MIT Open Access Articles

Delay analysis of the Max-Weight policy under heavy-tailed traffic via fluid approximations

Massachusetts Institute of Technology

Delay Analysis of the Max-Weight Policy under Heavy-Tailed Traffic via Fluid Approximations

Mihalis G. Markakis, Eytan Modiano, and John N. Tsitsiklis

Abstract—We consider a single-hop switched queueing network with a mix of heavy-tailed (i.e., arrival processes with infinite variance) and light-tailed traffic, and study the delay performance of the Max-Weight policy, known for its throughput optimality and asymptotic delay optimality properties. Classical results in queueing theory imply that heavy-tailed queues are delay unstable, i.e., they experience infinite expected delays in steady state. Thus, we focus on the impact of heavy-tailed traffic on the light-tailed queues, using delay stability as performance metric. Recent work has shown that this impact may come in the form of subtle rate-dependent phenomena, the stochastic analysis of which is quite cumbersome.

Our goal is to show how fluid approximations can facilitate the delay analysis of the Max-Weight policy under heavy-tailed traffic. More specifically, we show how fluid approximations can be combined with renewal theory in order to prove delay instability results. Furthermore, we show how fluid approximations can be combined with stochastic Lyapunov theory in order to prove delay stability results. We illustrate the benefits of the proposed approach in two ways: (i) analytically, by providing a sharp characterization of the delay stability regions of networks with disjoint schedules, significantly generalizing previous results; (ii) computationally, through a Bottleneck Identification algorithm, which identifies (some) delay unstable queues by solving the fluid model of the network from certain initial conditions.

I. Introduction

We study scheduling problems arising in single-hop **switched queueing networks**, a class of stochastic systems that are often used to model the dynamics and decisions in data communication networks, e.g., cellular networks [20], Internet routers [19], and wireless ad hoc networks [28], but also in flexible manufacturing systems [11] and cloud computing clusters [16]. The salient feature of switched networks is that not all queues can be served at the same time, e.g., due to wireless interference constraints or due to matching constraints in a switch. Thus, only certain subsets of queues, the so-called schedules, can be served simultaneously, giving rise to a fundamental scheduling problem: which schedule to activate at each time slot? Clearly, the overall performance of the network depends critically on the scheduling policy applied.

The focus of this paper is on a well-studied queue-length-based scheduling policy, the **Max-Weight policy**. A remarkable property of the Max-Weight policy is its throughput

optimality, i.e., the ability to stabilize a queueing network whenever this is possible, without any explicit information on the traffic statistics [26]. Thus, dynamic instability phenomena, such as the ones reported by Kumar & Seidman [15] or Rybko & Stolyar [22], are avoided. Moreover, Max-Weight-type policies achieve very good delay performance under light-tailed traffic, e.g., they achieve optimal or order-optimal average delay for specific network topologies [12], [20], optimal large deviations exponent [28], and are asymptotically delay optimal in the heavy-traffic regime [25]. For these reasons, Max-Weight has become the "benchmark" policy for scheduling in switched networks. However, the delay performance of the Max-Weight policy in the presence of heavy-tailed traffic is not well-understood.

Empirical evidence of high variability phenomena in data communication networks [21], manufacturing [9], and cloud computing [8] motivates us to study switched networks with a **mix of heavy-tailed and light-tailed traffic**. The impact of heavy tails has been analyzed extensively in single or multi-server queues, e.g., see the survey paper [6] and the references therein, and more recently in monotone separable networks [1], a class of stochastic systems that includes multi-server queues, Jackson networks, and polling systems as special cases. Closer to our work come the papers by Borst *et al.* [4] and by Jagannathan *et al.* [14], both of which consider a system with two parallel queues, receiving heavy-tailed and light-tailed traffic, respectively, while sharing a single server. The authors determine the queue-length asymptotics of the Generalized Processor Sharing policy [4] and of the Generalized Max-Weight policy [14]. Also related to our work is the paper by Boxma *et al.* [5], which analyzes a M/G/2 queue with a heavy-tailed and a light-tailed server, and shows a dependence of the queue-length asymptotics on the arrival rate. The present paper builds upon our earlier work [18], which considers a single-hop switched queueing network with a mix of heavy-tailed and light-tailed traffic, under the Max-Weight policy (a brief discussion of the findings of [18] and the contributions of the present work can be found below).

For the purposes of this paper, heavy-tailed traffic is defined as an arrival process with infinite variance. Classical results in queueing theory (e.g., the Pollaczek-Khinchin formula, Kingman's bounds) imply that queues receiving heavy-tailed traffic are delay unstable, i.e., they experience infinite expected delays in steady state. Thus, we focus on the (policy-dependent) impact of heavy-tailed traffic on light-tailed queues, using delay stability as performance metric.

Motivating Example

We motivate the subsequent developments by presenting the main findings of [18] through the queueing system of Figure 1. This very simple switched network has two schedules, $\{1, 2\}$ and $\{3\}$, and queues are served at unit rate whenever the respective schedules are activated. We assume that the arrivals to queue 1 are heavy-tailed whereas the arrivals to queues 2 and 3 are light-tailed. In this setting the Max-Weight policy compares the length of queue 3 to the sum of the lengths of queues 1 and 2, and serves the "heavier" schedule.


Fig. 1. **Delay performance of the Max-Weight policy under heavy-tailed traffic through a simple example**. Queue 1 receives heavy-tailed traffic whereas queues 2 and 3 receive light-tailed traffic. Queues 1 and 2 are served simultaneously, whereas queue 3 can only be served alone. Max-Weight compares the length of queue 3 to the sum of the lengths of queues 1 and 2, and serves the "heavier" schedule. Queue 1 is delay unstable under any scheduling policy. Queue 3 is delay unstable under the Max-Weight policy. Finally, queue 2 may or may not be delay stable under the Max-Weight policy, depending on the arrival rates.

The delay stability of queue 1 does not depend on the specifics of the queueing system of Figure 1 or on the scheduling policy applied, but merely on the fact that it receives heavy-tailed traffic. Consider the best case scenario for queue 1, i.e., it is served at all times. In that case, queue 1 is equivalent to a M/G/1 queue with infinite second moment of service time. Then, classical results from queueing theory, e.g., the Pollaczek-Khinchin formula, imply that the expected steady-state delay is infinite, namely queue 1 is delay unstable even in the best case. This observation is the main idea behind Theorem 1 of [18]: in a switched queueing network, every heavy-tailed queue is delay unstable under any policy.

Coming to queue 3, notice that, under the Max-Weight policy, it receives no service unless its length is greater than or equal to the length of queue 1. However, queue 1 is, occasionally, very long due to its heavy-tailed arrivals. On those occasions, queue 3 has to build up to a similar length, leading to delay instability. This observation can be generalized (Theorem 2 of [18]): in a switched queueing network under the Max-Weight policy, every light-tailed queue that "conflicts" with a heavy-tailed queue is delay unstable.

An interesting finding of [18] concerns the delay stability of queue 2. One would expect that this queue is delay

stable: it is light-tailed itself, and is served together with a heavy-tailed queue, which should result in more service opportunities under Max-Weight. Surprisingly, there exist arrival rates in the stability region of the system, such that queue 2 is delay unstable. The key observation is as follows: even though queue 2 does not conflict with a heavy-tailed queue, it does conflict with queue 3, which is delay unstable because it conflicts with a heavy-tailed queue. Conversely, queue 2 is delay stable if its arrival rate is sufficiently low.

Proposition 1: (Rate-Dependent Delay Instability [18]) Consider the queueing system of Figure 1 under the Max-Weight policy. If the arrival rates satisfy $\lambda_2 > \left(1 + \lambda_1 - \lambda_3\right)/2$, then queue 2 is delay unstable.

Proposition 2: (Rate-Dependent Delay Stability [18]) Consider the queueing system of Figure 1, with arrival rates in the stability region and under the Max-Weight policy. If $\lambda_2 < \left(1 + \lambda_1 - \lambda_3\right)/2$ and the arrivals to queues 2 and 3 are exponential-type, then queue 2 is delay stable.

Propositions 1 and 2 provide a sharp characterization of the **delay stability region** of queue 2, i.e., the set of arrival rates for which queue 2 is delay stable. Previous proofs of these results were based on purely stochastic arguments, and were somewhat long and tedious. However, the main ideas behind them are rather simple and intuitive, and were presented in [18] through an informal "fluid approximation" to the stochastic system. The goal of this paper is to formalize this approach: we will show that the formal use of fluid approximations simplifies delay stability analysis, allowing us to generalize considerably the findings of [18].

Main Contributions

The main contribution of this paper is to show how fluid approximations can facilitate the delay analysis of switched networks with heavy-tailed traffic. More specifically, we show how fluid approximations can be combined with renewal theory in order to prove delay instability results (Theorem 1). Furthermore, we show how fluid approximations can be combined with stochastic Lyapunov theory in order to prove delay stability results (Theorem 2). Finally, we identify a class of piecewise linear Lyapunov functions, whose drift analysis can provide exponential bounds on queue-length asymptotics, in the presence of heavy-tailed traffic.

We illustrate the benefits of the proposed methodology in two ways:

(i) analytically, by providing an explicit characterization of the delay stability regions of switched networks with disjoint schedules, under the Max-Weight policy (Theorem 3). Moreover, we show that the propagation of delay instability to light-tailed queues is exacerbated close to certain boundaries of the stability region (Theorem 4). This is in sharp contrast to the asymptotic delay optimality of Max-Weight in the heavy-traffic regime and under diffusion scaling [25];

(ii) computationally, through a Bottleneck Identification algorithm, which identifies (some) delay unstable queues by solving, perhaps numerically, the fluid model of the network from specific initial conditions of interest.

Preview of the Paper

The remainder of the paper is organized as follows. We begin with a high-level discussion of the methodological challenges of the scheduling problem at hand, and the rationale behind our approach, in Section II. In Section III we provide a detailed description of a switched queueing network under the Max-Weight policy, together with some useful definitions and notation. In Section IV we present the fluid approximation of this queueing system. Sections V and VI summarize the technical contributions of the paper, namely how fluid approximations can facilitate the delay analysis of networks with heavy-tailed traffic. In Section VII we analyze the delay stability of networks with disjoint schedules. In Section VIII we introduce the Bottleneck Identification algorithm, accompanied by examples illustrating its applicability. We conclude the paper in Section IX with some brief remarks.

II. METHODOLOGICAL CHALLENGES AND CONTRIBUTIONS

The problem of delay analysis of the Max-Weight policy in the presence of heavy-tailed traffic poses a number of methodological challenges. First of all, Dynamic Programming or Markov Decision Problem formulations of scheduling problems in queueing systems are known to be analytically intractable, and to have prohibitive computational requirements. Moreover, Monte Carlo methods can be very slow to converge, or may even fail to converge at all, due to the very nature of heavy tails.

Regarding classical approaches in the queueing literature, a standard way of showing that queues exhibit large delays (e.g., lower bounds on queue-length/delay asymptotics or on the corresponding expected values) relies on sample path arguments. However, tracking the evolution of sample paths can be hard when the system exhibits complex dynamics, such as the ones imposed by Max-Weight. This also hinders the use of transform methods, at least as a way to obtain analytical results. The main idea of our approach is as follows: even when we are not able to analyze sample paths directly, we might still be able to do so approximately, in terms of the solution to the fluid model from certain initial conditions of interest. Then, we can use renewal theory to translate sample path analysis to lower bounds on queue-length asymptotics or steady-state moments.

On the other hand, showing that queues exhibit low delays (e.g., upper bounds on queue-length/delay asymptotics or on the corresponding expected values) is usually based on drift analysis of suitable Lyapunov functions, since coupling arguments can only be applied to systems with relatively simple dynamics. Unfortunately, popular candidates such as standard piecewise linear functions [3], [7], quadratic functions [26], and norms [23], [28] cannot be used under heavy-tailed traffic. This is because the steady-state expectation of these functions is infinite, rendering their drift analysis uninformative. Additionally, drift analysis can be a challenge by itself under stochastic dynamics. Our approach to this problem is as follows: we identify a class of piecewise linear Lyapunov functions that are nonincreasing in the length of the heavy-tailed queues, and which are suitable for performance analysis of queueing systems with a mix of heavy-tailed and exponential-type traffic (more specifically, for obtaining exponential upper bounds on queue-length asymptotics). Moreover, we show how fluid approximations can simplify the drift analysis of this class of Lyapunov functions. Critical to the latter is a connection that we establish between fluid approximations and Lyapunov theory: suppose $V(\cdot)$ is continuous, piecewise linear, and a "Lyapunov function" for the fluid model; then $V(\cdot)$ is also a "Lyapunov function" for the original queueing system.[1] Moreover, if $V(\cdot)$ has exponential-type "upward jumps" in the stochastic system, then the results of [13] imply an exponential upper bound for its steady-state distribution.

III. THE MODEL

In this section we give a detailed presentation of the queueing model under consideration, together with some necessary definitions and notation.

We denote by $\mathbb{R}_+$, $\mathbb{Z}_+$, and $\mathbb{N}$ the sets of nonnegative reals, nonnegative integers, and positive integers, respectively. The cartesian products of M copies of $\mathbb{R}_+$ and $\mathbb{Z}_+$ are denoted by $\mathbb{R}_+^M$ and $\mathbb{Z}_+^M$, respectively.

We consider a discrete time switched queueing network, where arrivals occur at the end of each time slot. Let $\mathcal{F} = \{1, \ldots, F\}$, $F \in \mathbb{N}$. Central to our model is the notion of a **traffic flow** $f \in \mathcal{F}$, which is a long-lived stream of traffic that arrives to the network according to a discrete time stochastic arrival process $\{A_f(t); \ t \in \mathbb{Z}_+\}$. We assume that all arrival processes take values in $\mathbb{Z}_+$, and are independent and identically distributed (IID) over time. Furthermore, different arrival processes are mutually independent. We denote by $\lambda_f = \mathbb{E}[A_f(0)] > 0$ the arrival rate of traffic flow f and by $\lambda = (\lambda_f; \ f = 1, \ldots, F)$ the vector of arrival rates of all traffic flows.

Definition 1: (Heavy/Light Tails) A nonnegative random variable X is heavy-tailed if $\mathbb{E}[X^2]$ is infinite, and is light-tailed otherwise. Moreover, X is exponential-type if there exists $\theta > 0$ such that $\mathbb{E}[\exp(\theta X)] < \infty$.

We define similarly a heavy-tailed/light-tailed/exponential-type traffic flow.

There are several definitions of heavy/light tails in the literature. In fact, a random variable is often defined as light-tailed if it is exponential-type, and heavy-tailed otherwise. The definition adopted in this paper has been used in the area of data communication networks, e.g., see [21], due to its close connection to long-range dependence.

For technical reasons we assume the existence of some $\gamma \in (0, 1)$ such that $\mathbb{E}[A_f^{1+\gamma}(0)] < \infty$, for all $f \in \mathcal{F}$.

---

[1]This is to be contrasted to the more common use of fluid approximations for stability analysis, where a Lyapunov function $V(\cdot)$ for the fluid model implies the existence of **another** Lyapunov function $G(\cdot)$ for the stochastic model.

We consider single-hop traffic flows, i.e., the traffic of flow f is buffered in a dedicated single-server queue (queue f and server f, henceforth), eventually gets served, and then exits the system. Our modeling assumptions imply that the set of traffic flows can be identified with the set of queues and the set of servers of the network. The service discipline within each queue is assumed to be "First Come, First Served." The stochastic process $\{Q_f(t);\ t \in \mathbb{Z}_+\}$ captures the evolution of the length of queue f. Since our main motivation comes from data communication networks, $A_f(t)$ will be interpreted as the number of packets that queue f receives at the end of time slot t, and $Q_f(t)$ as the total number of packets in queue f at the beginning of time slot t. The arrivals and the lengths of the various queues at time slot t are captured by the vectors $A(t) = \big(A_f(t);\ f = 1, \ldots, F\big)$ and $Q(t) = \big(Q_f(t);\ f = 1, \ldots, F\big)$, respectively.

In the context of data communication networks, a batch of packets arriving to a queue at any given time slot can be viewed as a single entity, e.g., as a file that needs to be transmitted. We define the **end-to-end delay of a file** of flow f to be the number of time slots that the file spends in the network, starting from the time slot right after it arrives at queue f, until the time slot that its last packet gets served. For $k \in \mathbb{N}$, we denote by $D_f(k)$ the end-to-end delay of the k^{th} file of flow f, and use the vector notation $D(k) = \big(D_f(k);\ f = 1, \ldots, F\big)$.

The salient feature of a switched queueing network is that not all servers can be simultaneously active, e.g., due to interference in wireless networks or matching constraints in a switch. Consequently, not all traffic flows can be served simultaneously. A set of traffic flows that can be served simultaneously is called a **schedule**. We denote by $\mathcal{S}$ the set of all schedules, which is assumed to be an arbitrary subset of the powerset of $\mathcal{F}$. For simplicity, we assume that all packets have the same size, and that the service rate of all servers is equal to one packet per time slot. We denote by $S_f(t) \in \{0, 1\}$ the number of packets that are scheduled for service from queue f at time slot t. Note that this is not necessarily equal to the number of packets that are actually served, because the queue may be empty. We use the vector notation $S(t) = \big(S_f(t);\ f = 1, \ldots, F\big)$. For convenience, we also identify schedules with vectors in $\{0, 1\}^F$.

Using the notation above, the **dynamics** of queue f take the form

$$Q_f(t + 1) = Q_f(t) + A_f(t) - S_f(t) \cdot 1_{\{Q_f(t) > 0\}},$$

for all $t \in \mathbb{Z}_+$, where $1_{\{Q_f(t) > 0\}}$ denotes the indicator function of the event $\{Q_f(t) > 0\}$. The vector of initial queue lengths $Q(0)$ is assumed to be an arbitrary element of $\mathbb{Z}_+^F$.

The service vector $S(t)$ is determined by the scheduling policy applied to the network. In this paper we focus on the **Max-Weight policy**, where the scheduling vector $S(t)$ satisfies

$$S(t) \in \arg \max_{(S_f) \in \mathcal{S}} \Big\{ \sum_{f \in \mathcal{F}} Q_f(t) \cdot S_f \Big\},$$

at any given time slot. If the set on the right-hand side includes multiple schedules, then one of them is chosen uniformly at random.

As alluded to in the Introduction, a very appealing property of the Max-Weight policy is throughput optimality. Before we can state this property in mathematical terms, we need to define formally the notion of stability of a queueing network.

Definition 2: (Stability) The switched queueing network described above is stable if the vector-valued sequences $\{Q(t);\ t \in \mathbb{Z}_+\}$ and $\{D(k);\ k \in \mathbb{N}\}$ converge in distribution, and their limiting distributions do not depend on the initial queue lengths $Q(0)$.

Under a stabilizing scheduling policy, we denote by $Q = \big(Q_f;\ f = 1, \ldots, F\big)$ and $D = \big(D_f;\ f = 1, \ldots, F\big)$ generic random vectors distributed according to the limiting distributions of $\{Q(t);\ t \in \mathbb{Z}_+\}$ and $\{D(k);\ k \in \mathbb{N}\}$, respectively. We refer to Q_f as the steady-state length of queue f. Similarly, we refer to D_f as the steady-state delay of a file of traffic flow f.

The ability to stabilize a switched queueing network depends on the arrival rates of the various traffic flows relative to the service rates of the servers, and on the scheduling constraints. This relation is captured by the stability region of the network.

Definition 3: (Stability Region) The stability region of the switched queueing network described above is the set of arrival rate vectors

$$\Big\{ \lambda \in \mathbb{R}_+^F \ \Big|\ \exists\ \zeta_s \in \mathbb{R}_+,\ s \in \mathcal{S} :\ \lambda \le \sum_{s \in \mathcal{S}} \zeta_s \cdot s,\ \sum_{s \in \mathcal{S}} \zeta_s < 1 \Big\}.$$

Lemma 1: (Throughput Optimality of Max-Weight) Consider the switched queueing network described above under the Max-Weight policy. The network is stable for any arrival rate vector in the stability region.

Proof: For the case of light-tailed traffic, this result follows from the findings in [26]; in the presence of heavy-tailed traffic, it follows from Proposition 2 of [25]. For a formal proof the reader is referred to Lemma 4.1 in [17]. ∎

Finally, we define the property that we use to evaluate the delay performance of Max-Weight.

Definition 4: (Delay Stability) Traffic flow f is delay stable, under a specific scheduling policy, if the switched queueing network is stable under that policy and $\mathbb{E}[D_f]$ is finite; otherwise, f is delay unstable.

IV. FLUID APPROXIMATION OF A SWITCHED NETWORK UNDER THE MAX-WEIGHT POLICY

In this section we give some background material on the natural fluid model of a single-hop switched queueing network under the Max-Weight policy.

The fluid model is a deterministic dynamical system, which aims to capture the evolution of its stochastic counterpart on longer time scales by taking advantage of Laws of Large Numbers. Initially, we give a brief description and

some useful properties of the fluid model. Then, we introduce the notion of fluid scaling, and establish a formal connection between the deterministic and the scaled stochastic system.

The **Fluid Model (FM)** of a single-hop switched queueing network under the Max-Weight policy is defined by the set of ordinary differential equations Eqs. (1)-(6), for every time $t \geq 0$ for which the derivatives exist (such t is often called a regular time):

$$\dot{q}_f(t) = \lambda_f - \sum_{\pi \in \mathcal{S}} \dot{s}_\pi(t)\pi_f + \dot{y}_f(t), \qquad \forall f \in \mathcal{F}; \qquad (1)$$

$$\dot{s}_\pi(t) \geq 0, \qquad \forall \pi \in \mathcal{S}; \qquad (2)$$

$$\sum_{\pi \in \mathcal{S}} \dot{s}_\pi(t) = 1; \qquad (3)$$

$$0 \leq \dot{y}_f(t) \leq \sum_{\pi \in \mathcal{S}} \dot{s}_\pi(t)\pi_f, \qquad \forall f \in \mathcal{F}; \qquad (4)$$

$$q_f(t) > 0 \implies \dot{y}_f(t) = 0, \qquad \forall f \in \mathcal{F}; \qquad (5)$$

$$\sum_{f \in \mathcal{F}} q_f(t)\pi_f < \max_{\sigma \in \mathcal{S}}\left\{ \sum_{f \in \mathcal{F}} q_f(t)\sigma_f \right\}$$
$$\implies \dot{s}_\pi(t) = 0, \qquad \forall \pi \in \mathcal{S}. \quad (6)$$

In the equations above, $q(t)$ represents the vector of queue lengths at time t, $y(t)$ represents the vector of cumulative idling/wasted service up to time t, and $s_\pi(t)$ represents the total amount of time that schedule π has been activated up to time t. Eq. (3) states that the scheduling policy is "work-conserving," and Eq. (5) that there can be no wasted service when queue lengths are positive. Finally, Eq. (6) is the natural analogue of the Max-Weight policy in the fluid domain: schedules that do not have maximum weight receive no service.

The above differential equations have appeared in previous studies, e.g., see [24], [25].

Fix arbitrary $T > 0$. A **Fluid Model Solution (FMS)** from initial condition $q(0) = q$ is a Lipschitz continuous function $x(\cdot) = \big(q(\cdot), y(\cdot), s(\cdot)\big)$ that satisfies: (i) $x(0) = (q, 0, 0)$; (ii) $q(t) \geq 0$, for all $t \in [0, T]$; (iii) Eqs. (1)-(6) over the interval $[0, T]$.

A FMS is differentiable almost everywhere (equivalently, almost every $t \in [0, T]$ is a regular time), since it is Lipschitz continuous by assumption.

Now we define the notion of fluid scaling, and we establish the existence of a fluid limit and of a FMS. Consider a sequence of initial queue lengths $\{Q^b(0); \ b \in \mathbb{N}\}$ for the queueing system of Section III, and the corresponding sequence of queue-length processes $\{Q^b(\cdot); \ b \in \mathbb{N}\}$.

We define the "fluid-scaled" queue-length process as

$$\tilde{q}^b(t) = \frac{Q^b(bt)}{b}, \qquad t \in [0, T], \qquad b \in \mathbb{N}.$$

We assume the existence of a vector $q \in \mathbb{R}_+^F$ and of a sequence of positive numbers $\{\epsilon_b; \ b \in \mathbb{N}\}$, converging to zero as b goes to infinity, that satisfy

$$\max_{f \in \mathcal{F}} \big| \tilde{q}_f^b(0) - q_f \big| \leq \epsilon_b, \qquad \forall b \in \mathbb{N}.$$

We recall our standing assumption from Section III that there exists $\gamma \in (0, 1)$ so that all traffic flows have $(1 + \gamma)$ moments. Fix some $\gamma' \in (0, \gamma)$ and consider the sequence of sets of sample paths of the arrival processes defined by

$$H_b = \left\{ \omega : \ \sup_{1 \leq t \leq bT} \frac{1}{bT} \max_{f \in \mathcal{F}} \Big| \sum_{\tau=0}^{t-1} A_f(\tau) - \lambda_f \Big| < (bT)^{-\frac{\gamma'}{1+\gamma}} \right\},$$

$b \in \mathbb{N}$. Intuitively, H_b contains those sample paths of the arrival processes that stay close to their average behavior over the time interval $[0, bT]$.

Lemma 2: (Existence of Fluid Limit and FMS) There exists a Lipschitz continuous function $z(t) = \big(z_1(t), \ldots, z_F(t)\big)$, $t \in [0, T]$, such that for every $\epsilon > 0$ there exists $b_0(\epsilon)$ so that

$$\mathbb{P}\big(H_b\big) \geq 1 - \epsilon, \quad \forall b \geq b_0(\epsilon),$$

and

$$\sup_{t \in [0,T]} \max_{f \in \mathcal{F}} \big| \tilde{q}_f^b(t) - z_f(t) \big| \leq \epsilon, \quad \forall \omega \in H_b, \quad \forall b \geq b_0(\epsilon).$$

Additionally, there exist Lipschitz continuous functions $v(\cdot)$ and $w(\cdot)$, such that $\big(z(\cdot), v(\cdot), w(\cdot)\big)$ is a FMS from initial condition $q(0) = q$.

Proof: The reader is referred to Lemma 5.1 of [17]. ∎

Lemma 3: (Uniqueness and Continuity of FMS) For any given $q \in \mathbb{R}_+^F$ there exists a unique Lipschitz continuous function $z(t) = \big(z_1(t), \ldots, z_F(t)\big)$, $t \in [0, T]$, such that the queue-length part of every FMS from initial condition q is $z(\cdot)$. Moreover, $z(\cdot)$ depends continuously on both the initial condition q and the arrival rate vector λ.

Proof: The reader is referred to Lemma 5.2 of [17]. ∎

V. DELAY INSTABILITY RESULTS VIA FLUID APPROXIMATIONS

Next, we illustrate how fluid approximations can be used for proving delay instability results. A standard way for showing large delays (e.g., lower bounds on the tail asymptotics or on the corresponding expected values) is through sample path arguments. However, tracking the evolution of sample paths can be hard when the system exhibits complex dynamics. The main idea of our approach is to analyze sample paths approximately, as the solution to the FM from certain initial conditions of interest, and then to translate the findings to delay instability results through renewal theory.

Our contribution is summarized in the following theorem, which provides a sufficient condition for the delay instability of traffic flows.

Theorem 1: Consider the single-hop switched queueing network of Section III under the Max-Weight policy, and its natural FM of Section IV. Let $h \in \mathcal{F}$ be a heavy-tailed traffic flow, and $q^*(\cdot)$ be the (unique) queue-length part of a FMS from initial condition $q_h^*(0) = 1$ and $q_f^*(0) = 0$, for all $f \neq h$. If there exists $\tau \in [0, T]$ such that $q_j^*(\tau) > 0$, then traffic flow j is delay unstable.

Proof: (Outline) Suppose that there exists $\tau \in [0,T]$ such that $q_j^*(\tau) > 0$. Then, the existence of a fluid limit, which also guarantees the existence of a FMS, and the uniqueness of the queue-length part of a FMS imply that, after a big arrival to queue h, queue j builds to the order of magnitude of the heavy-tailed queue with high probability. In turn, renewal theory and Little's Law provide the desired delay instability result. For a formal proof the reader is referred to Theorem 5.1 of [17]. ∎

Remark: Theorem 1 holds for any choice of $T > 0$ (the horizon of the FMS). However, the fact that a single-hop switched queueing network is stable under the Max-Weight policy (Lemma 1) implies the existense of some $T^* > 0$, proportional to the initial condition of the FMS, such that $q(t) = 0$, for all $t > T^*$. Consequently, the most effective application of Theorem 1 is when T is chosen large enough so that the FMS "drains" within this horizon.

VI. Delay Stability Results via Fluid Approximations

In this section we shift our attention to delay stability results in networks that receive a mix of heavy-tailed and exponential-type traffic. Typically, proving low delays is either based on coupling arguments, if the underlying dynamics are relatively simple, or, more often, on drift analysis of suitable Lyapunov functions. We focus on the latter approach. The presence of heavy-tailed traffic, though, introduces an additional complication: popular candidate Lyapunov functions such as standard piecewise linear functions [3], [7], quadratic functions [26], and norms [23], [28] cannot be used because the steady-state expectation of these functions is infinite under heavy-tailed traffic, rendering drift analysis uninformative.

We introduce a class of piecewise linear Lyapunov functions that are nonincreasing in the length of the heavy-tailed queues, and which can provide exponential upper bounds on queue-length asymptotics despite the presence of heavy-tailed traffic. However, drift analysis of piecewise linear Lyapunov functions is sometimes a challenge by itself, due to the fact that the stochastic descent property is often lost at locations where the function is nondifferentiable. This difficulty can be handled by either smoothing the Lyapunov function, e.g., as in [7], or by showing that the stochastic descent property still holds if we look ahead a sufficiently large number of time slots, e.g., as in [27]. We follow the second approach, and show how fluid approximations can simplify significantly drift analysis of this class of functions.

Theorem 2: Consider the single-hop switched queueing network of Section III under the Max-Weight policy, and its natural FM of Section IV. Consider a function $V : \mathbb{R}_+^F \to \mathbb{R}_+$ of the form

$$V(x) = \max_{j \in \mathcal{J}} \left\{ \sum_{f \in \mathcal{F}} c_{jf} x_f \right\},$$

where $\mathcal{J} = \{1, \ldots, J\}$, and $c_{jf} \in \mathbb{R}$, for all $j \in \mathcal{J}$, $f \in \mathcal{F}$. Suppose that:

(i) there exists $l > 0$ such that, for every initial condition $q(0)$ and regular time $t \geq 0$, we have $\dot{V}(q(t)) \leq -l$, whenever $V(q(t)) > 0$;

(ii) if $c_{jf} > 0$, for some $j \in \mathcal{J}$, then traffic flow f is exponential-type.

Then,

(i) the sequence $\{ V(Q(t)); \ t \in \mathbb{Z}_+ \}$ converges in distribution to random variable $V(Q)$;

(ii) $\mathbb{E}\left[\exp\left(\theta V(Q) \right) \right] < \infty$, for some $\theta > 0$.

Proof: See Theorem 5.2 of [17]. ∎

VII. Delay Stability Regions of Networks with Disjoint Schedules

In this section we consider a single-hop switched queueing network with disjoint schedules, i.e., each traffic flow belongs to exactly one schedule, and we provide a sharp characterization of the delay stability regions of the Max-Weight policy in the presence of heavy-tailed traffic. To achieve this, we apply the connections between fluid approximations and delay stability/instability that we established in Theorems 1 and 2.

More specifically, we consider a system with $(K + 1)$ schedules, which we denote by $\sigma_0, \sigma_1, \ldots, \sigma_K$. Schedule σ_0 includes $(F_0 + 1)$ queues that we denote by (σ_0, f), $f = 0, \ldots, F_0$; schedule σ_k, $k = 1, \ldots, K$, includes F_k queues that we denote by (σ_k, f), $f = 1, \ldots, F_k$. Note that $\sigma_k \cap \sigma_l = \emptyset$, if $k \neq l$. Since the system has single-hop traffic, we use the notions of queue and traffic flow interchangeably.

We denote the arrival rate to queue (σ_k, f) by $\lambda_f^{\sigma_k}$. The characterization of the delay stability regions hinges on the assumption that the arrival rates are distinct. Without loss of generality, we also assume that queues are indexed in descending order of arrival rates, except, possibly, for queue $(\sigma_0, 0)$ (the significance of the latter queue will become obvious shortly). In mathematical terms,

$$\lambda_f^{\sigma_k} > \lambda_{f+1}^{\sigma_k}, \quad f \in \{1, \ldots, F_k - 1\}, \ \forall k \in \{0, \ldots, K\}. \tag{7}$$

At each time slot at most one schedule can be activated. If a schedule is activated then one packet is removed from all nonempty queues of that schedule.

We assume that the arriving traffic is in the stability region of the system, which implies that

$$\max_{f=0,\ldots,F_0} \left\{ \lambda_f^{\sigma_0} \right\} + \sum_{k=1}^{K} \max_{f=1,\ldots,F_k} \left\{ \lambda_f^{\sigma_k} \right\} < 1. \tag{8}$$

Traffic flow $(\sigma_0, 0)$ is heavy-tailed whereas every other traffic flow is exponential-type. Theorem 2 of [18] implies that, under the Max-Weight policy, every traffic flow that does not belong to schedule σ_0 is delay unstable, for any positive arrival rate, because it conflicts with $(\sigma_0, 0)$. In contrast, we expect traffic flows (σ_0, f), $f = 1, \ldots, F_0$ to have nontrivial delay stability regions, since they do not conflict with the heavy-tailed flow $(\sigma_0, 0)$; much similar to the 3-queue system in Figure 1, where traffic flow 2 had a nontrivial delay stability region.

Theorem 3: Consider the single-hop switched queueing network with disjoint schedules described above under the Max-Weight policy, and arrival rates satisfying Eqs. (7) and (8). The conditions

$$\lambda_f^{\sigma_0} \leq \frac{1}{1+Kf}\Big(1 + K\sum_{l=0}^{f-1}\lambda_l^{\sigma_0} - \sum_{k=1}^{K}\lambda_1^{\sigma_k}\Big), \quad f = j,\ldots,F_0,$$
(9)

are necessary for the delay stability of traffic flow (σ_0, j), $j \in \{1,\ldots,F_0\}$. If all inequalities are strict, then these conditions are also sufficient for delay stability, and the steady-state length of the associated queue is exponential-type.

Proof: (Outline) The necessity part relies on Theorem 1, and on showing that the length of queue (σ_0, j), along the FMS from initial condition one for queue $(\sigma_0, 0)$, becomes positive if one of the above conditions is violated. The sufficiency part is based on drift analysis of the piecewise linear Lyapunov function

$$V(q) = \sum_{f=j}^{F_0} c_f q_f^{\sigma_0} + \max_{k=1,\ldots,K}\Big\{\Big[\sum_{f=1}^{F_k} q_f^{\sigma_k} - \sum_{f=0}^{F_0} q_f^{\sigma_0}\Big]^+\Big\},$$

where $c_f \in (0,1)$, for all $f \in \{j,\ldots,F_0\}$, and subsequent use of Theorem 2. For a formal proof the reader is referred to Theorem 3 in [17]. ∎

Remark: The delay stability of traffic flow (σ_0, j) when one of the above conditions holds with equality will depend, in general, on higher order moments of the arrivals, and not just the rates. To see this, suppose that a large batch of b packets arrives to the heavy-tailed queue $(\sigma_0, 0)$. A random-walk-type argument can show that queue (σ_0, j) will build up to $\Omega(\sqrt{b})$ during an $\Omega(b)$ time interval, assuming that the configuration corresponding to the equality condition is reached. Thus, the aggregate length of this queue over a busy period will be $\Omega(b^{3/2})$, which implies that the delay stability of traffic flow (σ_0, j) will depend on the 1.5 moment of the arrivals to the heavy-tailed queue.

Theorem 4: Consider the single-hop switched queueing network with disjoint schedules described above under the Max-Weight policy. There exist arrival rates satisfying Eqs. (8) such that all traffic flows in the network are delay unstable.

Proof: (Outline) The proof reduces to showing that the lengths of all light-tailed queues of schedule σ_0, along the FMS from initial condition one for queue $(\sigma_0, 0)$, become positive for the set of arrival rates

$$\lambda_0^{\sigma_0} = \lambda_f^{\sigma_k} = \epsilon > 0, \quad \forall f \in \{1,\ldots,F_k\}, \quad \forall k \in \{1,\ldots,K\},$$

and

$$\lambda_f^{\sigma_0} = 1 - (K+1)\epsilon, \qquad \forall f \in \{1,\ldots,F_0\},$$

when ϵ is sufficiently small. Then, the result is implied directly from Theorem 1. For a formal proof the reader is referred to Theorem 4 in [17]. ∎

VIII. THE BOTTLENECK IDENTIFICATION ALGORITHM

Theorem 1 provides a sufficient condition for the delay instability of traffic flows, based on the FMS from a specific initial condition. The following algorithmic procedure, which we term the **Bottleneck Identification (BI) algorithm**, tests this condition for all initial conditions of interest.

INITIALIZATION: $U = \emptyset$
REPEAT
For every heavy-tailed traffic flow $h \in \mathcal{F}$,
   (i) solve the FM with initial condition one for queue h, and zero for all other queues;
   (ii) find the set of queues that become positive at any point before the FMS drains, U_h;
   (iii) set $U = U \cup U_h$;
END

Clearly, upon termination of the algorithm, all queues/flows included in U are delay unstable.

Below we present concrete examples that illustrate the use of the BI algorithm.

3×3 Switch

Consider a 3×3 input-queued switch under the Max-Weight policy. This is a system of 9 queues indexed by (i,j), $i = 1,\ldots,3$, $j = 1,\ldots,3$, with index i representing the input port and index j the output port of the switch. A schedule is a matching between input and output ports. Namely, the set of schedules is as follows:

$$\mathcal{S} = \Big\{\{(1,1),(2,2),(3,3)\}, \{(1,1),(2,3),(3,2)\},$$
$$\{(1,2),(2,1),(3,3)\}, \{(1,2),(2,3),(3,1)\},$$
$$\{(1,3),(2,1),(3,2)\}, \{(1,3),(2,2),(3,1)\}\Big\}.$$

The 3×3 input-queued switch is a network with non-disjoint schedules, so an explicit characterization of its delay stability regions is not available.

Consider the set of arrival rates $\lambda_{11} = 0.1$, $\lambda_{12} = 0.1$, $\lambda_{13} = 0.1$, $\lambda_{21} = 0.1$, $\lambda_{22} = 0.38$, $\lambda_{23} = 0.4$, $\lambda_{31} = 0.1$, $\lambda_{32} = 0.42$, and $\lambda_{33} = 0.44$. Note that this set of rates satisfies $\sum_i \lambda_{ij} < 1$ and $\sum_j \lambda_{ij} < 1$, so that the system is stable under the Max-Weight policy [19].

We assume that traffic flow $(1,1)$ is heavy-tailed, while all other traffic flows are light-tailed. We are interested in the delay stability of flows (2,2), (2,3), (3,2), (3,3); these are the flows that do not conflict with flow (1,1). Figure 2 shows the FMS for the considered set of rates, and with initial condition one for queue (1,1), and zero for all other queues (we present only the queues of interest). The lengths of all queues of interest become positive before the FMS drains, so according to Theorem 1 they are delay unstable.

3×3 Grid Network

Consider the 3×3 grid network depicted in Figure 3 under the Max-Weight policy. This system represents a wireless network with interference constraints. Queues are identified with links and are indexed by $i = 1,\ldots,12$. As soon as a packet is transmitted through the respective wireless link,

Fig. 2. The FMS of a 3×3 switch, with initial condition one for queue (1,1) and zero for the other queues, and arrival rates $\lambda_{11} = 0.1$, $\lambda_{12} = 0.1$, $\lambda_{13} = 0.1$, $\lambda_{21} = 0.1$, $\lambda_{22} = 0.38$, $\lambda_{23} = 0.4$, $\lambda_{31} = 0.1$, $\lambda_{32} = 0.42$, $\lambda_{33} = 0.44$.

it exits the system. We assume the two-hop interference constraint model, i.e., if a wireless link is transmitting, all links in a two-hop distance must idle. This implies that the set of schedules is as follows:

$$\mathcal{S} = \Big\{ \{1, 11\}, \{1, 12\}, \{1, 10\}, \{2, 8\}, \{2, 11\},$$
$$\{2, 12\}, \{3, 5\}, \{3, 10\}, \{3, 12\}, \{4\},$$
$$\{5, 8\}, \{5, 11\}, \{6\}, \{7\}, \{8, 10\}, \{9\} \Big\}.$$

Again, this is a network with non-disjoint schedules, so an explicit characterization of its delay stability regions is not available.


Fig. 3. A 3×3 grid wireless network with two-hop interference constraints.

Consider the set of arrival rates $\lambda_1 = 0.01$, $\lambda_2 = 0.02$, $\lambda_3 = 0.03$, $\lambda_4 = 0.04$, $\lambda_5 = 0.05$, $\lambda_6 = 0.06$, $\lambda_7 = 0.07$, $\lambda_8 = 0.08$, $\lambda_9 = 0.09$, $\lambda_{10} = 0.1$, $\lambda_{11} = 0.11$, and $\lambda_{12} = 0.12$. It can be verified that this set of rates is in the stability region of the system.

We assume that traffic flow 1 is heavy-tailed, while all other traffic flows are light-tailed. We are interested in the delay stability of traffic flows 10, 11, and 12, since these flows do not conflict with flow 1. Figure 4 shows the FMS for the considered set of rates, and with initial condition one for queue 1, and zero for all other queues (we present only the queues of interest). The lengths of all queues of interest become positive before the FMS drains, so according

to Theorem 1 they are delay unstable.

Fig. 4. The FMS of the 3×3 grid network of Figure 3, with initial condition one for queue 1 and zero for the other queues, and arrival rates $\lambda_1 = 0.01$, $\lambda_2 = 0.02$, $\lambda_3 = 0.03$, $\lambda_4 = 0.04$, $\lambda_5 = 0.05$, $\lambda_6 = 0.06$, $\lambda_7 = 0.07$, $\lambda_8 = 0.08$, $\lambda_9 = 0.09$, $\lambda_{10} = 0.1$, $\lambda_{11} = 0.11$, $\lambda_{12} = 0.12$.

IX. DISCUSSION

This paper built on, and extended the results of, our earlier work [18]. More specifically, we studied a single-hop switched queueing network with a mix of heavy-tailed and light-tailed traffic, and carried out a delay stability analysis of the Max-Weight policy. Our goal was to showcase the use of fluid approximations in showing both delay insta-bility (combined with renewal theory) and delay stability (combined with stochastic Lyapunov theory). Moreover, we applied these results to get a sharp characterization of the delay stability regions of the Max-Weight policy in networks with disjoint schedules, generalizing the findings of [18].

We conclude the paper with some brief remarks.

The use of fluid approximations in delay analysis of queueing systems with heavy-tailed traffic is not new. For example, Baccelli *et al.* [2] have used fluid models to determine the precise tail asymptotics of the steady-state maximal dater (i.e., the time to clear all customers present at time t, assuming arrivals are stopped from that point on, in the limit as t goes to infinity) of generalized Jackson networks with subexponential service times. The tail asymptotics are determined through a sample path construction of the maximal dater, which preserves crucial monotonicity properties of Jackson networks. In contrast, our approach is based on renewal theory, which on one hand does not provide as refined results (moment bounds instead of the precise asymptotics), but on the other hand does not rely on any special structure, besides the regenerative property, making it easier to apply.

Theorems 1 and 2 are stated and proved in the context of a single-hop switched queueing network under the Max-Weight policy. However, a closer look at their proofs reveals that the properties that we are really leveraging are:

(i) the existence of a fluid limit;

(ii) the uniqueness of the fluid model solution;

(iii) the existence of a suitable Lyapunov function for the fluid model.

Thus, Theorems 1 and 2 can be easily extended to any Markovian queueing system for which properties (i)-(iii) hold. Regarding properties (i) and (ii), which would allow the generalization of Theorem 1, one could potentially take advantage of the extensive literature on fluid approximations that has developed over the last 20 years. Proving property (iii), though, brings up several questions: given a queueing system, is a there a systematic and efficient algorithm for constructing a Lyapunov function? And even if there exists such an algorithm, can we ensure that the Lyapunov function is piecewise linear and nonincreasing in the lengths of the heavy-tailed queues? The undecidability results in [10] suggest that the answers to the above questions could very well be negative in the general case. So, in order to extend Theorem 2, one would probably have to restrict to systems with special structure.

In Theorem 4 we showed that, in every network with disjoint schedules that receives heavy-tailed traffic and operates under the Max-Weight policy, there exists an arrival rate vector in the corresponding stability region for which all flows are delay unstable. We conjecture that the same is true for every single-hop switched queueing network, i.e., with, possibly, overlapping schedules. Moreover, the rate vectors of interest are always close to certain boundaries of the stability regions, suggesting "heavily loaded" systems. Thus, at a conceptual level, our result is in sharp contrast to the asymptotic optimality of the Max-Weight policy in the heavy-traffic regime [25]. The latter result, of course, concerns the diffusion-scaled queue-length processes, whereas our findings apply to the unscaled processes.

Finally, we conjecture that, for the Max-Weight policy, the BI algorithm identifies all delay unstable traffic flows in the network, except, possibly, for the case of arrival rate vectors on the boundaries of delay stability regions. A closer look at the proof of Theorem 3 reveals that this conjecture is true for networks with disjoint schedules. The general case is expected to be more challenging, since certain monotonicity properties of networks with disjoint schedules no longer hold. However, affirmative resolution of this conjecture would reduce the problem of delay stability analysis to solving (perhaps numerically) a system of ordinary differential equations from certain initial conditions.

REFERENCES

[1] F. Baccelli, S. Foss (2004). Moments and tails in monotone-separable stochastic networks. *The Annals of Applied Probability*, 14(2), pp. 612-650.

[2] F. Baccelli, S. Foss, M. Lelarge (2005). Tails in generalized Jackson networks with subexponential service-time distributions. *Journal of Applied Probability*, 42(2), pp. 513-530.

[3] D. Bertsimas, D. Gamarnik, J.N. Tsitsiklis (2001). Performance of multiclass Markovian queueing networks via piecewise linear Lyapunov functions. *The Annals of Applied Probability*, 11(4), pp. 1384-1428.

[4] S. Borst, M. Mandjes, M. van Uitert (2003). Generalized processor sharing with light-tailed and heavy-tailed input. *IEEE/ACM Transactions on Networking*, 11(5), pp. 821-834.

[5] O. Boxma, Q. Deng, B. Zwart (2002). Waiting-time asymptotics for the M/G/2 queue with heterogeneous servers. *Queueing Systems*, 40(1), pp. 5-31.

[6] O. Boxma, B. Zwart (2007). Tails in scheduling. *Performance Evaluation Review*, 34(4), 13-20.

[7] D. Down, S.P. Meyn (1997). Piecewise linear test functions for stability and instability of queueing networks. *Queueing Systems*, 27(3-4), pp. 205-226.

[8] D. Ersoz, M. Yousif, and C. Das (2007). Characterizing network traffic in a cluster-based, multi-tier data center. In: *Proceedings of the 27^{th} International Conference on Distributed Computing Systems*.

[9] M. Fisher, A. Raman (1996). Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research*, 44(1), pp. 87-99.

[10] D. Gamarnik, D. Katz (2009). On deciding stability of multiclass queueing networks under buffer priority scheduling policies. *The Annals of Applied Probability*, 19(5), pp. 2008-2037.

[11] N. Gans, G. van Ryzin (1997). Optimal control of a multi-class, flexible queueing system. *Operations Research*, 45(5), pp. 677-693.

[12] A. Ganti, E. Modiano, J.N. Tsitsiklis (2007). Optimal transmission scheduling in symmetric communication models with intermittent connectivity. *IEEE Transactions on Information Theory*, 53(3), pp. 998-1008.

[13] B. Hajek (1982). Hitting-time and occupation-time bounds implied by drift analysis with applications. *Advances in Applied Probability*, 14(3), pp. 502-525.

[14] K. Jagannathan, M.G. Markakis, E. Modiano, J.N. Tsitsiklis (2012). Queue-length asymptotics for generalized Max-Weight scheduling in the presence of heavy-tailed traffic. *IEEE/ACM Transactions on Networking*, 20(4), pp. 1096-1111.

[15] P.R. Kumar, T. Seidman (1990). Dynamic instabilities and stabilization methods in distributed real-time scheduling of manufacturing systems. *IEEE Transactions on Automatic Control*, 35(3), pp. 289-298.

[16] S.T. Maguluri, R. Srikant, L. Ying (2012). Stochastic models of load balancing and scheduling in cloud computing clusters. In: *Proceedings of the 31^{st} IEEE Infocom Conference*.

[17] M.G. Markakis (2013). Scheduling in switched queueing networks with heavy-tailed traffic. *PhD Thesis, Massachusetts Institute of Technology*.

[18] M.G. Markakis, E. Modiano, J.N. Tsitsiklis (2013). Max-Weight scheduling in queueing networks with heavy-tailed traffic. *IEEE/ACM Transactions on Networking*, 22(1).

[19] N. McKeown, A. Mekkittikul, V. Anantharam, J. Walrand (1999). Achieving 100% throughput in an input-queued switch. *IEEE Transactions on Communications*, 47(8), pp. 1260-1267.

[20] M.J. Neely (2008). Order optimal delay for opportunistic scheduling in multi-user wireless uplinks and downlinks. *IEEE/ACM Transactions on Networking*, 16(5), pp. 1188-1199.

[21] K. Park, W. Willinger (2000). Self-similar network traffic: an overview. In: *Self-Similar Network Traffic and Performance Evaluation*, Wiley Inc.

[22] A. Rybko, A. Stolyar (1992). Ergodicity of stochastic processes describing the operation of open queueing networks. *Problemy Peredachi Informatsii.*, 28(3), pp. 3-26.

[23] D. Shah, J.N. Tsitsiklis, Y. Zhong (2010). Qualitative properties of α-weighted scheduling policies. In: *Proceedings of the ACM Sigmetrics 2010 Conference*.

[24] D. Shah, D. Wischik (2012). Switched networks with maximum weight policies: fluid approximation and multiplicative state space collapse. *The Annals of Applied Probability*, 22(1), pp. 70-127.

[25] A. Stolyar (2004). Maxweight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic. *The Annals of Applied Probability*, 14(1), pp. 1-53.

[26] L. Tassiulas, A. Ephremides (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12), pp. 1936-1948.

[27] J.N. Tsitsiklis (1987). Analysis of a multiaccess control scheme. *IEEE Transactions on Automatic Control*, 32(11), pp. 1017-1020.

[28] V.J. Venkataraman, X. Lin, L. Ying, S. Shakkottai (2010). On scheduling for minimizing end-to-end buffer usage over multihop wireless networks. In: *Proceedings of the 29^{th} IEEE Infocom Conference*.