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

Krishna Jagannathan, Mihalis Markakis, Eytan Modiano, *Fellow, IEEE*, and John N. Tsitsiklis, *Fellow, IEEE*

***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 at least as heavy as 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 coefficient of the light queue distribution, among all nonidling 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.**

***Index Terms*—Heavy-tailed traffic, scheduling, throughput optimality.**

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 timescales, and only an average behavior is perceived at longer timescales. 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 timescales than can be modeled by any finite-state Markovian point process. Empirical evidence, such as the famous Bellcore study on self-similarity and long-range dependence in ethernet traffic [18], led 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. Furthermore, theoretical work such as [15], 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 [18] and elsewhere.

Many of the early queueing theoretic results for heavy-tailed traffic were obtained for the single server queue; see [5], [6], and [23] for surveys of these results. In [7], the authors study the tail behavior of the waiting time in an M/G/2 system, when one of the service time distributions is heavy-tailed and the other is exponential.

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 nonpreemptive service discipline leads to infinite expected delay when the traffic is sufficiently heavy-tailed. Furthermore, the asymptotic behavior of 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 [5], [28]. This is important, for example, in the context of scheduling jobs in server farms [14].

In the context of communication networks, a subset of the traffic flows may be well modeled using heavy-tailed processes, and the rest better modeled as light-tailed processes. For example, an Internet user might generate occasional file download requests with highly variable file sizes that can be modeled as being heavy-tailed. On the other hand, routine Web page loading, e-mail, and Twitter traffic are likely to be far less variable and thus better modeled as being 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. An important paper in this category is [4], 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. In a related paper [3], the authors obtain the asymptotic workload behavior under a general coupled-queues framework, which includes GPS as a special case.

Manuscript received February 03, 2011; revised August 11, 2011; accepted October 13, 2011; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor T. Bonald. Date of publication November 18, 2011; date of current version August 14, 2012. This work was supported by the NSF under Grants CNS-0626781, CNS-0915988, and CCF-0728554 and the ARO Muri under Grant W911NF-08-1-0238. This work was carried out as a part of K. Jagannathan's doctoral dissertation at MIT.

K. Jagannathan is with the Department of Electrical Engineering, Indian Institute of Technology (IIT) Madras, Chennai 600 036, India (e-mail: krishnaj@ee.iitm.ac.in).

M. Markakis, E. Modiano, and J. N. Tsitsiklis are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: mihalis@mit.edu; modiano@mit.edu; jnt@mit.edu).

Digital Object Identifier 10.1109/TNET.2011.2173553

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 [26], [27] 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 [20], satellites [21], wireless [22], and optical networks [8].

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 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 [19], 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. Furthermore, 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 steady-state queue lengths*.

A surprising outcome of our asymptotic characterization is that the "plain" max-weight scheduling policy induces the worst possible decay rate on the light queue tail distribution. 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 an often 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.

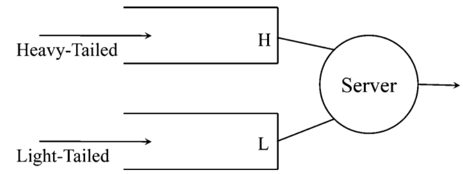

Fig. 1. System of two parallel queues, with one of them fed with heavy-tailed traffic.

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 policy, 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 VII respectively contain our asymptotic results for max-weight-α scheduling and the LMW policy. We conclude the paper in Section VIII. A shorter version of this work appeared in [17].

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. In the sequel, we will sometimes refer to the collection of packets that arrive at a given time as a *burst*. We will refer to the queues H and L as the heavy and light queues, respectively. The queues are assumed to be always connected to the server. Let $q_H(t)$ and $q_L(t)$, respectively, denote the number of packets in H and L at the beginning of slot t (i.e., before the slot t arrivals), 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 begin by defining some properties of tail distributions of nonnegative 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 subexponentials, have been defined and studied in the past [25]. We now review some definitions and existing results on some relevant classes of heavy-tailed distributions. In the remainder of this section, X will denote a nonnegative random variable, with complementary distribution function $\bar{F}(x) = \mathbb{P}\{X > x\}$. For the most part, we adhere to the terminology in [2] and [9].

Notation: If $f(x)$ and $g(x)$ are positive functions defined on $[0,\infty)$, we write $f(x) \sim g(x)$ to mean

$$\lim_{x\to\infty} \frac{f(x)}{g(x)} = 1.$$

Similarly, $f(x) \gtrsim g(x)$ means

$$\liminf_{x\to\infty} \frac{f(x)}{g(x)} \geq 1.$$

Definition 2:

1) $\bar{F}(x)$ has a *regularly varying* tail of index ν, denoted by $\bar{F} \in \mathcal{R}(\nu)$, if

$$\lim_{x\to\infty} \frac{\bar{F}(kx)}{\bar{F}(x)} = k^{-\nu} \qquad \forall\, k > 0.$$

2) $\bar{F}(x)$ is *extended-regularly varying*, denoted by $\bar{F} \in \mathcal{ER}$, if for some real $c, d > 0$, and $\Gamma > 1$

$$k^{-d} \leq \liminf_{x\to\infty} \frac{\bar{F}(kx)}{\bar{F}(x)} \leq \limsup_{x\to\infty} \frac{\bar{F}(kx)}{\bar{F}(x)} \leq k^{-c} \qquad \forall k \in [1,\Gamma].$$

3) $\bar{F}(x)$ is *intermediate-regularly varying*, denoted by $\bar{F} \in \mathcal{IR}$, if

$$\lim_{k\downarrow 1} \liminf_{x\to\infty} \frac{\bar{F}(kx)}{\bar{F}(x)} = \lim_{k\downarrow 1} \limsup_{x\to\infty} \frac{\bar{F}(kx)}{\bar{F}(x)} = 1.$$

4) $\bar{F}(x)$ is *order-regularly varying*, denoted by $\bar{F} \in \mathcal{OR}$, if for some $\Gamma > 1$

$$\liminf_{x\to\infty} \frac{\bar{F}(kx)}{\bar{F}(x)} > 0 \qquad \forall k \in [1,\Gamma].$$

It is easy to see from the definitions that $\mathcal{R} \subset \mathcal{ER} \subset \mathcal{IR} \subset \mathcal{OR}$. In fact, the containments are proper, as shown in [9]. Intuitively, $\mathcal{R}$ is the class of distributions with tails that decay according to a power-law with parameter ν. Indeed, it can be shown [11] that

$$\bar{F} \in \mathcal{R}(\nu) \iff \bar{F}(x) = U(x)x^{-\nu}$$

where $U(x)$ is a *slowly varying* function, i.e., a function that satisfies $U(kx) \sim U(x)$, $\forall k > 0$. The other three classes are increasingly more general, but as we shall see, they all correspond to distributions that are asymptotically heavier than some power-law curve. In what follows, a statement such as $X \in \mathcal{IR}$ should be construed to mean $\mathbb{P}\{X > x\} \in \mathcal{IR}$.

Next, we define the lower and upper orders of a distribution.

Definition 3:

1) The *lower order* of $\bar{F}(x)$ is defined by

$$\xi(\bar{F}) = \liminf_{x\to\infty} -\frac{\log \bar{F}(x)}{\log x}.$$

2) The *upper order* of $\bar{F}(x)$ is defined by

$$\rho(\bar{F}) = \limsup_{x\to\infty} -\frac{\log \bar{F}(x)}{\log x}.$$

It can be shown that for regularly varying distributions, the upper and lower orders coincide with the index ν. It also turns out that both orders are finite for the class $\mathcal{OR}$, as asserted in the following.

Proposition 1: $\rho(\bar{F}) < \infty$ for every $\bar{F} \in \mathcal{OR}$.

Proof: Follows from [2, Theorem 2.1.7 and Proposition 2.2.5].

The following result, which is a consequence of Proposition 1, states that every $\bar{F} \in \mathcal{OR}$ is asymptotically heavier than a power-law curve.

Proposition 2: Let $\bar{F} \in \mathcal{OR}$. Then, for each $\rho > \rho(\bar{F})$, we have $x^{-\rho} = o(\bar{F}(x))$ as $x \to \infty$.

Proof: See [24, eq. (2.4)].

Definitions 2 and 3 deal with asymptotic tail probabilities of a random variable. Next, we introduce the notion of tail coefficient, which is a *moment* property.

Definition 4: The *tail coefficient* of a random variable X is defined by

$$C_X = \sup\{c \geq 0 \mid \mathbb{E}[X^c] < \infty\}.$$

In other 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). The next result shows that the tail coefficient and order are, in fact, closely related parameters.

Proposition 3: The tail coefficient of X is equal to the lower order of $\bar{F}(x)$.[1]

Proof: Suppose first that the lower order is infinite, so that for any $s > 0$, we can find an x large enough such that

$$-\frac{\log \mathbb{P}\{X > x\}}{\log x} > s.$$

Thus, for large enough x, we have

$$\mathbb{P}\{X > x\} < x^{-s} \qquad \forall\, s > 0.$$

This implies $\mathbb{E}[X^c] < \infty$ for all $c > 0$. Therefore, the tail coefficient of X is also infinite.

Next, suppose that $\xi(\bar{F}) \in [0, \infty)$. We will show that: 1) $\mathbb{E}[X^c] < \infty$ for all $c < \xi(\bar{F})$; and 2) $\mathbb{E}[X^c] = \infty$ for all $c > \xi(\bar{F})$. Regarding 1), we note that there is nothing to be shown if $\xi(\bar{F}) = 0$. If $\xi(\bar{F}) > 0$, we argue as above that for large enough x, we have $\mathbb{P}\{X > x\} < x^{-s}$ when $s < \xi(\bar{F})$. Thus, $\mathbb{E}[X^c] < \infty$ for all $c < \xi(\bar{F})$. To show 2), let us consider some s such that $c > s > \xi(\bar{F})$. By the definition of $\xi(\bar{F})$, there exists a sequence $\{x_i\}$ that increases to infinity as $i \to \infty$, such that

$$-\frac{\log \mathbb{P}\{X > x_i\}}{\log x_i} \le s \qquad \forall\, i, \text{ i.e., } \mathbb{P}\{X > x_i\} \ge {x_i}^{-s},\ \forall\, i.$$

Therefore

$$\begin{aligned}\mathbb{E}[X^c] = \int_0^\infty x^c \mathrm{d}F_X(x) &\ge \int_{x_i}^\infty x^c \mathrm{d}F_X(x)\\ &\ge x_i^c \mathbb{P}\{X > x_i\} \ge x_i^c x_i^{-s} \qquad \forall\, i\end{aligned}$$

from which it follows that $\mathbb{E}[X^c] = \infty$. Therefore, the tail coefficient of X is equal to $\xi(\bar{F})$. □

We remark that Proposition 3 holds for *any* random variable, regardless of its regularity properties. Finally, we show that any distribution in $\mathcal{OR}$ necessarily has a finite tail coefficient.

Proposition 4: If $X \in \mathcal{OR}$, then X has a finite tail coefficient.

Proof: From Proposition 1, the upper order is finite: $\rho(\bar{F}) < \infty$. Thus, the lower order $\xi(\bar{F})$ is also finite. Since the lower order equals the tail coefficient (Proposition 3), the result follows. □

B. Assumptions on the Arrival Processes

We are now ready to state the precise assumptions on the arrivals processes.

1) The arrival processes to the two queues are independent of each other. Furthermore, $H(t)$ and $L(t)$ are independent of the past history until time t.
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 stable[2] under any nonidling 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 nonzero probability mass at zero since $\lambda_H < 1$. We define H_+ as a random variable distributed according to the conditional distribution of $H(t)$, given that $H(t) > 0$. Specifically

$$\mathbb{P}\{H_+ = m\} = \frac{\mathbb{P}\{H(\cdot) = m\}}{1 - \mathbb{P}\{H(\cdot) = 0\}}, \qquad m = 1, 2, \ldots.$$

Note that H_+ has tail coefficient equal to C_H and inherits any regularity property of $H(\cdot)$.

Now consider a discrete-time renewal process with interrenewal times distributed as H_+. Let $H_{\mathrm{R}} \in \{1, 2, \ldots\}$ denote the residual random variable, and $H_{\mathrm{A}} \in \{0, 1, \ldots\}$ the age of the renewal process, in steady state [12].[3] The joint distribution of the residual and the age can be derived using basic renewal theory

$$\mathbb{P}\{H_{\mathrm{R}} = k, H_{\mathrm{A}} = l\} = \frac{\mathbb{P}\{H_+ = k + l\}}{\mathbb{E}[H_+]} \tag{1}$$

$k \in \{1, 2 \cdots\}$, $l \in \{0, 1, \ldots\}$. The marginals of H_{R} and H_{A} can be derived from (1)

$$\mathbb{P}\{H_{\mathrm{R}} = k\} = \frac{\mathbb{P}\{H_+ \ge k\}}{\mathbb{E}[H_+]}, \qquad k \in \{1, 2, \ldots\}. \tag{2}$$

$$\mathbb{P}\{H_{\mathrm{A}} = k\} = \frac{\mathbb{P}\{H_+ > k\}}{\mathbb{E}[H_+]}, \qquad k \in \{0, 1, \ldots\}. \tag{3}$$

Next, let us invoke a useful result from the literature.

Lemma 1: If $H(\cdot) \in \mathcal{OR}$, then $H_{\mathrm{R}} \in \mathcal{ER}$ and

$$\sup_n \frac{n\mathbb{P}\{H_+ > n\}}{\mathbb{P}\{H_{\mathrm{R}} > n\}} < \infty. \tag{4}$$

A corresponding result also holds for the age H_{A}.

Proof: See [9, Lemma 4.2(i)]. □

Using the above, we prove the important result that the residual distribution is *one order heavier* than the original distribution.

Proposition 5: If $H(\cdot) \in \mathcal{OR}$ has tail coefficient equal to C_H, then H_{R} and H_{A} have tail coefficient equal to $C_H - 1$.

Proof: According to (4), we have, for all a and some real χ

$$-\log \mathbb{P}\{H_{\mathrm{R}} > a\} \le -\log a - \log \mathbb{P}\{H_+ > a\} + \chi.$$

[1]The first author is grateful to J. Nair (Caltech, Pasadena, CA) for suggesting a proof of Proposition 3 via a personal communication.

[2]The notion of stability used here is the positive recurrence of the system occupancy Markov chain.

[3]We define the residual time and age so 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.

Let us now consider the lower order of H_{R}

$$\liminf_{a\to\infty} -\frac{\log \mathbb{P}\{H_{\mathrm{R}} > a\}}{\log a} \leq \liminf_{a\to\infty} \frac{-\log a - \log \mathbb{P}\{H_+ > a\} + \chi}{\log a} = C_H - 1.$$

In the last step above, we have used the tail coefficient of H_+. Since the lower order of H_{R} equals its tail coefficient (Proposition 3), the above relation shows that the tail coefficient of H_{R} is *at most* $C_H - 1$.

Next, to show the opposite inequality, let us consider the *duration* random variable, defined as

$$H_{\mathrm{D}} = H_{\mathrm{R}} + H_{\mathrm{A}}.$$

Using the joint distribution (1), we can obtain the marginal of H_{D} as

$$\mathbb{P}\{H_{\mathrm{D}} = k\} = \frac{k\mathbb{P}\{H_+ = k\}}{\mathbb{E}[H_+]}, \qquad k \in \{1, 2, \ldots\}.$$

Thus, for any $\epsilon > 0$, the $C_H - 1 - \epsilon$ moment of H_{D} is finite

$$\mathbb{E}\left[H_{\mathrm{D}}^{C_H-1-\epsilon}\right] = \sum_{k\geq 1} \frac{k^{C_H-\epsilon}\mathbb{P}\{H_+ = k\}}{\mathbb{E}[H_+]} = \frac{\mathbb{E}\left[H_+^{C_H-\epsilon}\right]}{\mathbb{E}[H_+]} < \infty.$$

Since H_{R} is stochastically dominated by H_{D}, it is immediate that $\mathbb{E}\left[H_{\mathrm{R}}^{C_H-1-\epsilon}\right] < \infty$. Therefore, the tail coefficient of H_{R} is *at least* $C_H - 1$, and the proposition is proved. □

IV. 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. Furthermore, we show that giving priority to the light queue ensures the best possible asymptotic decay for *both* queue-length distributions.

A. Priority for H

In this policy, H receives service whenever it is nonempty, and L receives service only when H is empty. It should be intuitively clear at the outset that this policy is bound to have 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.

Our first result shows that the steady-state heavy queue occupancy is one order heavier than its input distribution.

Theorem 1: Under priority scheduling for H, the steady-state queue occupancy distribution of the heavy queue satisfies the following bounds.

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

$$\mathbb{P}\{q_H > b\} < \kappa_H(\epsilon) b^{-(C_H-1-\epsilon)} \qquad \forall\, b. \tag{5}$$

2)

$$\mathbb{P}\{q_H > b\} \geq \lambda_H \mathbb{P}\{H_{\mathrm{R}} > b\} \qquad \forall\, b. \tag{6}$$

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

$$\mathbb{E}\left[q_H^{C_H-1-\epsilon}\right] < \infty \tag{7}$$

and

$$\mathbb{E}\left[q_H^{C_H-1+\epsilon}\right] = \infty. \tag{8}$$

Proof: Equation (7) can be shown using a straightforward Lyapunov argument, along the lines of [19, Proposition 6]. Equation (5) follows from (7) and the Markov inequality.

Next, to show (6), we consider a time instant t at steady state, and write

$$\mathbb{P}\{q_H(t) > b\} = \mathbb{P}\{q_H(t) > b \mid q_H(t) > 0\}\,\mathbb{P}\{q_H(t) > 0\} = \lambda_H \mathbb{P}\{q_H(t) > b \mid q_H(t) > 0\}.$$

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 $\mathbb{P}\{q_H(t) > b \mid q_H(t) > 0\}$. Conditioned on H being nonempty, denote by $\tilde{B}(t)$ the number of packets that belong to the burst in service that are still in queue at time t. Then, clearly, $q_H(t) \geq \tilde{B}(t)$, from which $\mathbb{P}\{q_H(t) > b \mid q_H(t) > 0\} \geq \mathbb{P}\left\{\tilde{B}(t) > b\right\}$. Now, since the H queue receives service whenever it is nonempty, 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, $\mathbb{P}\{q_H(t) > b \mid q_H(t) > 0\} \geq \mathbb{P}\{H_{\mathrm{R}} > b\}$ and (6) follows. Finally, (8) follows from (6) and Proposition 5. □

When the distribution of $H(\cdot)$ is regularly varying, the lower bound (6) takes on a power-law form that agrees with the upper bound (5).

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

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

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 6: Under any scheduling policy, q_H is heavy-tailed with tail coefficient at most $C_H - 1$. That is, (8) holds for all scheduling policies.

Proof: The queue occupancy q_H under any policy stochastically dominates the queue occupancy under priority for H. Therefore, the lower bounds (6) and (8) hold for all 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.

Theorem 2: Under priority for H, q_L is heavy-tailed with tail coefficient $C_H - 1$. Furthermore, the tail distribution $\mathbb{P}\{q_L > b\}$ satisfies the following asymptotic bounds.

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

$$\mathbb{P}\{q_L > b\} < \kappa_L(\epsilon) b^{-(C_H - 1 - \epsilon)}. \tag{9}$$

2) If $H(\cdot) \in \mathcal{OR}$, then

$$\mathbb{P}\{q_L > b\} \gtrsim \lambda_H \mathbb{P}\left\{H_{\mathrm{A}} > \frac{b}{\lambda_L}\right\}. \tag{10}$$

Proof: The upper bound (9) is a special case of Theorem 4 given in the next section. Let us show (10). Notice first that the lower bound (10) is asymptotic, unlike (6), which is exact. As before, let us consider a time t at steady state and write using Little's law

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

Let us denote by $\tilde{A}(t)$ the number of slots that the current head-of-line burst has been in service. Clearly then, L has not received any service in the interval $[t - \tilde{A}(t), t]$ and has kept all the arrivals that occurred during the interval. Thus, conditioned on H being nonempty, $q_L(t) \geq \sum_{\sigma = t - \tilde{A}(t)}^{t} L(\sigma)$. Next, it can be seen that in steady state, $\tilde{A}(t)$ is distributed as the age variable H_{A}. Putting everything together, we can write

$$\begin{aligned}\mathbb{P}\{q_L > b\} &\geq \lambda_H \mathbb{P}\{q_L(t) > b \mid q_H(t) > 0\} \\ &\geq \lambda_H \mathbb{P}\left\{\sum_{i=1}^{H_{\mathrm{A}}} L(i) > b\right\}. \end{aligned} \tag{11}$$

Next, since $H(\cdot) \in \mathcal{OR}$, Lemma 1 implies that $H_{\mathrm{A}} \in \mathcal{ER} \subset \mathcal{IR}$. We can therefore invoke Lemma 4 in the Appendix to write

$$\mathbb{P}\left\{\sum_{i=1}^{H_{\mathrm{A}}} L(i) > b\right\} \sim \mathbb{P}\left\{H_{\mathrm{A}} > \frac{b}{\lambda_L}\right\}. \tag{12}$$

Finally, (10) follows from (11) and (12). □

When $H(\cdot)$ is regularly varying, the lower bound (10) takes on a power-law form that matches the upper bound (9).

B. Priority for L

We now study the policy that serves L whenever it is nonempty 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 nonidling 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 [13] 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

$$E_L = \sup\{\theta \mid \Lambda_L(\theta) - \theta < 0\}. \tag{13}$$

Proposition 7: Under priority for L, q_L satisfies the large deviation principle (LDP)

$$\lim_{b \to \infty} -\frac{1}{b} \log \mathbb{P}\{q_L > b\} = E_L. \tag{14}$$

In other 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

$$E_L = \inf_{a > 0} \frac{1}{a} \Lambda_L^*(1 + a) \tag{15}$$

where $\Lambda_L^*(\cdot)$ is the Fenchel–Legendre transform [13] of $\Lambda_L(\theta)$.

It is clear that the priority policy for L gives the best possible asymptotic behavior for the light queue and the worst possible treatment for the heavy queue. Perhaps surprisingly however, the heavy queue tail under priority for L is asymptotically no worse than that under priority for H.

Proposition 8: Under priority for L, q_H is heavy-tailed with tail coefficient $C_H - 1$.

Proof: This is a special case of Theorem 4, given in Section V. □

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

Proposition 9: Under any nonidling scheduling policy, q_H has a tail coefficient of at least $C_H - 1$. That is, (7) holds for all nonidling scheduling policies.

Proof: The queue occupancy q_H under any other policy is stochastically dominated by the queue occupancy under priority for L. □

Propositions 6 and 9 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 3: Under *any* nonidling scheduling policy, q_H is heavy-tailed with tail coefficient equal to $C_H - 1$. Furthermore, $\mathbb{P}\{q_H > b\}$ satisfies bounds of the form (5) and (6) under all nonidling 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 nonidling, 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

$$q_L(t)^{\alpha_L} \gtrless q_H(t)^{\alpha_H}$$

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. The case where $\alpha_L/\alpha_H > 1$ 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. Furthermore, our distributional bounds on the light queue occupancy shed further light and refine the moment results derived in [19] for max-weight-α scheduling.

A. Upper Bound

In this section, we derive two different upper bounds on the overflow probability $\mathbb{P}\{q_L > b\}$ (Theorems 4 and 6) that both hold under max-weight-α scheduling. Depending on the values of α_H and α_L, either bound can be the tighter one. The first upper bound holds for all nonidling policies, including max-weight-α scheduling.

Theorem 4: Under *any* nonidling policy, and for every $\epsilon > 0$, there exists a constant $\kappa_1(\epsilon) > 0$, such that

$$\mathbb{E}\left[q_L^{C_H-1-\epsilon}\right] < \infty \tag{16}$$

and

$$\mathbb{P}\{q_L > b\} < \kappa_1(\epsilon) b^{-(C_H-1-\epsilon)}. \tag{17}$$

Proof: Let us combine the two queues into one, and consider the sum input process $H(t) + L(t)$ feeding the composite queue. The server serves one packet from the composite queue in each slot. Under any nonidling policy in the original system, the occupancy of the composite queue is given by $q = q_H + q_L$. Lemma 2 in the Appendix shows that the combined input has tail coefficient equal to C_H. The composite queue is therefore a G/D/1 queue with input tail coefficient C_H. For such a queue, it can be shown that

$$\mathbb{E}\left[q^{C_H-1-\epsilon}\right] < \infty. \tag{18}$$

This is, in fact, a direct consequence of Theorem 1.

Thus, in terms of the queue lengths in the original system, we have

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

from which it is immediate that $\mathbb{E}\left[q_L^{C_H-1-\epsilon}\right] < \infty$. This proves (16). To show (17), we use the Markov inequality to write

$$\begin{aligned}\mathbb{P}\{q_L > b\} &= \mathbb{P}\left\{q_L^{C_H-1-\epsilon} > b^{C_H-1-\epsilon}\right\} \\ &< \frac{\mathbb{E}\left[q_L^{C_H-1-\epsilon}\right]}{b^{C_H-1-\epsilon}} < \kappa_1(\epsilon) b^{-(C_H-1-\epsilon)}.\end{aligned}$$

□

The above result asserts that the tail coefficient of q_L is at least $C_H - 1$ under any nonidling 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 [19]. We first make an elementary but useful observation.

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 [19], which is proved therein using a suitable Lyapynov function.

Theorem 5: If max-weight-α scheduling is performed with $0 < \alpha_H < C_H - 1$, then, for *any* $\alpha_L > 0$, we have $\mathbb{E}\left[q_L^{\alpha_L}\right] < \infty$.

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 6: Define

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

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

$$\mathbb{E}\left[q_L^{\gamma-\epsilon}\right] < \infty \tag{19}$$

and

$$\mathbb{P}\{q_L > b\} < \kappa_2(\epsilon) b^{-(\gamma-\epsilon)}. \tag{20}$$

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 our earlier observation, this policy is identical to the original max-weight-α policy. Next, since $\beta\alpha_H < C_H - 1$, Theorem 5 applies, and we have $\mathbb{E}\left[q_L^{\beta\alpha_L}\right] = \mathbb{E}\left[q_L^{\gamma-\epsilon}\right] < \infty$, which proves (19). Finally, (20) can be proved using (19) and the Markov inequality. □

The above theorem asserts that the tail coefficient of q_L is at least γ under the max-weight-α policy. We remark that Theorems 4 and 6 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.

1) $\alpha_L/\alpha_H \leq 1$: This is the regime where the light queue is given lesser priority, when compared to the heavy queue. In this case, Theorem 4 yields a stronger bound.
2) $\alpha_L/\alpha_H > 1$: This is the regime where the light queue is given more priority compared to the heavy queue. In this case, Theorem 6 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 Section V-B.

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 7: Let $H(\cdot) \in \mathcal{OR}$. Then, 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 $\alpha_L/\alpha_H < 1$

$$\mathbb{P}\{q_L \geq b\} \gtrsim \lambda_H \mathbb{P}\left\{H_{\mathrm{R}} \geq \frac{b}{\lambda_L}\right\}. \tag{21}$$

2) If $\alpha_L/\alpha_H = 1$

$$\mathbb{P}\{q_L \geq b\} \gtrsim \lambda_H \mathbb{P}\left\{H_{\mathrm{R}} \geq b\left(1+\frac{1}{\lambda_L}\right)\right\}. \tag{22}$$

3) If $\alpha_L/\alpha_H > 1$

$$\mathbb{P}\{q_L \geq b\} \gtrsim \lambda_H \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}. \tag{23}$$

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 (17) and (20).

Corollary 2: Suppose that $H(\cdot) \in \mathcal{R}(C_H)$. 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 $\alpha_L/\alpha_H \leq 1$

$$\mathbb{P}\{q_L \geq b\} \gtrsim U(b) b^{-(C_H-1)}. \tag{24}$$

2) If $\alpha_L/\alpha_H > 1$

$$\mathbb{P}\{q_L \geq b\} \gtrsim U(b) b^{-\gamma} \tag{25}$$

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

Corollary 2 follows from Theorem 7 together with Karamata's theorem for regularly varying functions [2, Section 1.6].

It takes several steps to prove Theorem 7; we start by defining and studying a related fictitious queueing system.

C. Fictitious System

We introduce a fictitious system that consists of two queues fed by the *same input processes* that feed the original system. In the fictitious system, let us call the queues fed by heavy-tailed and light-tailed traffic $\tilde{H}$ and $\tilde{L}$, respectively. The fictitious system operates under the following service discipline.

Service for the Fictitious System: The queue $\tilde{H}$ receives service in every time slot. The queue $\tilde{L}$ receives service at time t if and only if $q_{\tilde{L}}(t)^{\alpha_L} \geq q_{\tilde{H}}(t)^{\alpha_H}$.

Note that if $\tilde{L}$ receives service and $\tilde{H}$ is nonempty, *two* packets are served from the fictitious system. Also, $\tilde{H}$ is just a discrete time $G/D/1$ queue since it receives service at every time slot. We now show a simple result that asserts that the light queue in the original system is "longer" than in the fictitious system.

Proposition 10: Suppose that we drive both the original and the fictitious system with a common sample path of the arrival processes. Then, $q_{\tilde{L}}(t) \leq q_L(t)$, for all t. In particular, for every $b > 0$, we have

$$\mathbb{P}\{q_L \geq b\} \geq \mathbb{P}\{q_{\tilde{L}} \geq b\}.$$

Proof: We will assume the contrary and arrive at a contradiction. Suppose that $q_{\tilde{L}}(0) = q_L(0)$, and that for some time $t > 0$, $q_{\tilde{L}}(t) > q_L(t)$. Let $\tau > 0$ be the first time when $q_{\tilde{L}}(\tau) > q_L(\tau)$. It is then necessary that $q_{\tilde{L}}(\tau-1) = q_L(\tau-1)$ since no more than one packet is served from a queue in each slot. Next, $q_{\tilde{L}}(\tau-1) = q_L(\tau-1)$ and $q_{\tilde{L}}(\tau) > q_L(\tau)$ together imply that L received service at time $\tau-1$, but $\tilde{L}$ did not. This is possible only if $q_H(\tau-1) < q_{\tilde{H}}(\tau-1)$, which is a contradiction, since $\tilde{H}$ receives service in each slot. □

Next, we show that the distribution of $q_{\tilde{L}}$ satisfies the lower bounds in (21)–(23). Theorem 7 then follows, in light of Proposition 10.

Theorem 8: In the fictitious system, the distribution of $q_{\tilde{L}}$ is asymptotically lower-bounded as follows.

1) If $\alpha_L/\alpha_H < 1$

$$\mathbb{P}\{q_{\tilde{L}} \geq b\} \gtrsim \lambda_H \mathbb{P}\left\{H_{\mathrm{R}} \geq \frac{b}{\lambda_L}\right\}. \tag{26}$$

2) If $\alpha_L/\alpha_H = 1$

$$\mathbb{P}\{q_{\tilde{L}} \geq b\} \gtrsim \lambda_H \mathbb{P}\left\{H_{\mathrm{R}} \geq b\left(1+\frac{1}{\lambda_L}\right)\right\}. \tag{27}$$

3) If $\alpha_L/\alpha_H > 1$

$$\mathbb{P}\{q_{\tilde{L}} \geq b\} \gtrsim \lambda_H \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}. \tag{28}$$

Proof: Consider the fictitious system in steady state, and let us fix a particular time t. Since the heavy queue in the fictitious system receives service in each slot, the steady-state distribution satisfies $\mathbb{P}\{q_{\tilde{H}} > 0\} = \lambda_H$ by Little's law. Therefore, we have the lower bound

$$\mathbb{P}\{q_{\tilde{L}} \geq b\} \geq \lambda_H \mathbb{P}\{q_{\tilde{L}} \geq b | q_{\tilde{H}} > 0\}.$$

In the rest of the proof, we will lower-bound the above conditional probability.

Indeed, conditioned on $q_{\tilde{H}} > 0$, denote as before by $\tilde{B}(t)$ the number of packets that belong to the head-of-line burst that still remain in queue $\tilde{H}$ at time t. Similarly, denote by $\tilde{A}(t)$ the number of packets from the head-of-line (HoL) burst that have already been served by time t. Since $\tilde{H}$ is served in every time slot, $\tilde{A}(t)$ also denotes the number of time slots that the HoL burst has been in service at $\tilde{H}$.

The remainder of the proof shows that $q_{\tilde{L}}(t)$ stochastically dominates a particular heavy-tailed random variable. Indeed, at the instant t, there are two possibilities:

a) $q_{\tilde{L}}(t)^{\alpha_L} \geq \tilde{B}(t)^{\alpha_H}$;

b) $q_{\tilde{L}}(t)^{\alpha_L} < \tilde{B}(t)^{\alpha_H}$.

Let us take a closer look at case b) in the following proposition.

Proposition 11: Suppose that

$$q_{\tilde{L}}(t)^{\alpha_L} < \tilde{B}(t)^{\alpha_H}.$$

Let $\sigma \leq t$ be the instant before t that $\tilde{L}$ last received service. Then, the head-of-line burst at time t in $\tilde{H}$ arrived *after* the instant σ.

Proof: We have

$$q_{\tilde{H}}(\sigma)^{\alpha_H} \leq q_{\tilde{L}}(\sigma)^{\alpha_L} \leq q_{\tilde{L}}(t)^{\alpha_L} < \tilde{B}(t)^{\alpha_H}.$$

The first inequality holds because $\tilde{L}$ received service at σ, the second inequality is true since $\tilde{L}$ does not receive service between σ and t, and the final inequality is from the hypothesis.

We have shown that $q_{\tilde{H}}(\sigma) < \tilde{B}(t)$, and hence the HoL burst could not have arrived by the time slot σ. ☐

The above proposition implies that if case b) holds, $\tilde{L}$ has not received service ever since the HoL burst arrived at $\tilde{H}$. In particular, $\tilde{L}$ has not received service for $\tilde{A}(t)$ time slots, and it accumulates all arrivals that occur during the interval $[t - \tilde{A}(t), t]$. Let us denote the number of arrivals to $\tilde{L}$ during this interval as

$$S_{\tilde{A}} = \sum_{i=t-\tilde{A}(t)}^{t} L(i).$$

In this notation, our argument above implies that if case b) holds, then $q_{\tilde{L}}(t) \geq S_{\tilde{A}}$. Putting this together with case a), we can conclude that

$$q_{\tilde{L}}(t) \geq \min(\tilde{B}(t)^{\alpha_H/\alpha_L},\ S_{\tilde{A}}). \tag{29}$$

Therefore

$$\mathbb{P}\left\{q_{\tilde{L}}(t) \geq b\right\} \geq \lambda_H \mathbb{P}\left\{\tilde{B}(t)^{\alpha_H/\alpha_L} \geq b,\ S_{\tilde{A}} \geq b\right\}. \tag{30}$$

Recall now that in steady state, $\tilde{B}(t)$ is distributed as H_{R}, and $\tilde{A}(t)$ is distributed as H_{A}. Therefore, the above bound can be written as

$$\mathbb{P}\left\{q_{\tilde{L}} \geq b\right\} \geq \lambda_H \mathbb{P}\left\{H_{\mathrm{R}}^{\alpha_H/\alpha_L} \geq b,\ \sum_{i=1}^{H_{\mathrm{A}}} L(i) \geq b\right\}. \tag{31}$$

Next, Lemma 5 (in the Appendix) shows that

$$\mathbb{P}\left\{H_{\mathrm{R}}^{\alpha_H/\alpha_L} \geq b,\ \sum_{i=1}^{H_{\mathrm{A}}} L(i) \geq b\right\} \sim \begin{cases} \mathbb{P}\left\{H_{\mathrm{R}} \geq \frac{b}{\lambda_L}\right\}, & \frac{\alpha_L}{\alpha_H} < 1 \\ \mathbb{P}\left\{H_{\mathrm{R}} \geq b + \frac{b}{\lambda_L}\right\}, & \frac{\alpha_L}{\alpha_H} = 1 \\ \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}, & \frac{\alpha_L}{\alpha_H} > 1. \end{cases}$$

Notice that the assumption $H(\cdot) \in \mathcal{OR}$ is used in the proof of Lemma 5.

Theorem 8 now follows from the above asymptotic relation and (31). ☐

Proof of Theorem 7: The result follows from Theorem 8 and Proposition 10. ☐

VI. Tail Coefficient of q_L

In this section, we characterize the exact tail coefficient of the light queue distribution under max-weight-α scheduling. In particular, we show that the upper bound (16) is tight if $\alpha_L/\alpha_H \leq 1$ and (19) is tight if $\alpha_L/\alpha_H > 1$.

Theorem 9: The tail coefficient of the steady-state queue length q_L of the light queue is given by the following:

i) $C_H - 1$ if $\alpha_L/\alpha_H \leq 1$;

ii) $\gamma = (C_H - 1)\alpha_L/\alpha_H$ if $\alpha_L/\alpha_H > 1$.

Proof: Consider first the case $\alpha_L/\alpha_H \leq 1$. The lower order (Definition 3) of q_L can be upper bounded using (21) or (22) as follows:

$$\begin{aligned}\liminf_{b\to\infty} -\frac{\log \mathbb{P}\{q_L > b\}}{\log b} &\leq \liminf_{b\to\infty} -\frac{\log \lambda_H \mathbb{P}\left\{H_{\mathrm{R}} \geq \frac{b}{\lambda_L}\right\}}{\log b} \\ &= \liminf_{a\to\infty} -\frac{\log \mathbb{P}\{H_{\mathrm{R}} \geq a\}}{\log a} = C_H - 1.\end{aligned}$$

The last step is from Proposition 5. The above equation shows that the tail coefficient of q_L is at most $C_H - 1$. However, it is evident from (16) that the tail coefficient of q_L is *at least* $C_H - 1$. Therefore, the tail coefficient of q_L equals $C_H - 1$ for $\frac{\alpha_L}{\alpha_H} \leq 1$. This proves case i) of the theorem.

Next, suppose that $\alpha_L/\alpha_H > 1$. Using (23), we can upper-bound the lower order of q_L as

$$\begin{aligned}\liminf_{b\to\infty} -\frac{\log \mathbb{P}\{q_L > b\}}{\log b} &\leq \liminf_{b\to\infty} -\frac{\log \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}}{\log b} \\ &= \frac{\alpha_L}{\alpha_H} \liminf_{a\to\infty} \frac{-\log \mathbb{P}\{H_{\mathrm{R}} \geq a\}}{\log a} \\ &= \frac{\alpha_L}{\alpha_H}(C_H - 1). \end{aligned} \tag{32}$$

Equation (32) shows that the tail coefficient of q_L is at most γ. However, it is evident from (19) that the tail coefficient of q_L is *at least* γ. Therefore, the tail coefficient of q_L equals $\gamma = (C_H - 1)\alpha_L/\alpha_H$ if $\alpha_L/\alpha_H > 1$. This proves case ii) of the theorem. ☐

In Fig. 2, we show the tail coefficient of q_L as a function of the ratio α_L/α_H. The tail coefficient stays constant at the value $C_H - 1$ as α_L/α_H varies from 0 to 1. Recall that $\alpha_L/\alpha_H = 1$ 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. In other words, max-weight scheduling leads to the *worst possible* asymptotic behavior for the light queue among all nonidling policies in the sense that it leads to the smallest possible tail coefficient for q_L. 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 9 that the expected

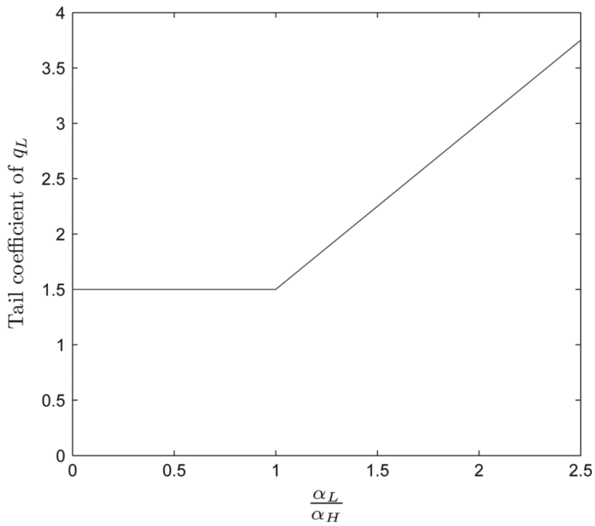

Fig. 2. Tail coefficient of q_L under max-weight α scheduling, as a function of α_L/α_H, for $C_H = 2.5$.

delay in the light queue is infinite under max-weight scheduling. Thus, our result can be viewed as a generalization of [19, Proposition 5].

VII. Log-Max-Weight Scheduling

We showed in Theorem 9 that the light queue occupancy distribution is necessarily heavy-tailed with a finite tail coefficient under max-weight-α scheduling. On the other hand, the priority for L policy that ensures the best possible asymptotic behavior for both queues suffers from possible instability effects in more general queueing networks.

In this section, we analyze the log-max-weight (LMW) scheduling policy. We show that the *light queue distribution is light-tailed* under LMW scheduling, i.e., that $\mathbb{P}\{q_L > b\}$ decays exponentially fast in b. However, unlike the priority for L policy, LMW scheduling can be shown to be throughput-optimal in very general settings [10]. 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 deemphasis of the heavy queue with respect to the light queue is sufficient to ensure an exponentially decaying tail 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 deemphasizes the heavy queue, it does not *ignore* it, unlike the policy that gives priority to L. That is, if the H queue occupancy gets overwhelmingly large compared to L, the LMW policy will serve the H queue. In contrast, the priority for L policy will ignore any buildup in H as long as L is nonempty. This property turns out to be crucial in more complex queueing models, where throughput optimality is nontrivial to obtain. For example, when the queues have time-varying connectivity to the server, the LMW policy will stabilize both queues for all traffic rates within the rate region, whereas priority for L leads to a strictly smaller stability region [16].

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

Theorem 10: Under log-max-weight scheduling, q_L is light-tailed. Specifically

$$\liminf_{b\to\infty} -\frac{1}{b}\log \mathbb{P}\{q_L \geq b\} \geq \min(E_L,\ C_H - 1) \tag{33}$$

where E_L is the intrinsic exponent, given by (13) and (15).

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

$$\begin{aligned}\mathbb{P}\{q_L \geq b\} = &\underbrace{\mathbb{P}\{q_L \geq b,\ \log(1+q_H) < \delta b\}}_{\text{(i)}}\\ &+ \underbrace{\mathbb{P}\{q_L \geq b,\ (1-\delta)b \geq \log(1+q_H) \geq \delta b\}}_{\text{(ii)}}\\ &+ \underbrace{\mathbb{P}\{q_L \geq b,\ \log(1+q_H) > (1-\delta)b\}}_{\text{(iii)}}. \end{aligned} \tag{34}$$

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

(i) $\mathbb{P}\{q_L \geq b,\ \log(1+q_H) < \delta b\}$: Intuitively, this event corresponds to an overflow of the light queue, when the heavy 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 $-\tau \leq 0$ 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 $-u \leq 0$ satisfying

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

Therefore

$$\begin{aligned}&\mathbb{P}\{q_L \geq b,\ \log(1+q_H) < \delta b\}\\ &\quad\leq \mathbb{P}\left\{\exists u \geq 0 \,\middle|\, \sum_{i=-u+1}^{0}(L(i)-1) > (1-\delta)b\right\}.\end{aligned}$$

Letting $S_u = \sum_{i=-u+1}^{0} L(i)$, the above inequality can be written as

$$\begin{aligned}&\mathbb{P}\{q_L \geq b,\ \log(1+q_H) < \delta b\}\\ &\quad\leq \mathbb{P}\left\{\sup_{u\geq 0}(S_u - u) > (1-\delta)b\right\}.\end{aligned} \tag{35}$$

The right-hand side of (35) 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 [13, Lemma 1.5], we get

$$\begin{aligned}\liminf_{b\to\infty} -\frac{1}{b}\log \mathbb{P}\left\{\sup_{u\geq 0}(S_u - u) > (1-\delta)b\right\} &\geq \inf_{u>0} u\Lambda_L^*\left(1+\frac{1-\delta}{u}\right)\\ &= \inf_{a>0}\frac{1-\delta}{a}\Lambda_L^*(1+a) = (1-\delta)E_L. \end{aligned}\tag{36}$$

From (35) and (36), we see that for every $\epsilon > 0$, there exists a positive κ_3 (in our notation, we suppress the dependence of κ_3 on ϵ), such that for all large enough b

$$\mathbb{P}\left\{q_L \geq b,\ \log(1+q_H) < \delta b\right\} < \kappa_3 e^{-b(1-\delta)(E_L-\epsilon)}. \tag{37}$$

(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

$$\begin{aligned}\mathbb{P}\left\{q_L \geq b, \log(1+q_H) > (1-\delta)b\right\} &= \mathbb{P}\left\{q_L \geq b, q_H > e^{(1-\delta)b} - 1\right\}\\ &\leq \mathbb{P}\left\{q_L + q_H > e^{(1-\delta)b}\right\}.\end{aligned}$$

We have shown earlier, in the proof of Theorem 4, that for any nonidling policy and any $\epsilon > 0$

$$\mathbb{P}\left\{q_L + q_H > M\right\} < \kappa_1 M^{-(C_H - 1 - \epsilon)}$$

for and some $\kappa_1 > 0$. Therefore

$$\begin{aligned}&\mathbb{P}\left\{q_L \geq b,\ \log(1+q_H) > (1-\delta)b\right\}\\ &\quad < \kappa_1 \exp\left(-(1-\delta)b(C_H - 1 - \epsilon)\right) \qquad \forall\, \epsilon > 0.\end{aligned}\tag{38}$$

(iii) We now deal with the second term in (34). Let us call this event $\mathcal{E}_2$. Suppose this event occurs at time 0. Denote by $-\tau \leq 0$ the last time during the current busy period that H received service, and define

$$\eta = \log(1 + q_H(-\tau)). \tag{39}$$

If 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 $\eta = 0$. We can deduce that $\eta \leq (1-\delta)b$ because H receives no service in $[-\tau+1, 0]$. Since H received service at time $-\tau$, it follows from (39) that $q_L(-\tau) \leq \eta$. Therefore, during the time interval $[-\tau+1, 0]$, the queue length of L grows from at most η to more than b, in spite of receiving all the service. Additionally, it is evident from (39) that $q_H(-\tau) + q_L(-\tau) \geq e^\eta - 1$. The last two statements imply that every time the event $\mathcal{E}_2$ occurs, there necessarily exists $-u \leq 0$ and $\xi \geq 0$ such that $S_u - u > b - \xi$ and $q_H(-u) + q_L(-u) \geq e^\xi - 1$. This leads to the bound in (40), which can be further upper-bounded using a union bound, as shown in (40) and (41) at the bottom of the page. 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, the right-hand side of (41) equals

$$\begin{aligned}&\sum_{\xi,\, u}\mathbb{P}\left\{S_u - u > b - \xi\right\}\mathbb{P}\left\{q_H(-u) + q_L(-u) \geq e^\xi - 1\right\}\\ &\leq \sum_{\xi=0}^{(1-\delta)b}\sum_{u\geq 0}\mathbb{P}\left\{S_u - u > b - \xi\right\}\kappa_1 e^{-(C_H-1-\epsilon)\xi} \qquad (42)\\ &\geq \sum_{\xi=0}^{(1-\delta)b}\kappa_3 e^{-(E_L-\epsilon)(b-\xi)}\kappa_1 e^{-(C_H-1-\epsilon)\xi}. \qquad (43)\end{aligned}$$

Equation (42) follows from Theorem 4, and (43) is a classical large deviation bound that follows, for example, from [13, Lemma 1.5]. (These inequalities are valid for all $\epsilon > 0$ and for some κ_1 and κ_3 that depend on ϵ). Thus, for every $\epsilon > 0$

$$\mathbb{P}\left\{\mathcal{E}_2\right\} \leq \sum_{\xi=0}^{(1-\delta)b}\kappa_1\kappa_3 e^{-[(C_H-1-\epsilon)\xi + (E_L-\epsilon)(b-\xi)]}. \tag{44}$$

Let us now distinguish two cases.

a) $C_H - 1 > E_L$: In this case, we can bound the probability in (44) as follows:

$$\begin{aligned}\mathbb{P}\left\{\mathcal{E}_2\right\} &\leq \kappa_1\kappa_3 e^{-b(E_L-\epsilon)}\sum_{\xi=0}^{(1-\delta)b} e^{-(C_H-1-E_L)\xi}\\ &< \frac{\kappa_1\kappa_3}{1 - e^{-(C_H-1-E_L)}}e^{-b(E_L-\epsilon)}\\ &= \kappa e^{-b(E_L-\epsilon)} \qquad \forall \epsilon > 0\end{aligned}\tag{45}$$

where $\kappa > 0$ is some constant that depends on ϵ.

b) $C_H - 1 \leq E_L$: Manipulating similar to case a), we get

$$\mathbb{P}\left\{\mathcal{E}_2\right\} < \kappa e^{-b(C_H-1-\epsilon)(1-\delta)} \qquad \forall \epsilon > 0. \tag{46}$$

$$\begin{aligned}\mathbb{P}\left\{\mathcal{E}_2\right\} &\leq \mathbb{P}\left\{\exists\, \xi \leq (1-\delta)b,\ \exists\, u \geq 0 \text{ such that } S_u - u > b - \xi,\ q_H(-u) + q_L(-u) \geq e^\xi - 1\right\} \qquad (40)\\ &\leq \sum_{\xi=0}^{(1-\delta)b}\sum_{u\geq 0}\mathbb{P}\left\{S_u - u > b - \xi,\ q_H(-u) + q_L(-u) \geq e^\xi - 1\right\}. \qquad (41)\end{aligned}$$

Let us now put together the bounds on terms (i)–(iii) into (34).

1) If $C_H - 1 > E_L$, we get from (37), (38), and (45)

$$\mathbb{P}\{q_L \geq b\} < e^{-b(1-\delta)(E_L-\epsilon)} \cdot \left[\kappa_3 + \kappa_1 e^{-((1-\delta)b(C_H-1-E_L))} + \kappa\right]$$

from which it is immediate that

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

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

$$\liminf_{b\to\infty} -\frac{1}{b}\log \mathbb{P}\{q_L \geq b\} \geq E_L. \tag{47}$$

2) If $C_H - 1 \leq E_L$, we get from (37), (38), and (46)

$$\mathbb{P}\{q_L \geq b\} < e^{-b(1-\delta)(C_H-1-\epsilon)} \cdot \left[\kappa_3 e^{-((1-\delta)b(E_L-C_H+1))} + \kappa_1 + \kappa\right]$$

from which it is immediate that

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

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

$$\liminf_{b\to\infty} -\frac{1}{b}\log \mathbb{P}\{q_L \geq b\} \geq C_H - 1. \tag{48}$$

Theorem 10 now follows from (47) and (48). □

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 10 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 large deviation lower bound that matches the upper bound in Theorem 10.

Theorem 11: 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

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

Proof: In light of Theorem 10, it is enough to prove that

$$\limsup_{b\to\infty} -\frac{1}{b}\log \mathbb{P}\{q_L \geq b\} \leq \min(E_L,\ C_H - 1).$$

Let us denote by $q_L^{(p)}$ the length of the light queue, when it is given complete priority over H. Note that $\mathbb{P}\left\{q_L^{(p)} > b\right\}$ is a lower bound on the overflow probability under *any* policy, including LMW. Therefore, for all $b > 0$, $\mathbb{P}\{q_L \geq b\} \geq \mathbb{P}\left\{q_L^{(p)} > b\right\}$. This implies

$$\limsup_{b\to\infty} -\frac{1}{b}\log \mathbb{P}\{q_L \geq b\} \leq \limsup_{b\to\infty} -\frac{1}{b}\log \mathbb{P}\left\{q_L^{(p)} > b\right\} = E_L \tag{49}$$

where the last step is from (14).

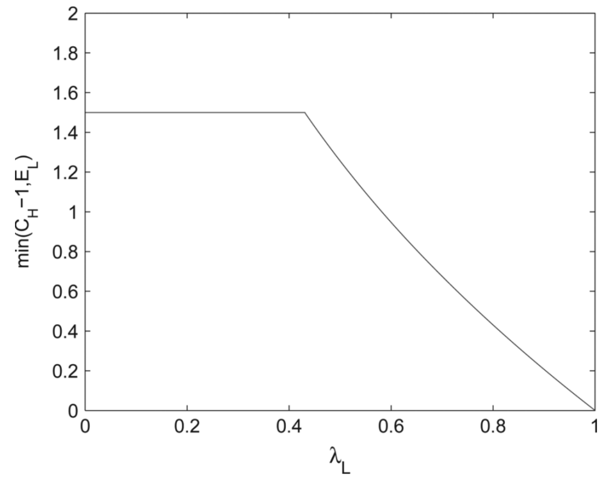

Fig. 3. Large deviation exponent for q_L under LMW scheduling, as a function of λ_L. The light queue is fed by Poisson bursts, and $C_H = 2.5$.

Next, we can show, following the arguments in the proofs of Proposition 10 and Theorem 8 that

$$\mathbb{P}\{q_L \geq b\} \geq \lambda_H \mathbb{P}\left\{H_{\mathrm{R}} \geq e^b - 1,\ \sum_{i=1}^{H_{\mathrm{A}}} L(i) \geq b\right\}.$$

Arguing as in the proof of Lemma 5, we can show that

$$\mathbb{P}\left\{H_{\mathrm{R}} \geq e^b - 1,\ \sum_{i=1}^{H_{\mathrm{A}}} L(i) \geq b\right\} \sim \mathbb{P}\left\{H_{\mathrm{R}} \geq e^b - 1\right\}.$$

Thus

$$\mathbb{P}\{q_L \geq b\} \gtrsim \mathbb{P}\left\{H_{\mathrm{R}} \geq e^b - 1\right\}.$$

Next, since $H(\cdot)$ is regularly varying with tail coefficient C_H, H_{R} is also regularly varying with tail coefficient $C_H - 1$, so that $\mathbb{P}\left\{H_{\mathrm{R}} \geq e^b - 1\right\} = U(e^b)e^{-b(C_H-1)}$. Finally, we can write

$$\begin{aligned}\limsup_{b\to\infty} -\frac{1}{b}\log \mathbb{P}\{q_L \geq b\} &\geq \limsup_{b\to\infty} -\frac{1}{b}\log \mathbb{P}\left\{H_{\mathrm{R}} \geq e^b - 1\right\} \\ &= C_H - 1 - \limsup_{b\to\infty} \frac{\log U(e^b)}{b}.\end{aligned}$$

The final limit supremum is shown to be zero in Lemma 6 in the Appendix, using a representation theorem for slowly varying functions. Thus

$$\limsup_{b\to\infty} -\frac{1}{b}\log \mathbb{P}\{q_L \geq b\} \leq C_H - 1. \tag{50}$$

Equations (49) and (50) imply the theorem. □

Fig. 3 shows the large deviation exponent given by Theorem 11 as a function of λ_L, for $C_H = 2.5$ and Poisson inputs 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 were absent. This mode of overflow is more likely for larger values of λ_L, which explains the diminishing exponent in this regime.

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 takes over the server, 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 determined entirely by the "burstiness" of the heavy-tailed traffic, as reflected in the tail coefficient.

VIII. 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 tail coefficient of the light queue distribution among all nonidling policies.

Another main 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 focused on a very simple queueing network in this paper, we believe that the insights obtained 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 [19], in the context of expected delays. Regarding the asymptotic distribution of the steady-state delays, we can intuitively expect similar behavior to that of the respective queue lengths, as long as the queues are served in an FCFS fashion. However, when the queueing discipline is not FCFS, the delay asymptotics could be more complex. Furthermore, results analogous to the ones derived in this paper are also expected to hold in continuous time.

We also believe that the LMW policy occupies a "sweet spot" in the context of scheduling light-tailed traffic in the presence of heavy-tailed traffic. This is because the LMW policy deemphasizes 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 multihop networks and time-varying channel models.

Appendix
Technical Lemmata

Lemma 2: The tail coefficient of $H(\cdot)+L(\cdot)$ is C_H.

Proof: Clearly, $\mathbb{E}\left[(H+L)^{C_H+\delta}\right] \geq \mathbb{E}\left[H^{C_H+\delta}\right] = \infty$, for every $\delta > 0$. We next need to show that $\mathbb{E}\left[(H+L)^{C_H-\delta}\right] < \infty$, for every $\delta > 0$. For a random variable X and event E, let us introduce the notation $\mathbb{E}[X;E] = \mathbb{E}[X1_E]$, where 1_E is the indicator of E. (Thus, for example, $\mathbb{E}[X] = \mathbb{E}[X;E] + \mathbb{E}[X;E^c]$). Now

$$\begin{aligned}\mathbb{E}\left[(H+L)^{C_H-\delta}\right] &= \mathbb{E}\left[(H+L)^{C_H-\delta}; H>L\right] \\ &\quad + \mathbb{E}\left[(H+L)^{C_H-\delta}; H\leq L\right] \\ &\leq \mathbb{E}\left[(2H)^{C_H-\delta}; H>L\right] \\ &\quad + \mathbb{E}\left[(2L)^{C_H-\delta}; H\leq L\right] \\ &< 2^{C_H-\delta}\left\{\mathbb{E}\left[H^{C_H-\delta}\right]+\mathbb{E}\left[L^{C_H-\delta}\right]\right\}<\infty\end{aligned}$$

where the last step follows from the tail coefficient of $H(\cdot)$ and the light-tailed nature of $L(\cdot)$. □

Lemma 3: $\mathbb{P}\{H_{\mathrm{R}}\geq m,\ H_{\mathrm{A}}\geq n\} = \mathbb{P}\{H_{\mathrm{R}}\geq m+n\}$.

Proof: Using (1) and (2)

$$\begin{aligned}\mathbb{P}\{H_{\mathrm{R}}\geq m,\ H_{\mathrm{A}}\geq n\} &= \sum_{k\geq m}\sum_{l\geq n}\frac{\mathbb{P}\{H_+=k+l\}}{\mathbb{E}[H_+]} \\ &= \sum_{k\geq m}\sum_{p=k+n}^{\infty}\frac{\mathbb{P}\{H_+=p\}}{\mathbb{E}[H_+]} \\ &= \sum_{k\geq m}\mathbb{P}\{H_{\mathrm{R}}=k+n\} \\ &= \mathbb{P}\{H_{\mathrm{R}}\geq m+n\}.\end{aligned}$$

□

Lemma 4: Let $N \in \mathcal{IR}$ be a nonnegative integer-valued random variable. Let $X_i i \geq 1$ be i.i.d. nonnegative light-tailed random variables, with mean μ, independent of N. Define

$$S_N = \sum_{i=1}^{N} X_i.$$

Then

$$\mathbb{P}\{S_N > b\} \sim \mathbb{P}\{N > b/\mu\}.$$

Proof: For notational ease, we will prove the result for $\mu = 1$, although the result and its proof technique remain valid for any $\mu > 0$. First, for a fixed $\delta > 0$, we have

$$\begin{aligned}\mathbb{P}\{S_N>b\} &= \mathbb{P}\{S_N>b,\, N\leq b(1-\delta)\} \\ &\quad + \mathbb{P}\{S_N>b,\, N>b(1-\delta)\} \\ &< \mathbb{P}\left\{S_{\lfloor b(1-\delta)\rfloor}>b\right\} + \mathbb{P}\{N>b(1-\delta)\}. \quad (51)\end{aligned}$$

Next, we write a lower bound

$$\begin{aligned}\mathbb{P}\{S_N > b\} &\geq \mathbb{P}\{S_N > b, N > b(1+\delta)\}\\ &= \mathbb{P}\{N > b(1+\delta)\} - \mathbb{P}\{S_N \leq b, N > b(1+\delta)\}\\ &\geq \mathbb{P}\{N > b(1+\delta)\} - \mathbb{P}\left\{S_{\lceil b(1+\delta)\rceil} \leq b\right\}. \end{aligned}\tag{52}$$

Since the X_i have a well-defined moment-generating function, their sample average satisfies an exponential concentration inequality around the mean. Specifically, we can show using the Chernoff bound that there exist positive constants κ, η such that

$$\mathbb{P}\left\{S_{\lfloor b(1-\delta)\rfloor} > b\right\} < \kappa e^{-b\eta}.$$

Thus

$$\mathbb{P}\left\{S_{\lfloor b(1-\delta)\rfloor} > b\right\} = o(\mathbb{P}\{N > b\})\tag{53}$$

as $b \to \infty$. Similarly

$$\mathbb{P}\left\{S_{\lfloor b(1+\delta)\rfloor} \leq b\right\} = o(\mathbb{P}\{N > b\}).\tag{54}$$

Next, getting back to (51)

$$\begin{aligned}\limsup_{b\to\infty} \frac{\mathbb{P}\{S_N > b\}}{\mathbb{P}\{N > b\}} \leq \limsup_{b\to\infty} &\frac{\mathbb{P}\left\{S_{\lfloor b(1-\delta)\rfloor} > b\right\}}{\mathbb{P}\{N > b\}}\\ &+ \limsup_{b\to\infty} \frac{\mathbb{P}\{N > b(1-\delta)\}}{\mathbb{P}\{N > b\}}.\end{aligned}\tag{55}$$

The first term on the right-hand side is zero in view of (53), so that for all δ, we have

$$\limsup_{b\to\infty} \frac{\mathbb{P}\{S_N > b\}}{\mathbb{P}\{N > b\}} \leq \limsup_{b\to\infty} \frac{\mathbb{P}\{N > b(1-\delta)\}}{\mathbb{P}\{N > b\}}.$$

Taking the limit as $\delta \downarrow 0$

$$\limsup_{b\to\infty} \frac{\mathbb{P}\{S_N > b\}}{\mathbb{P}\{N > b\}} \leq \lim_{\delta\downarrow 0} \limsup_{b\to\infty} \frac{\mathbb{P}\{N > b(1-\delta)\}}{\mathbb{P}\{N > b\}} = 1.\tag{56}$$

The final limit is unity, by the definition of the class $\mathcal{IR}$. Similarly, we can show using (52), (54), and the intermediate-regular variation of the tail of N that

$$\liminf_{b\to\infty} \frac{\mathbb{P}\{S_N > b\}}{\mathbb{P}\{N > b\}} \geq 1.\tag{57}$$

Equations (56) and (57) imply the result. □

The above lemma can be proved under more general assumptions than stated here; see [24].

Lemma 5: If $H(\cdot) \in \mathcal{OR}$, we have

$$\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ \sum_{i=1}^{H_{\mathrm{A}}} L(i) \geq b\right\} \sim \begin{cases} \mathbb{P}\left\{H_{\mathrm{R}} \geq \frac{b}{\lambda_L}\right\}, & \text{if } \frac{\alpha_L}{\alpha_H} < 1\\ \mathbb{P}\left\{H_{\mathrm{R}} \geq b + \frac{b}{\lambda_L}\right\}, & \text{if } \frac{\alpha_L}{\alpha_H} = 1\\ \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}, & \text{if } \frac{\alpha_L}{\alpha_H} > 1.\end{cases}\tag{58}$$

Proof: In this proof, let us take $\lambda_L = 1$ for notational simplicity, although the same proof works without this assumption. Denote $S_n = \sum_{i=1}^{n} L(i)$.

We first get an upper bound. For every $\delta > 0$, we have

$$\begin{aligned}&\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b\right\}\\ &\quad= \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b,\ H_{\mathrm{A}} < b(1-\delta)\right\}\\ &\qquad+ \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b,\ H_{\mathrm{A}} > b(1-\delta)\right\}\\ &\quad< \mathbb{P}\left\{S_{H_{\mathrm{A}}} \geq b,\ H_{\mathrm{A}} < b(1-\delta)\right\}\\ &\qquad+ \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ H_{\mathrm{A}} > b(1-\delta)\right\}\\ &\quad\leq \mathbb{P}\left\{S_{H_{\mathrm{A}}} \geq b,\ H_{\mathrm{A}} < b(1-\delta)\right\}\\ &\qquad+ \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H} + b(1-\delta)\right\}\\ &\quad\leq \mathbb{P}\left\{S_{\lfloor b(1-\delta)\rfloor} \geq b\right\} + \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H} + b(1-\delta)\right\}.\end{aligned}\tag{59}$$

In (59), we have utilized Lemma 3.

Next, let us derive a lower bound.

$$\begin{aligned}&\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b\right\}\\ &\quad\geq \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b,\ H_{\mathrm{A}} > b(1+\delta)\right\}\\ &\quad= \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ H_{\mathrm{A}} > b(1+\delta)\right\}\\ &\qquad- \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} < b,\ H_{\mathrm{A}} > b(1+\delta)\right\}\\ &\quad\geq \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ H_{\mathrm{A}} > b(1+\delta)\right\}\\ &\qquad- \mathbb{P}\left\{S_{H_{\mathrm{A}}} < b,\ H_{\mathrm{A}} > b(1+\delta)\right\}\\ &\quad\geq \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H} + b(1+\delta)\right\} - \mathbb{P}\left\{S_{\lceil b(1+\delta)\rceil} < b\right\}.\end{aligned}\tag{60}$$

Equation (60) uses Lemma 3. Now, observe that the terms $\mathbb{P}\left\{S_{\lfloor b(1-\delta)\rfloor} \geq b\right\}$ in (59) and $\mathbb{P}\left\{S_{\lceil b(1+\delta)\rceil} < b\right\}$ in (60) decay exponentially fast as $b \to \infty$, for any $\delta > 0$. This is because $L(\cdot)$ is light-tailed, and its sample average satisfies an exponential concentration inequality around the mean (unity). More precisely, a Chernoff bound can be used to show that

$$\mathbb{P}\left\{S_{\lfloor b(1-\delta)\rfloor} \geq b\right\} = o\left(\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H} + b\right\}\right)\tag{61}$$

and

$$\mathbb{P}\left\{S_{\lceil b(1+\delta)\rceil} < b\right\} = o\left(\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H} + b\right\}\right).\tag{62}$$

Case (i): $\alpha_L/\alpha_H < 1$. Using (59), we write

$$\begin{aligned}&\limsup_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}, S_{H_{\mathrm{A}}} \geq b\right\}}{\mathbb{P}\{H_{\mathrm{R}} \geq b\}}\\ &\quad\leq \mathrm{ls}\, \frac{\mathbb{P}\left\{S_{\lfloor b(1-\delta)\rfloor} \geq b\right\}}{\mathbb{P}\{H_{\mathrm{R}} \geq b\}} + \mathrm{ls}\, \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H} + b(1-\delta)\right\}}{\mathbb{P}\{H_{\mathrm{R}} \geq b\}}\end{aligned}$$

where we have used "ls" to abbreviate $\limsup_{b\to\infty}$. The first term on the right is zero in view of (61). Since $\alpha_L/\alpha_H < 1$

$$\begin{aligned}&\limsup_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b\right\}}{\mathbb{P}\{H_{\mathrm{R}} \geq b\}}\\ &\quad\leq \limsup_{b\to\infty} \frac{\mathbb{P}\{H_{\mathrm{R}} \geq b(1-\delta)\}}{\mathbb{P}\{H_{\mathrm{R}} \geq b\}} \qquad \forall \delta > 0.\end{aligned}$$

Thus

$$\limsup_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b\right\}} \leq \lim_{\delta\downarrow 0}\limsup_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b(1-\delta)\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b\right\}} = 1. \tag{63}$$

To justify the final step, recall that according to Lemma 1, $H(\cdot) \in \mathcal{OR}$ implies $H_{\mathrm{R}} \in \mathcal{ER}$. Since $\mathcal{ER} \subset \mathcal{IR}$, the final limit in (63) is unity by the definition of intermediate-regular variation (Definition 2).

Along similar lines, we can use (60), (62), and the fact that $H_{\mathrm{R}} \in \mathcal{IR}$ to show that

$$\liminf_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b\right\}} \geq \lim_{\delta\downarrow 0}\liminf_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b(1+2\delta)\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b\right\}} = 1. \tag{64}$$

Equations (63) and (64) imply that

$$\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b\right\} \sim \mathbb{P}\left\{H_{\mathrm{R}} \geq b\right\}$$

which implies Lemma 5 for $\alpha_L/\alpha_H < 1$, and $\lambda_L = 1$.

Case (ii): $\alpha_L/\alpha_H = 1$. The proof is similar to the previous case. Here, we get

$$\mathbb{P}\left\{H_{\mathrm{R}} \geq b,\ S_{H_{\mathrm{A}}} \geq b\right\} \sim \mathbb{P}\left\{H_{\mathrm{R}} \geq 2b\right\}.$$

Case (iii): $\alpha_L/\alpha_H > 1$.

For the upper bound, we have from (59) and (61)

$$\limsup_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}} \leq \limsup_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H} + b(1-\delta)\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}} \leq 1.$$

Similarly, for the lower bound, we have from (60) and (62)

$$\begin{aligned}\liminf_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}} &\geq \liminf_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H} + b(1+\delta)\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}} \\ &\geq \liminf_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}(1+\delta)\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}} \qquad \forall \delta > 0.\end{aligned}$$

Thus

$$\liminf_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}} \geq \lim_{\delta\downarrow 0}\liminf_{b\to\infty} \frac{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}(1+\delta)\right\}}{\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}} = 1$$

where the last limit is unity due to the intermediate-regular variation of H_{R}. Therefore, for $\alpha_L/\alpha_H > 1$, we can conclude that

$$\mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H},\ S_{H_{\mathrm{A}}} \geq b\right\} \sim \mathbb{P}\left\{H_{\mathrm{R}} \geq b^{\alpha_L/\alpha_H}\right\}.$$

Lemma 5 is now proved. □

Lemma 6: For any slowly varying function $U(\cdot)$

$$\lim_{a\to\infty} \frac{\log U(a)}{\log a} = 0.$$

Proof: We use the representation theorem for slowly varying functions derived in [11]. For every slowly varying function $U(\cdot)$, there exists a $B > 0$ such that for all $x \geq B$, the function can be written as

$$U(x) = \exp\left(v(x) + \int_B^x \frac{\zeta(y)}{y}\mathrm{d}y\right)$$

where $v(x)$ converges to a finite constant, and $\zeta(x) \to 0$ as $x \to \infty$. Therefore

$$\lim_{a\to\infty} \frac{\log U(a)}{\log a} = \lim_{a\to\infty} \frac{v(a) + \int_B^a \frac{\zeta(y)}{y}\mathrm{d}y}{\log a} = \lim_{a\to\infty} \frac{\int_B^a \frac{\zeta(y)}{y}\mathrm{d}y}{\log a}$$

where the last step is because $v(a)$ converges to a constant. Next, given any $\epsilon > 0$, choose $C(\epsilon)$ such that $|\zeta(a)| < \epsilon,\ \forall\, a > C(\epsilon)$. Then, we have

$$\begin{aligned}\lim_{a\to\infty} \frac{\left|\int_B^a \frac{\zeta(y)}{y}\mathrm{d}y\right|}{\log a} &\leq \lim_{a\to\infty} \frac{\int_B^{C(\epsilon)} \frac{|\zeta(y)|}{y}\mathrm{d}y + \int_{C(\epsilon)}^a \frac{|\zeta(y)|}{y}\mathrm{d}y}{\log a} \\ &\leq \lim_{a\to\infty} \frac{\epsilon \log \frac{a}{C(\epsilon)}}{\log a} = \epsilon.\end{aligned}$$

Since the above is true for every $\epsilon > 0$, the result follows. □

REFERENCES

[1] V. Anantharam, "Scheduling strategies and long-range dependence," *Queue. Syst.*, vol. 33, no. 1, pp. 73–89, 1999.

[2] N. Bingham, C. Goldie, and J. Teugels, *Regular Variation*. Cambridge, U.K.: Cambridge Univ. Press, 1987.

[3] S. Borst, O. Boxma, and M. Van Uitert, "The asymptotic workload behavior of two coupled queues," *Queue. Syst.*, vol. 43, no. 1, pp. 81–102, 2003.

[4] S. Borst, M. Mandjes, and M. Van Uitert, "Generalized processor sharing with light-tailed and heavy-tailed input," *IEEE/ACM Trans. Netw.*, vol. 11, no. 5, pp. 834–834, Oct. 2003.

[5] S. Borst, O. Boxma, R. N. Núñez-Queija, and A. Zwart, "The impact of the service discipline on delay asymptotics," *Perform. Eval.*, vol. 54, no. 2, pp. 175–206, 2003.

[6] O. Boxma and B. Zwart, "Tails in scheduling," *Perform. Eval. Rev.*, vol. 34, no. 4, p. 20, 2007.

[7] O. Boxma, Q. Deng, and A. Zwart, "Waiting-time asymptotics for the M/G/2 queue with heterogeneous servers," *Queue. Syst.*, vol. 40, no. 1, pp. 5–31, 2002.

[8] A. Brzezinski and E. Modiano, "Dynamic reconfiguration and routing algorithms for IP-over-WDM networks with stochastic traffic," *J. Lightw. Technol.*, vol. 23, no. 10, pp. 3188–3188, Oct. 2005.

[9] D. Cline, "Intermediate regular and Π variation," *Proc. London Math. Soc.*, vol. 3, no. 3, pp. 594–594, 1994.

[10] A. Eryilmaz, R. Srikant, and J. Perkins, "Stable scheduling policies for fading wireless channels," *IEEE/ACM Trans. Netw.*, vol. 13, no. 2, pp. 411–424, Apr. 2005.

[11] J. Galambos and E. Seneta, "Regularly varying sequences," *Proc. Amer. Math. Soc.*, vol. 41, no. 1, pp. 110–116, 1973.
[12] R. Gallager, *Discrete Stochastic Processes*. Norwell, MA: Kluwer, 1996.
[13] A. Ganesh, N. O'Connell, and D. Wischik, *Big Queues*. New York: Springer-Verlag, 2004.
[14] M. Harchol-Balter, A. Scheller-Wolf, and A. Young, "Surprising results on task assignment in server farms with high-variability workloads," in *Proc. ACM SIGMETRICS*, 2009, pp. 287–298.
[15] D. Heath, S. Resnick, and G. Samorodnitsky, "Heavy tails and long range dependence in on/off processes and associated fluid models," *Math. Oper. Res.*, vol. 23, no. 1, pp. 145–165, 1998.
[16] K. Jagannathan, M. Markakis, E. Modiano, and J. Tsitsiklis, "Throughput optimal scheduling in the presence of heavy-tailed traffic," in *Proc. 48th Annu. Allerton Conf. Commun., Control, Comput.*, Monticello, IL, 2010, pp. 953–960.
[17] K. Jagannathan, M. Markakis, E. Modiano, and J. Tsitsiklis, "Queue length asymptotics for generalized max-weight scheduling in the presence of heavy-tailed traffic," in *Proc. IEEE INFOCOM*, Shanghai, China, 2011, pp. 2318–2326.
[18] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of ethernet traffic (Extended version)," *IEEE/ACM Trans. Netw.*, vol. 2, no. 1, pp. 1–15, Feb. 1994.
[19] M. Markakis, E. Modiano, and J. Tsitsiklis, "Scheduling policies for single-hop networks with heavy-tailed traffic," in *Proc. Allerton Conf.*, Monticello, IL, 2009, pp. 112–120.
[20] N. McKeown, A. Mekkittikul, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input-queued switch," *IEEE Trans. Commun.*, vol. 47, no. 8, pp. 1260–1267, Aug. 1999.
[21] M. Neely, E. Modiano, and C. Rohrs, "Power and server allocation in a multi-beam satellite with time varying channels," in *Proc. IEEE INFOCOM*, 2002, vol. 3, pp. 1451–1460.
[22] M. Neely, E. Modiano, and C. Rohrs, "Dynamic power allocation and routing for time varying wireless networks," in *Proc. IEEE INFOCOM*, 2003, vol. 1, pp. 745–755.
[23] K. Park and W. Willinger, *Self-Similar Network Traffic and Performance Evaluation*. New York: Wiley, 2000.
[24] C. Robert and J. Segers, "Tails of random sums of a heavy-tailed number of light-tailed terms," *Insurance, Math. Econ.*, vol. 43, no. 1, pp. 85–92, 2008.
[25] K. Sigman, "A primer on heavy-tailed distributions," *Queue. Syst.*, vol. 33, no. 1–3, pp. 261–275, 1999.
[26] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.
[27] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Trans. Inf. Theory*, vol. 39, no. 2, pp. 466–478, Mar. 1993.
[28] A. Wierman and B. Zwart, "Is tail-optimal scheduling possible?," *Oper. Res.* 2009 [Online]. Available: http://www.cs.caltech.edu/~adamw/papers/impossibility.pdf, submitted for publication

Krishna Jagannathan received the B. Tech. degree in electrical engineering from the Indian Institute of Technology (IIT) Madras, Chennai, India, in 2004, and the S.M. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 2006 and 2010, respectively.

He is an Assistant Professor with the Department of Electrical Engineering, IIT Madras. From 2010 to 2011, he was a Visiting Post-Doctoral Scholar in computing and mathematical sciences with the California Institute of Technology (Caltech), Pasadena, and an off-campus Post-Doctoral Fellow with MIT. He worked as a Consultant with the Mathematical Sciences Research Center, Bell Laboratories, Murray Hill, NJ, in 2005; an Engineering Intern with Qualcomm, Campbell, CA, in 2007; and a Summer Research Fellow with the Indian Institute of Science (IISc), Bangalore, India, in 2003. His research interests lie in communication networks, queuing theory, and network control.

Mihalis Markakis received the B.S. degree from the National Technical University of Athens, Athens, Greece, in 2005, and the M.S. degree from the University of Southern California, Los Angeles, in 2008, both in electrical engineering, and is currently pursuing the Ph.D. degree at the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge.

He held an internship position with Alcatel-Lucent Bell Labs, Murray Hill, NJ, in 2010. His research interests are in the areas of modeling, analysis, and control of dynamic and stochastic systems, with emphasis in queueing theory and networks.

Eytan Modiano (S'90–M'93–SM'00–F'12) received the B.S. degree in electrical engineering and computer science from the University of Connecticut, Storrs, in 1986, and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, in 1989 and 1992, respectively.

He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Post Doctoral Fellow from 1992 to 1993. Between 1993 and 1999, he was with the Massachusetts Institute of Technology (MIT) Lincoln Laboratory, Lexington, MA, where he was a project leader for its Next Generation Internet (NGI) project. Since 1999, he has been on the faculty at MIT; where he is a Professor with the Department of Aeronautics and Astronautics and the Laboratory for Information and Decision Systems (LIDS). His research is on communication networks and protocols with emphasis on satellite, wireless, and optical networks.

Prof. Modiano is currently an Associate Editor for the IEEE/ACM TRANSACTIONS ON NETWORKING. He had served as Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY and as Guest Editor for the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS special issue on WDM network architectures; the *Computer Networks Journal* special issue on broadband Internet access; the *Journal of Communications and Networks* special issue on wireless ad hoc networks; and the IEEE/OSA JOURNAL OF LIGHTWAVE TECHNOLOGY special issue on optical networks. He was the Technical Program Co-Chair for IEEE WiOpt 2006, IEEE INFOCOM 2007, and ACM MobiHoc 2007. He is the co-recipient of the SIGMETRICS 2006 Best Paper Award for the paper "Maximizing Throughput in Wireless Networks via Gossiping" and the WiOpt 2005 Best Student Paper Award for the paper "Minimum Energy Transmission Scheduling Subject to Deadline Constraints."

John N. Tsitsiklis (M'83–SM'97–F'99) received the B.S. degree in mathematics, and the B.S., M.S., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, in 1980, 1980, 1981, and 1984, respectively.

He is currently a Clarence J. Lebel Professor with the Department of Electrical Engineering, MIT. He has served as a Co-Director of the MIT Operations Research Center from 2002 to 2005, and in the National Council on Research and Technology in Greece (2005–2007, 2011–). His research interests are in systems, optimization, communications, control, and operations research. He has coauthored four books and several journal papers in these areas.

Dr. Tsitsiklis is a member of the National Academy of Engineering. In 2008, he was conferred the title of Doctor honoris causa from the Université Catholique de Louvain, Leuven, Belgium. He received the Outstanding Paper Award from the IEEE Control Systems Society in 1986, the MIT Edgerton Faculty Achievement Award in 1989, the Bodossakis Foundation Prize in 1995, and the INFORMS/CSTS Prize in 1997.