheduling policies. We showed that the occupancy distribution of the heavy queue is asymptotically insensitive to the scheduling policy used, and inevitably heavy-tailed. In contrast, the light queue occupancy distribution can be heavy-tailed or light-tailed depending on the scheduling policy.

The major contribution of the paper is in the derivation of an asymptotic characterization of the light queue occupancy distribution, under max-weight- α scheduling. We showed that the light queue distribution is heavy-tailed with a finite tail coefficient under max-weight- α scheduling, for any values of the scheduling parameters. However, the tail coefficient can be improved by choosing the scheduling parameters to favor the light queue. We also observed that ‘plain’ max-weight scheduling leads to the worst possible asymptotic behavior of the light queue distribution, among all non-idling policies.

Another important contribution of the paper is the log-max-weight policy, and the corresponding asymptotic analysis. We showed that the light queue occupancy distribution is light-tailed under LMW scheduling, and explicitly derived an exponentially decaying upper bound on the tail of the light

queue distribution. Additionally, the LMW policy also has the desirable property of being throughput optimal in a general queueing network.

Although we study a very simple queueing network in this paper, we believe that the insights obtained from this study are valuable in much more general settings. For instance, in a general queueing network with a mix of light-tailed and heavy-tailed traffic flows, we expect that the celebrated max-weight policy has the tendency to ‘infect’ competing light-tailed flows with heavy-tailed asymptotics. A similar effect was also noted in [15], in the context of expected delay.

We also believe that the LMW policy occupies a unique ‘sweet spot’ in the context of scheduling light-tailed traffic in the presence of heavy-tailed traffic. This is because the LMW policy de-emphasizes the heavy-tailed flow sufficiently to maintain good light queue asymptotics, while also ensuring network-wide stability.

For future work, we propose the extension of the results in this paper to more general single-hop and multi-hop networks, with time-varying channel models.

APPENDIX A PROOF OF THEOREM 7

Fix a small $\delta > 0$. We first write the equality

$$\begin{aligned} \mathbb{P}\{q_L \geq b\} &= \underbrace{\mathbb{P}\{q_L \geq b; \log(1 + q_H) < \delta b\}}_{(i)} \\ &+ \underbrace{\mathbb{P}\{q_L \geq b; (1 - \delta)b \geq \log(1 + q_H) \geq \delta b\}}_{(ii)} \\ &+ \underbrace{\mathbb{P}\{q_L \geq b; \log(1 + q_H) > (1 - \delta)b\}}_{(iii)} \quad (21) \end{aligned}$$

We will next upper bound each of the above three terms on the right.

- (i) Intuitively, this event corresponds to an overflow of the light queue, when the light queue is not ‘exponentially large’ in b , i.e., $q_H < \exp(\delta b) - 1$. Suppose without loss of generality that this event happens at time 0. Denote by $-\tau < 0$ the last instant when the heavy queue received service. Since H has not received service since $-\tau$, it is clear that $\log(1 + q_H(-\tau)) < \delta b$. Thus, $q_L(-\tau) < \delta b$. In the time interval $[-\tau + 1, 0]$ the light queue receives service in each slot. In spite of receiving all the service, it grows from less than δb to overflow at time 0. This implies that every time the event in (i) occurs, there necessarily exists a $-u \leq 0$ satisfying

$$\sum_{i=-u+1}^0 (L(i) - 1) > (1 - \delta)b.$$

Therefore,

$$\begin{aligned} (i) &\leq \mathbb{P}\left\{\exists u \geq 0 \mid \sum_{i=-u+1}^0 (L(i) - 1) > (1 - \delta)b\right\} \\ &= \mathbb{P}\left\{\sup_{u \geq 0} (S_u - u) > (1 - \delta)b\right\}, \quad (22) \end{aligned}$$

where $S_u = \sum_{i=-u+1}^0 L(i)$. The right hand side of (22) is precisely the probability of a single server queue fed by the process $L(\cdot)$ reaching the level $(1 - \delta)b$. Standard large deviation bounds are known for such an event. Specifically, from [10, Lemma 1.5], we get

$$\begin{aligned} \liminf_{b \rightarrow \infty} -\frac{1}{b} \log \mathbb{P}\left\{\sup_{u \geq 0} S_u - u > (1 - \delta)b\right\} &\geq \\ \inf_{a > 0} \frac{1 - \delta}{a} \Lambda_L^*(1 + a) &= (1 - \delta)E_L. \quad (23) \end{aligned}$$

From (22) and (23), we see that for every $\epsilon > 0$ and for large enough b ,

$$\mathbb{P}\{q_L \geq b; \log(1 + q_H) < \delta b\} < \kappa_1 e^{-b(1-\delta)(E_L - \epsilon)}. \quad (24)$$

- (ii) Let us deal with the term (iii) before (ii). This is the regime where the overflow of L occurs, along with H becoming exponentially large in b . We have

$$(iii) \leq \mathbb{P}\{q_L + q_H > e^{(1-\delta)b}\}$$

We have shown earlier in the proof of Theorem 2 that for any non-idling policy,

$$\mathbb{P}\{q_L + q_H > M\} < \kappa_2 M^{-(C_H - 1 - \epsilon)}$$

for every $\epsilon > 0$ and some $\kappa_2 > 0$. Therefore,

$$\mathbb{P}\{q_L \geq b; \log(1 + q_H) > (1 - \delta)b < \kappa_2 \exp(-(1 - \delta)b(C_H - 1 - \epsilon)), \forall \epsilon > 0. \quad (25)$$

- (ii) Let us now deal with the second term, $\mathbb{P}\{q_L \geq b; (1 - \delta)b \geq \log(1 + q_H) \geq \delta b\}$.

Suppose this event occurs at time 0. Denote by $-\tau \leq 0$ the last time during the current busy period that H received service³, and define $\eta = \log(1 + q_H(-\tau))$. We can deduce that $\eta \leq (1 - \delta)b$, because H receives no service in $[-\tau + 1, 0]$. It is also clear that $q_L(-\tau) < \eta$. Therefore, L grows from less than η to more than b , in spite of receiving all the service in $[-\tau + 1, 0]$. Using u and ξ as ‘dummy’ variables that represent the possible values taken by τ and η respectively, we can bound the term (ii) as shown in (26) in the next page. The last two steps in (26) are by the union bound. Notice now that for every $u \geq 0$, the event $S_u - u > b - \xi$ is independent of the value of $q_H(-u) + q_L(-u)$, since these are determined by arrivals in disjoint intervals. Therefore, continuing from (26),

$$\begin{aligned} &= \sum_{\xi=0}^{(1-\delta)b} \sum_{u \geq 0} \mathbb{P}\{S_u - u > b - \xi\} \mathbb{P}\{q_H(-u) + q_L(-u) \geq e^\xi\} \\ &= \sum_{\xi=0}^{(1-\delta)b} \sum_{u \geq 0} \mathbb{P}\{S_u - u > b - \xi\} \kappa_2 e^{-(C_H - 1 - \epsilon)\xi}, \forall \epsilon > 0 \quad (27) \end{aligned}$$

$$\leq \sum_{\xi=0}^{(1-\delta)b} \kappa_1 e^{-(E_L - \epsilon)(b - \xi)} \kappa_2 e^{-(C_H - 1 - \epsilon)\xi}, \forall \epsilon > 0. \quad (28)$$

³If H never received service during the current busy period, we take τ to be equal to the last instant that the system was empty, and $\eta = 0$.

$$\begin{aligned}
(ii) &\leq \mathbb{P} \{ \exists \xi \leq (1-\delta)b, u \geq 0 \mid S_u - u > b - \xi; q_H(-u) + q_L(-u) \geq e^\xi \} \leq \\
&\sum_{\xi=0}^{(1-\delta)b} \mathbb{P} \{ \exists u \geq 0 \mid S_u - u > b - \xi; q_H(-u) + q_L(-u) \geq e^\xi \} \leq \sum_{\xi=0}^{(1-\delta)b} \sum_{u \geq 0} \mathbb{P} \{ S_u - u > b - \xi; q_H(-u) + q_L(-u) \geq e^\xi \}
\end{aligned} \tag{26}$$

Equation (27) follows from (11), and (28) is a classical large deviation bound that follows, for example, from [10, Lemma 1.5]. Thus, for every $\epsilon > 0$,

$$(ii) \leq \sum_{\xi=0}^{(1-\delta)b} \kappa_1 \kappa_2 e^{-[(C_H-1-\epsilon)\xi + (E_L-\epsilon)(b-\xi)]}. \tag{29}$$

We now distinguish two cases:

- $C_H - 1 > E_L$: In this case, we can bound the above probability as

$$(ii) \leq \kappa e^{-b(E_L-\epsilon)}, \quad \forall \epsilon > 0, \tag{30}$$

where $\kappa > 0$ is some constant.

- $C_H - 1 \leq E_L$: In this case,

$$(ii) \leq \kappa e^{-b(C_H-1-\epsilon)(1-\delta)}, \quad \forall \epsilon > 0. \tag{31}$$

Let us now put together the bounds on terms (i), (ii) and (iii) into Equation (21).

- 1) If $C_H - 1 > E_L$, we can show using (24), (25), and (30) that

$$\liminf_{b \rightarrow \infty} -\frac{1}{b} \log \mathbb{P} \{ q_L \geq b \} \geq (1-\delta)(E_L - \epsilon).$$

Since the above is true for each ϵ and δ , we get

$$\liminf_{b \rightarrow \infty} -\frac{1}{b} \log \mathbb{P} \{ q_L \geq b \} \geq E_L. \tag{32}$$

- 2) If $C_H - 1 \leq E_L$, we can show using (24), (25), and (31) that

$$\liminf_{b \rightarrow \infty} -\frac{1}{b} \log \mathbb{P} \{ q_L \geq b \} \geq (1-\delta)(C_H - 1 - \epsilon).$$

Since the above is true for each ϵ and δ , we get

$$\liminf_{b \rightarrow \infty} -\frac{1}{b} \log \mathbb{P} \{ q_L \geq b \} \geq C_H - 1. \tag{33}$$

Theorem 7 now follows from (32) and (33). \square

REFERENCES

- [1] V. Anantharam, "Scheduling strategies and long-range dependence," *Queueing Systems*, vol. 33, no. 1, pp. 73–89, 1999.
- [2] N. Bingham, C. Goldie, and J. Teugels, *Regular variation*. Cambridge university press Cambridge, 1987.
- [3] S. Borst, M. Mandjes, and M. van Uitert, "Generalized processor sharing with light-tailed and heavy-tailed input," *IEEE/ACM Transactions on Networking*, vol. 11, no. 5, p. 834, 2003.
- [4] S. Borst, O. Boxma, R. Núñez-Queija, and A. Zwart, "The impact of the service discipline on delay asymptotics," *Performance Evaluation*, vol. 54, no. 2, pp. 175–206, 2003.
- [5] O. Boxma and B. Zwart, "Tails in scheduling," *ACM SIGMETRICS Performance Evaluation Review*, vol. 34, no. 4, p. 20, 2007.
- [6] A. Brzezinski and E. Modiano, "Dynamic reconfiguration and routing algorithms for IP-over-WDM networks with stochastic traffic," *Journal of Lightwave Technology*, vol. 23, no. 10, p. 3188, 2005.
- [7] D. Cline, "Intermediate Regular and Π Variation," *Proceedings of the London Mathematical Society*, vol. 3, no. 3, p. 594, 1994.
- [8] A. Eryilmaz, R. Srikant, and J. Perkins, "Stable scheduling policies for fading wireless channels," *IEEE/ACM Transactions on Networking*, vol. 13, no. 2, pp. 411–424, 2005.
- [9] R. Gallager, *Discrete stochastic processes*. Kluwer, 1996.
- [10] A. Ganesh, N. O'Connell, and D. Wischik, *Big queues*. Springer Verlag, 2004.
- [11] M. Harchol-Balter, A. Scheller-Wolf, and A. Young, "Surprising results on task assignment in server farms with high-variability workloads," in *SIGMETRICS*. ACM, 2009, pp. 287–298.
- [12] D. Heath, S. Resnick, and G. Samorodnitsky, "Heavy tails and long range dependence in on/off processes and associated fluid models," *Mathematics of Operations Research*, vol. 23, no. 1, pp. 145–165, 1998.
- [13] K. Jagannathan, M. Markakis, E. Modiano, and J. Tsitsiklis, "Queue Length Asymptotics for Generalized Max-Weight Scheduling in the presence of Heavy-Tailed Traffic," *ArXiv e-prints*, Jul. 2010, arXiv:1007.4286v1 [cs.NI].
- [14] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Transactions on networking*, vol. 2, no. 1, pp. 1–15, 1994.
- [15] M. Markakis, E. Modiano, and J. Tsitsiklis, "Scheduling policies for single-hop networks with heavy-tailed traffic," in *Allerton Conference, Monticello IL*, 2009.
- [16] N. McKeown, A. Mekkittikul, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input-queued switch," *IEEE Transactions on Communications*, vol. 47, no. 8, pp. 1260–1267, 1999.
- [17] M. Neely, E. Modiano, and C. Rohrs, "Power and server allocation in a multi-beam satellite with time varying channels," in *IEEE INFOCOM*, vol. 3, 2002.
- [18] —, "Dynamic power allocation and routing for time varying wireless networks," in *IEEE INFOCOM*, vol. 1, 2003.
- [19] K. Park and W. Willinger, *Self-similar network traffic and performance evaluation*. John Wiley & Sons, Inc. New York, NY, USA, 2000.
- [20] C. Robert and J. Segers, "Tails of random sums of a heavy-tailed number of light-tailed terms," *Insurance: Mathematics and Economics*, vol. 43, no. 1, pp. 85–92, 2008.
- [21] K. Sigman, "A primer on heavy-tailed distributions," *Queueing Systems*, vol. 33, no. 1-3, pp. 261–275, 1999.
- [22] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936–1948, 1992.
- [23] —, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Transactions on Information Theory*, vol. 39, no. 2, pp. 466–478, 1993.
- [24] A. Wierman and B. Zwart, "Is tail-optimal scheduling possible?" Under Submission.