Max-Weight Scheduling in Queueing Networks with Heavy-Tailed Traffic

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

***Abstract*—We consider the problem of scheduling in a single-hop switched network with a mix of heavy-tailed and light-tailed traffic, and analyze the impact of heavy-tailed traffic on the performance of Max-Weight scheduling. As a performance metric we use the delay stability of traffic flows: a traffic flow is delay stable if its expected steady state delay is finite, and delay unstable otherwise. First, we show that a heavy-tailed traffic flow is delay unstable under any scheduling policy. Then, we focus on the celebrated Max-Weight scheduling policy, and show that a light-tailed flow that conflicts with a heavy-tailed flow is also delay unstable. This is true irrespective of the rate or the tail distribution of the light-tailed flow, or other scheduling constraints in the network. Surprisingly, we show that a light-tailed flow can become delay unstable, even when it does not conflict with heavy-tailed traffic. Delay stability in this case may depend on the rate of the light-tailed flow. Finally, we turn our attention to the class of Max-Weight-α scheduling policies. We show that if the α-parameters are chosen suitably, then the sum of the α-moments of the steady-state queue lengths is finite. We provide an explicit upper bound for the latter quantity, from which we derive results related to the delay stability of traffic flows, and the scaling of moments of steady state queue lengths with traffic intensity.**

I. Introduction

We consider a single-hop switched network, a queueing system where the traffic of each flow is buffered in a dedicated single-server queue, eventually gets served, and then exits the system. This model has been used to capture the dynamics and decisions in data communication networks (e.g., wireless networks [10], input-queued switches [20]), flexible manufacturing systems [8], and cloud computing facilities [16]. In all of these application areas, 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 subsets of servers can be simultaneously active, giving rise to a fundamental scheduling problem: which subset of servers to activate, and at which point in time? Clearly, the overall performance of the network depends critically on the scheduling policy applied.

The focus of this paper is on a well-studied class of scheduling policies, commonly refered to as Max-Weight policies. This class of policies was introduced in the seminal work of Tassiulas and Ephremides [32], and since then, numerous studies have analyzed the performance of such policies in very general settings, e.g., see [1], [10], and the references therein. A remarkable property of Max-Weight policies is their throughput optimality, i.e., their ability to stabilize a queueing network whenever this is possible, without any explicit information on the arriving traffic. Moreover, it has been shown that policies from this class achieve low, or even optimal, average delay for specific network topologies under light-tailed traffic, and are asymptotically delay optimal in the heavy traffic regime [9], [22], [27], [31], [33]. [1] However, the performance of Max-Weight scheduling in the presence of heavy-tailed traffic is not well-understood.

We are motivated to study networks with heavy-tailed traffic by empirical evidence that traffic in data communication networks exhibits strong correlations and statistical similarity over different time scales. This observation was first made by Leland *et al.* [15] through analysis of Ethernet traffic traces. Subsequent empirical studies have documented this phenomenon in other networks, while accompanying theoretical studies have associated it with arrival processes that have heavy tails; see [24] for an overview. Although the impact of heavy tails has been analyzed extensively in single or multi-server queues, e.g., see the survey papers [2], [4], the related work for more complex queueing systems, with a mix of heavy-tailed and light-tailed traffic, is rather limited. Notable exceptions are the papers by Borst *et al.* [3], by Jagannathan *et al.* [13], and by Nair *et al.* [21], all of which study a system with two parallel queues and a single server, with a mix of heavy-tailed and light-tailed traffic, under different scheduling policies.

The present paper aims to fill a gap in the literature, by analyzing the performance of Max-Weight scheduling in the context of a switched queueing network, with a mix of heavy-tailed and light-tailed traffic. In particular, we study the delay stability of traffic flows: a traffic flow is delay stable if its expected steady state delay is finite, and delay unstable otherwise. Relative to the existing literature, our **main contributions** are the following: (i) we show that under the Max-Weight scheduling policy, any light-tailed flow that conflicts with a heavy-tailed flow is delay unstable; (ii) surprisingly, we also show that for certain admissible arrival rates, a light-tailed flow can be delay unstable even if it does not conflict with heavy-tailed traffic; (iii) we analyze the Max-Weight-α scheduling policy, and show that if the α-parameters are chosen suitably, then the α-moments of the steady state

The authors are with the Laboratory for Information and Decision Systems, at the Massachusetts Institute of Technology, Cambridge, MA, USA. Emails: {mihalis,modiano,jnt}@mit.edu. They gratefully acknowledge financial support from NSF Grants CNS-0915988 and CCF-0728554, and ARO MURI Grant W911NF-08-1-0238

The material in this paper was presented in part at the 47^{th} Allerton Conference on Communication, Control, and Computing, Monticello, IL, in 2009, and at the 31^{st} IEEE INFOCOM Conference, Orlando, FL, in 2012.

[1] On the other hand, when Max-Weight scheduling is combined with Back-Pressure routing in the context of multi-hop networks, there is evidence that delay performance can be poor, e.g., see the discussion in [5].

queue lengths is finite. We use this result to prove that, by proper choice of the α-parameters, all light-tailed flows are delay stable. Moreover, we show that Max-Weight-α achieves the optimal scaling of higher moments of steady state queue lengths with traffic intensity.

The rest of the paper is organized as follows. Section II includes a detailed presentation of the queueing model considered in this paper, as well as formal definitions of heavy-tailed and light-tailed traffic and of delay stability. In Section III we motivate the subsequent development by presenting, informally and through simple examples, the main results of the paper. In Section IV we analyze the performance of the celebrated Max-Weight scheduling policy. Section V contains the analysis of the parameterized Max-Weight-α scheduling policy and of the performance that it achieves, in terms of delay stability. This section also includes results about the scaling of moments of steady state queue lengths with the traffic intensity and the size of the network. We conclude with a discussion of our findings and future research directions in Section VI. The appendices contain some background material and most of the proofs of our results.

II. Model and Problem Formulation

We start with a detailed presentation of the queueing model considered in this paper, 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. Central to our model is the notion of a **traffic flow** $f \in \{1, \ldots, F\}$, 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 rate of traffic flow f and by $\lambda = (\lambda_f;\ f = 1, \ldots, F)$ the vector of the rates of all traffic flows.

Definition 1: (Heavy Tails) A random variable X is heavy-tailed if $\mathbb{E}[X^2]$ is infinite, and is light-tailed otherwise. We define similarly a heavy/light-tailed IID 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 of 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 [24]).

In this paper 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 $A(t) = (A_f(t);\ f = 1, \ldots, F)$ and $Q(t) = (Q_f(t);\ f = 1, \ldots, F)$, 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 queue f. The vector $D(k) = (D_f(k);\ f = 1, \ldots, F)$ captures the end-to-end delay of the k^{th} files of the different traffic flows.

In a switched network, 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 **feasible schedule**. We denote by $\mathcal{S}$ the set of all feasible schedules, which is assumed to be an arbitrary subset of the powerset of $\{1, \ldots, F\}$. 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 $S(t) = (S_f(t);\ f = 1, \ldots, F)$.

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 $\mathbb{Z}_+^F$.

Let us now define formally the notion of a **scheduling policy**. The past history and present state of the system at time slot $t \in \mathbb{N}$ is captured by the vector

$$H(t) = (Q(0), A(0), \ldots, Q(t-1), A(t-1), Q(t)).$$

At time slot 0, we have $H(0) = (Q(0))$. A (causal) scheduling policy is a sequence $\pi = (\mu_0, \mu_1, \ldots)$ of functions $\mu_t : H(t) \to \mathcal{S},\ t \in \mathbb{Z}_+$, used to determine scheduling decisions, according to $S(t) = \mu_t(H(t))$.

We restrict our attention to scheduling policies that are **regenerative**, i.e., policies under which the network starts afresh probabilistically at certain time slots. More precisely, under a regenerative policy there exists a sequence of stopping times $\{\tau_n;\ n \in \mathbb{Z}_+\}$ with the folowing properties. (i) The sequence $\{\tau_{n+1} - \tau_n;\ n \in \mathbb{Z}_+\}$ is IID. (ii) Let $X(t) = (Q(t), A(t), S(t))$, and consider the processes that describe the "cycles" of the network, namely, $C_0 = \{X(t);\ 0 \le t < \tau_0\}$, and $C_n = \{X(\tau_{n-1}+t);\ 0 \le t < \tau_n - \tau_{n-1}\},\ n \in \mathbb{N}$; then, $\{C_n;\ n \in \mathbb{N}\}$ is an IID sequence, independent of C_0. (iii) The (lattice) distribution of the cycle lengths, $\tau_{n+1} - \tau_n$, has span equal to one and finite expectation.

Properties (i) and (ii) imply that the switched network evolves like a (possibly delayed) regenerative process. Property (iii) states that this process is aperiodic and positive recurrent. We note that Max-Weight-type policies, which are the focus of this paper, are regenerative (this will be made precise later). Moreover, a number of other widely studied policies belong to this class, e.g., priority, round-robin, and randomized policies.

The following definition gives the precise notion of stability that we use in this paper.

Definition 2: (Stability) The switched network described above is stable under a specific scheduling policy 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 $Q(0)$.

Notice that our definition of stability is slightly different than the commonly used definition (positive recurrence of the Markov chain of queue lengths), since it includes the convergence of the sequence of file delays $\{D(k);\ k \in \mathbb{N}\}$. The reason is that in this paper we study properties of the limiting distribution of $\{D(k);\ k \in \mathbb{N}\}$ and, naturally, we need to ensure that this limiting distribution exists.

Under a stabilizing scheduling policy, we denote by $Q = (Q_f;\ f = 1,\dots,F)$ and $D = (D_f;\ f = 1,\dots,F)$ generic random vectors distributed according to the limiting distributions of $\{Q(t);\ t \in \mathbb{Z}_+\}$ and $\{D(k);\ k \in \mathbb{N}\}$, respectively. The dependence of these limiting distributions on the scheduling policy has been suppressed from the notation, but will be clear from the context. 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. We note that under a regenerative policy (if one exists), the queueing network is guaranteed to be stable. This is because the sequences of queue lengths and file delays are (possibly delayed) aperiodic and positive recurrent regenerative processes, and, hence, converge in distribution; see [30].

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

Definition 3: (Stability Region [32]) The stability region of the queueing network described above, denoted by Λ, is the set of rate vectors

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

In other words, a rate vector λ belongs to Λ if there exists a convex combination of feasible schedules that covers the rates of all traffic flows. If a rate vector is in the stability region of the network, then the traffic corresponding to this vector is called **admissible**, and there exists a scheduling policy under which the network is stable.

Definition 4: (Traffic Intensity) The traffic intensity of a rate vector $\lambda \in \Lambda$ is a real number defined as follows:

$$\rho(\lambda) = \inf\Big\{\sum_{s\in\mathcal{S}} \zeta_s \;\Big|\; \lambda \le \sum_{s\in\mathcal{S}} \zeta_s \cdot s;\ \zeta_s \in \mathbb{R}_+,\ \forall s \in \mathcal{S}\Big\}.$$

Clearly, arriving traffic with rate vector λ is admissible if and only if $\rho(\lambda) < 1$. **Throughout this paper we assume that the traffic is admissible**.

Let us now define the property that we use to evaluate the performance of scheduling policies, namely, the delay stability of a traffic flow.

Definition 5: (Delay Stability) A traffic flow f is delay stable under a specific scheduling policy, if the switched network is stable under that policy and $\mathbb{E}[D_f]$ is finite; otherwise, the traffic flow f is delay unstable.

The following lemma relates the steady state quantities $\mathbb{E}[Q_f]$ and $\mathbb{E}[D_f]$, and will help us prove delay stability results.

Lemma 1: Consider the switched network described above under a regenerative scheduling policy. Then,

$$\mathbb{E}[Q_f] < \infty \iff \mathbb{E}[D_f] < \infty, \qquad \forall f \in \{1,\dots,F\}.$$

Proof: see Appendix 1.1. ■

Theorem 1: (Delay Instability of Heavy Tails) Consider the switched network described above under a regenerative scheduling policy. Every heavy-tailed traffic flow is delay unstable.

Proof: see Appendix 2. ■

Since there is little we can do about the delay stability of heavy-tailed flows, we turn our attention to light-tailed traffic. The Pollaczek-Khinchine formula for the expected delay in a $M/G/1$ queue indicates that the intrinsic burstiness of light-tailed traffic is not sufficient to cause delay instability. However, scheduling in a queueing network couples the statistics of different traffic flows. We will see that this coupling can cause light-tailed flows to become delay unstable, giving rise to a form of **propagation of delay instability**.

It should be noted that Lemma 1, Theorem 1, and all subsequent results are proved under the assumption that the "First Come, First Served" discipline is used within each queue. Indeed, a heavy-tailed flow could be delay stable under other intra-queue service disciplines, e.g., "Last Come, First Served" or Processor-Sharing; see [4]. However, the focus of this paper is on the impact of heavy-tailed traffic on light-tailed flows, under Max-Weight-type scheduling. Since Max-Weight policies are queue length-based, the main findings of this paper that characterize this impact (Theorem 2, Propositions 1 and 2, Corollary 1) remain true irrespective of the service discipline within each queue.

III. Overview of Main Results

In this section we introduce, informally and through simple examples, the main results of the paper and the basic intuition behind them.

Let us start with the queueing system of Figure 1, which consists of two parallel queues and a single server. Traffic flow 1 is assumed to be heavy-tailed, whereas traffic flow 2 is light-tailed. Service is allocated according to the Max-Weight scheduling policy, which is equivalent to "Serve the Longest Queue" in this simple setting. Theorem 1 implies that traffic flow 1 is delay unstable. Our findings imply that **traffic flow 2 is also delay unstable, even though it is light-tailed**. The intuition behind this result is that queue 1 is occasionally very long (infinite, in steady state expectation) because of its heavy-tailed arrivals. When this happens, and under the Max-Weight policy, queue 2 has to build up to a similar length in order to receive service. A very long queue then implies very large delays for the files of that queue under "First Come, First Served," which leads to delay instability.

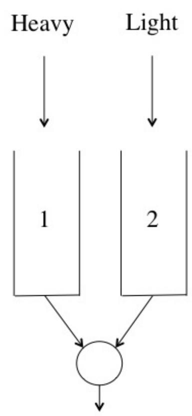

Fig. 1. Delay instability in parallel queues with heavy-tailed traffic.

Systems of parallel queues have been analyzed extensively in the literature. One of the main reasons is that their simple dynamics often lead to elegant analysis and simple results. In this paper we go beyond parallel queues, and analyze queueing networks with a more complicated structure. An example is the queueing network of Figure 2, where traffic flow 1 is assumed to be heavy-tailed, whereas traffic flows 2 and 3 are light-tailed. The server can serve either queue 1 alone, or queues 2 and 3 simultaneously. This example could represent a wireless network with interference constraints. In this setting the Max-Weight policy compares the length of queue 1 to the sum of the lengths of queues 2 and 3, and serves the "heavier" schedule.

The intuition from the previous example suggests that at least one of the queues 2 and 3 has to build up to the order of magnitude of queue 1, in order for these two queues to receive service. In other words, we expect that at least one of the traffic flows 2 and 3 will be delay unstable under Max-Weight. Our findings imply that, in fact, **both traffic flows are delay unstable**. The main idea behind this result is the following: with positive probability, the arrival processes to queues 2 and 3 exhibit their "average" behavior. In that case, the corresponding queues build up slowly and together, which implies that when they finally claim the server, they have both built up to the order of magnitude of queue 1.

The simple networks of Figures 1 and 2 illustrate special cases of a general result: every light-tailed flow that conflicts

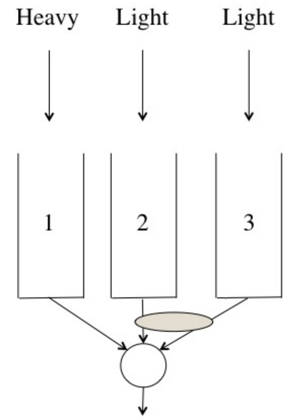

Fig. 2. Propagation of delay instability: conflicting with heavy-tailed traffic.

with a heavy-tailed flow is delay unstable. For more details, see Theorem 2 in Section IV.A.

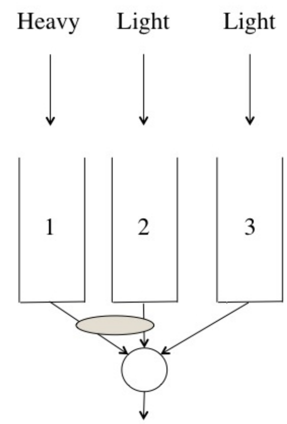

Fig. 3. Propagation of delay instability: non-conflicting with heavy-tailed traffic.

Going one step further, consider the queueing network of Figure 3. Traffic flow 1 is assumed to be heavy-tailed, whereas traffic flows 2 and 3 are light-tailed. The server can serve either queues 1 and 2 simultaneously, or queue 3 alone. 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. The intuition from the previous examples suggests that traffic flow 3 is delay unstable, but there is a nontrivial question regarding the delay stability of traffic flow 2. One would expect that this flow is delay stable: it is light-tailed itself, and is served together with a heavy-tailed flow, which should result in more service opportunities under Max-Weight. Surprisingly though, we show that **there exist arrival rates within the stability region of this network, such that traffic flow 2 is delay unstable**. The key observation here is that even though traffic flow 2 does not conflict with heavy-tailed traffic, it does conflict with traffic flow 3, which is delay unstable because it conflicts with heavy-tailed traffic. Conversely, we also show that traffic flow 2 is delay stable, if its rate is sufficiently low. For more details see Propositions 1 and 2 in Section IV.B.

The examples above suggest that in queueing networks with heavy-tailed traffic, delay instability not only appears but propagates through the network under the Max-Weight policy. Seeking a remedy to this situation, we turn to the more general Max-Weight-α scheduling policy. This policy assigns a positive α-parameter to each traffic flow, and instead of using the queue lengths to calculate the weight of a schedule, it uses the respective α-powers of the queue lengths. Our findings imply that in the network of Figure 1, we can guarantee that **traffic flow 2 is delay stable, provided the α-parameter for traffic flow 1 is sufficiently small**. In other words, we can prevent the propagation of delay instability. This is a special case of a general result: if the α-parameters of the Max-Weight-α policy are chosen suitably, then the sum of the α-moments of steady state queue lengths is finite (see Theorem 3 in Section V.A).

IV. Max-Weight Scheduling

In this section we evaluate the performance of the Max-Weight scheduling policy in terms of the delay stability of traffic flows. Informally speaking, the "weight" of a feasible schedule is the sum of the lengths of all queues included in it. As its name suggests, the Max-Weight policy activates a feasible schedule that has maximum weight, at any given time slot. More formally, under the Max-Weight policy, the scheduling vector $S(t)$ satisfies

$$S(t) \ \in \ \arg\max_{(S_f)\in\mathcal{S}} \Big\{ \sum_{f=1}^{F} Q_f(t)\cdot S_f \Big\}.$$

If the set on the right-hand side includes multiple feasible schedules, then one of them is chosen uniformly at random. The following lemma states that the network is stable under the Max-Weight policy. Essentially, this result is well-known, e.g., for light-tailed traffic, see [32]; for more general arrivals, see [31]. A subtle point is that in this paper we adopt a somewhat different definition for stability. So, we need to ensure that, apart from the sequences of queue lengths, the sequences of file delays converge as well.

Lemma 2: (Stability under Max-Weight) The switched network described in Section II is stable under the Max-Weight scheduling policy.

Proof: Consider the switched network of Section II under the Max-Weight scheduling policy. It can be verified that the sequence $\{Q(t);\ t\in\mathbb{Z}_+\}$ is a time-homogeneous, irreducible, and aperiodic Markov chain on the countable state-space $\mathbb{Z}_+^F$. Proposition 2 of [31] implies that this Markov chain is also positive recurrent. Hence, $\{Q(t);\ t\in\mathbb{Z}_+\}$ converges in distribution, and its limiting distribution does not depend on $Q(0)$. Based on this, it can be verified that the sequence $\{D(k);\ k\in\mathbb{N}\}$ is a (possibly delayed) aperiodic and positive recurrent regenerative process. Therefore, it also converges in distribution, and its limiting distribution does not depend on $Q(0)$; see [30]. ■

A. Conflicting with Heavy-Tailed Flows

Next, we state one of the main results of the paper, which generalizes our observations from the simple networks of Figures 1 and 2. Before we give the result, let us define precisely the notion of conflict between traffic flows.

Definition 6: Traffic flow f conflicts with f', and vice versa, if there exists no feasible schedule in $\mathcal{S}$ that includes both f and f'.

Theorem 2: (Conflicting with Heavy Tails) Consider the switched network described in Section II under the Max-Weight scheduling policy. Every light-tailed flow that conflicts with a heavy-tailed flow is delay unstable.

Proof: see Appendix 3. ■

We emphasize the generality of this result. Namely, a light-tailed flow that conflicts with heavy-tailed traffic is delay unstable, irrespective of (i) its arrival rate; (ii) the tail asymptotics of its arrivals; (iii) whether it is scheduled alone or with other traffic flows. Hence, we view Theorem 2 as capturing a universal phenomenon of instability propagation.

B. Non-Conflicting with Heavy-Tailed Flows

So far we have shown that (i) a heavy-tailed traffic flow is delay unstable under any regenerative scheduling policy; and (ii) a light-tailed traffic flow that conflicts with a heavy-tailed flow is delay unstable under the Max-Weight scheduling policy. It seems reasonable to assume that a light-tailed flow that does not conflict with heavy-tailed traffic is delay stable. Surprisingly, this is not always the case. We demonstrate this by means of a simple example.

Let us come back to the queueing system of Figure 3. The feasible schedules of this system are $\{1,2\}$ and $\{3\}$, and all queues are served at unit rate whenever the respective schedules are activated. The rate vector $\lambda=(\lambda_1,\lambda_2,\lambda_3)$ is assumed admissible. The following proposition shows that traffic flow 2 is delay unstable if its rate is sufficiently high.

Proposition 1: (Rate-Dependent Delay Instability) Consider the switched network of Figure 3 under the Max-Weight scheduling policy. If the arriving traffic is admissible and the rates satisfy $\lambda_2 > (1+\lambda_1-\lambda_3)/2$, then traffic flow 2 is delay unstable.

Proof: (Sketch) Our approach is based on tracking the evolution of the system on a particular set of "fluid" sample paths: assume that at time slot 0, queue 1 receives a very large file, consisting of b packets. For a long period of time after that, queue 3 does not receive service under the Max-Weight policy, and builds up. If the arrival processes of all traffic flows are close to their "average behavior," then at the time slot when the service switches from schedule $\{1,2\}$ to schedule $\{3\}$, the lengths of both queues 1 and 3 are proportional to b, whereas queue 2 is still small. From that point on, the Max-Weight policy will drain the weights of the two schedules at roughly the same rate, until one of the weights becomes zero.

Let μ_f be the average departure rate from queue f during the latter period. For the weights of the two schedules to be drained at the same rate, the departure rates have to satisfy:

$$\lambda_1 + \lambda_2 - \mu_1 - \mu_2 = \lambda_3 - \mu_3.$$

Moreover, the fact that Max-Weight is a work-conserving policy implies that

$$\mu_1 + \mu_3 = 1.$$

Finally, since queues 1 and 2 are served simultaneously, and queue 2 may be empty during parts of the draining period, we have that

$$\mu_1 \geq \mu_2.$$

The above equations and some simple algebra imply that

$$\frac{1 + \lambda_1 + \lambda_2 - \lambda_3}{3} \geq \mu_2.$$

Now suppose that the arrival rates satisfy

$$\lambda_2 > \frac{1 + \lambda_1 - \lambda_3}{2}.$$

Then,

$$\lambda_2 > \frac{1 + \lambda_1 + \lambda_2 - \lambda_3}{3} \geq \mu_2.$$

This implies that queue 2 builds up at a roughly constant rate during a period of time whose duration is proportional to b. Thus, queue 2 eventually builds up to size $O(b)$, and the integral of $Q_2(t)$ over a busy period of the process becomes of order $O(b^2)$. In that case $\mathbb{E}[Q_2]$ is infinite, because b is drawn from a heavy-tailed distribution (see Lemma 3 in Appendix 1.2). Finally, Lemma 1 implies the delay instability of queue 2. A detailed proof of Proposition 1 can be found in [18]. ■

Now we establish that when $\lambda_2 < (1 + \lambda_1 - \lambda_3)/2$ queue 2 is delay stable, achieving, thus, an exact characterization of the "delay stability region" of queue 2. In order to do that, we further assume that light-tailed traffic has exponentially decaying tails. Formally, a nonnegative random variable X is **exponential-type**, if there exists some $\theta > 0$ such that $\mathbb{E}[\exp(\theta X)]$ is finite.

Proposition 2: (Rate-Dependent Delay Stability) Consider the switched network of Figure 3 under the Max-Weight scheduling policy, with admissible arriving traffic. Suppose that $A_2(0)$ and $A_3(0)$ are exponential-type, and that there exists some $\gamma > 0$ such that $\mathbb{E}[A_1^{\gamma+1}(0)]$ is finite. If the arrival rates satisfy $\lambda_2 < (1+\lambda_1-\lambda_3)/2$, then queue 2 is delay stable, and the steady state length of queue 2 is exponential-type.

Proof: (Sketch) The proof of Proposition 2 is based on drift analysis of the following piecewise linear Lyapunov function, which is nonincreasing in the length of the heavy-tailed queue, and has a negative drift only when $\lambda_2 < (1 + \lambda_1 - \lambda_3)/2$:

$$G(t) = 3Q_2(t) + \Big[Q_3(t) - Q_1(t) - Q_2(t)\Big]^+,$$

where $[x]^+$ stands for $\max\{x, 0\}$, the nonnegative part of a scalar x. In particular, we show that for sufficiently large (but fixed) $T \in \mathbb{N}$, there exist positive constants α and ϵ, such that

$$\mathbb{E}\Big[G(t+T) - G(t) + \epsilon \ ; \ G(t) > \alpha \ \Big| \ \mathcal{F}_t\Big] \leq 0,$$

where $\mathcal{F}_t$ is the σ-algebra generated by $Q(0), A(0), \ldots, Q(t-1), A(t-1), Q(t)$, and $\mathbb{E}[X; A \mid \mathcal{H}]$ is a shorthand notation for $\mathbb{E}[X \cdot 1_A \mid \mathcal{H}]$.

The above drift condition, and the fact that the arrivals to queues 2 and 3 are exponential-type, imply that the (steady state) random variable

$$3Q_2 + [Q_3 - Q_1 - Q_2]^+$$

is exponential-type (see Theorem 2.3 in [11]). A detailed proof of Proposition 2 can be found in [18].

■

V. Max-Weight-α Scheduling

The results of the previous section suggest that Max-Weight scheduling performs poorly in the presence of heavy-tailed traffic. The reason is that by treating heavy-tailed and light-tailed flows equally, there are very long stretches of time during which heavy-tailed traffic dominates the service. This leads some light-tailed flows to experience very large delays. Intuitively, by discriminating against heavy-tailed flows, one should be able to improve the overall performance of the network, namely to mitigate the propagation of delay instability. One way to do this is by giving preemptive priority to the light-tailed flows. However, priority-based scheduling policies are undesirable because of fairness considerations, and also because they can be unstable in many network settings, e.g., see [14], [25].

Instead, we focus on the Max-Weight-α scheduling policy: given constants $\alpha_f > 0, \ f \in \{1, \ldots, F\}$, the scheduling vector $S(t)$ satisfies

$$S(t) \ \in \ \arg\max_{(S_f)\in\mathcal{S}} \Big\{ \sum_{f=1}^{F} Q_f^{\alpha_f}(t) \cdot S_f \Big\}.$$

If the set on the right-hand side includes multiple feasible schedules, one of them is chosen uniformly at random. By choosing smaller values for the α-parameters of heavy-tailed flows and larger values for light-tailed flows, we give a form of partial priority to light-tailed traffic.

A. The Main Result

Let us start with a preview of the **main result** of this section: if the α-parameters of the Max-Weight-α policy are chosen so that $\mathbb{E}[A_f^{\alpha_f+1}(0)]$ is finite, for all $f \in \{1, \ldots, F\}$, then the network is stable and the steady state queue lengths satisfy:

$$\mathbb{E}[Q_f^{\alpha_f}] < \infty, \qquad \forall f \in \{1, \ldots, F\}.$$

An earlier work by Eryilmaz *et al.* has given a similar result for the case of parallel queues with a single server; see Theorem 1 of [6]. In this paper we extend their result to a single-hop switched network setting. Moreover, we provide an explicit upper bound on the sum of the α-moments of the steady state queue lengths. Before we do that we need the following definition.

Definition 7: (Covering Number of Feasible Schedules) The covering number k^* of the set of feasible schedules is defined as the smallest number k for which there exist $s^1, \ldots, s^k \in S$ with $\bigcup_{i=1}^{k} s^i = \{1, \ldots, F\}$.

Theorem 3: (Max-Weight-α Scheduling) Consider the switched network described in Section II under the Max-Weight-α scheduling policy. Let the intensity of the arriving traffic be $\rho < 1$. If $\mathbb{E}[A_f^{\alpha_f+1}(0)]$ is finite, for all $f \in \{1, \ldots, F\}$, then the network is stable and the steady state queue lengths satisfy

$$\sum_{f=1}^{F} \mathbb{E}[Q_f^{\alpha_f}] \leq \sum_{f=1}^{F} H\Big(\rho, k^*, \alpha_f, \mathbb{E}[A_f^{\alpha_f+1}(0)]\Big),$$

where

$$H\Big(\rho, k^*, \alpha_f, \mathbb{E}[A_f^{\alpha_f+1}(0)]\Big) = \begin{cases} \frac{2k^*}{1-\rho} \cdot \Big(\mathbb{E}[A_f^{\alpha_f+1}(0)] + 1\Big), & \alpha_f \leq 1, \\ \Big(\frac{2k^*}{1-\rho}\Big)^{\alpha_f} \cdot K^{\alpha_f} + \frac{2k^*}{1-\rho} \cdot K, & \alpha_f > 1, \end{cases}$$

and $K = 2^{\alpha_f - 1} \cdot \alpha_f \cdot \Big(\mathbb{E}[A_f^{\alpha_f+1}(0)] + 1\Big)$.

Proof: see Appendix 4. ∎

It is known that bounds derived from conventional Lyapunov arguments are, in general, loose. The bound provided in Theorem 3 is, probably, no exception to this rule, e.g., see Corollary 3 and the subsequent discussion, in Section V.C. In this light, the value of Theorem 3 lies on the following: (i) it gives a feel for which structural parameters of the network and which characteristics of the arriving traffic may affect the actual performance of Max-Weight-type policies; (ii) it provides the correct scaling of higher order queue length moments with traffic intensity (see Corollary 2, in Section V.B).

B. Traffic Burstiness and Delay Stability

A first corollary of Theorem 3 relates to the delay stability of light-tailed flows.

Corollary 1: (Delay Stability under Max-Weight-α) Consider the switched network described in Section II under the Max-Weight-α scheduling policy. Suppose that there exists some $\gamma > 0$ such that $\mathbb{E}[A_f^{\gamma+1}(0)]$ is finite, for all $f \in \{1, \ldots, F\}$. If the α-parameters of all light-tailed flows are equal to 1, and the α-parameters of heavy-tailed flows are sufficiently small, then all light-tailed flows are delay stable.

Proof: With the particular choice of α-parameters, Theorem 3 guarantees that the expected steady state queue length of all light-tailed flows is finite. Lemma 1 relates this result to delay stability. ∎

Combining this with Theorem 1, we conclude that, when the α-parameters are chosen suitably, **the Max-Weight-α policy delay-stabilizes a traffic flow whenever this is possible**.

Max-Weight-α turns out to perform well in terms of another criterion too. Theorem 3 implies that by choosing the α-parameters so that $\mathbb{E}[A_f^{\alpha_f+1}(0)]$ is finite, for all $f \in \{1, \ldots, F\}$, the steady state queue length moment $\mathbb{E}[Q_f^{\alpha_f}]$ is finite, for all $f \in \{1, \ldots, F\}$. The following result suggests that, for traffic flows with polynomially decaying tails, this is the best we can do under any regenerative scheduling policy.

Theorem 4: Consider the switched network described in Section II under a regenerative scheduling policy. If, for any given $f \in \{1, \ldots, F\}$ and $\gamma > 0$, the moment $\mathbb{E}[A_f^{\gamma+1}(0)]$ is infinite, then $\mathbb{E}[Q_f^{\gamma}]$ is infinite.

Proof: This result is well-known in the context of a $M/G/1$ queue, e.g., see Section 3.2 of [4]. It can be proved similarly to Theorem 1. ∎

Thus, when the α-parameters are chosen suitably, the Max-Weight-α policy guarantees the finiteness of the highest possible moments for traffic flows with polynomially decaying tails.

C. Scaling Results under Light-Tailed Traffic

Although this paper focuses on heavy-tailed traffic and its consequences, some implications of Theorem 3 are of general interest. In this section we assume that all traffic flows in the network are light-tailed, and analyze how the sum of the α-moments of steady state queue lengths scales with traffic intensity and the size of the network.

Corollary 2: (Scaling with Traffic Intensity) Let us fix a switched network and constants $\alpha \geq 1$ and $B > 0$. The Max-Weight-α scheduling policy is applied with $\alpha_f = \alpha$, for all $f \in \{1, \ldots, F\}$. Assume that the traffic arriving to the network is admissible, and that $\mathbb{E}[A_f^{\alpha+1}(0)] \leq B$, for all f. Then,

$$\sum_{f=1}^{F} \mathbb{E}[Q_f^{\alpha}] \leq \frac{M(k^*, \alpha, B)}{(1-\rho)^{\alpha}},$$

where $M(k^*, \alpha, B)$ is a constant that depends only on k^*, α, and B. Moreover, under any stabilizing scheduling policy,

$$\sum_{f=1}^{F} \mathbb{E}[Q_f^{\alpha}] \geq \frac{M'(\alpha)}{(1-\rho)^{\alpha}},$$

where $M'(\alpha)$ is a constant that depends only on α.

Proof: The first part of the result follows directly from Theorem 3. Regarding the second part, Theorem 2.1 of [28] implies that under any stabilizing scheduling policy there exists an absolute constant $\tilde{M}$, such that

$$\sum_{f=1}^{F} \mathbb{E}[Q_f] \geq \frac{\tilde{M}}{1-\rho}.$$

Using Jensen's inequality, we have

$$\sum_{f=1}^{F} \mathbb{E}[Q_f^{\alpha}] \geq \sum_{f=1}^{F} (\mathbb{E}[Q_f])^{\alpha} \geq \frac{1}{F^{\alpha}} \Big(\sum_{f=1}^{F} \mathbb{E}[Q_f]\Big)^{\alpha}.$$

Consequently, there exists a constant $M'(\alpha)$ that depends only on α, such that

$$\sum_{f=1}^{F} \mathbb{E}[Q_f^{\alpha}] \geq \frac{M'(\alpha)}{(1-\rho)^{\alpha}},$$

under any stabilizing scheduling policy. ∎

Similar scaling results appear in queueing theory, mostly in the context of single-server queues, e.g., see Chapter 3 of [12]. More recently, results of this flavor have been shown for particular queueing networks, such as input-queued switches [26], [28]. All related work, though, concerns the scaling of first moments. Corollary 2 gives the precise scaling of higher order steady state queue length moments with traffic intensity, and shows that Max-Weight-α achieves the **optimal scaling of the α-moments**.

Finally, we turn our attention to the performance of the Max-Weight scheduling policy under Bernoulli traffic, i.e., when each of the arrival processes $\{A_f(t);\ t \in \mathbb{Z}_+\}$ is an independent Bernoulli process with parameter $\lambda_f > 0$. We denote by $S_{\max}$ the maximum number of traffic flows that any feasible schedule $s \in \mathcal{S}$ can serve.

The following bound characterizes the performance of Max-Weight in terms of structural parameters of the network and the traffic intensity.

Corollary 3: (Scaling under Bernoulli Traffic) Consider the switched network described in Section II under the Max-Weight scheduling policy. Assume that the traffic arriving to the network is Bernoulli, with traffic intensity $\rho < 1$. Then,

$$\sum_{f=1}^{F} \mathbb{E}[Q_f] \leq 2k^* S_{\max}\Big(\frac{1+\rho}{1-\rho}\Big).$$

Proof: If all traffic flows are light-tailed and all α-parameters are equal to one, a more careful accounting in the proof of Theorem 3 provides the following tighter upper bound:

$$\sum_{f=1}^{F} \mathbb{E}[Q_f] \leq \frac{2k^*}{1-\rho}\Big(S_{\max} + \sum_{f=1}^{F} \mathbb{E}[A_f^2(0)]\Big).$$

If the traffic arriving to the network is Bernoulli, then $\mathbb{E}[A_f^2(0)] = \lambda_f$, for all $f \in \{1, \dots, F\}$. Moreover, the fact that the arriving traffic has intensity ρ, implies the existence of nonnegative real numbers ζ_s, for $s \in \mathcal{S}$, such that

$$\lambda_f \leq \sum_{s \in S} \zeta_s \cdot s_f, \qquad \forall f \in \{1, \dots, F\},$$

with $\sum_{f=1}^{F} \zeta_s = \rho$. Consequently,

$$\begin{aligned}\sum_{f=1}^{F} \mathbb{E}[A_f^2(0)] = \sum_{f=1}^{F} \lambda_f \leq \sum_{f=1}^{F}\sum_{s \in S} \zeta_s \cdot s_f &\leq \sum_{s \in S} \zeta_s \cdot S_{\max} \\ &= \rho \cdot S_{\max},\end{aligned}$$

and the result follows. ∎

Example 1: (n **Parallel Queues**) Consider a single-server system with n parallel queues. The arriving traffic is assumed to be Bernoulli, with traffic intensity $\rho < 1$. In this case $k^* = n$ and $S_{\max} = 1$. Corollary 3 implies that under the Max-Weight scheduling policy, the sum of the steady state queue lengths is bounded from above as follows:

$$\sum_{i=1}^{n} \mathbb{E}[Q_i] \leq \frac{4n}{1-\rho}.$$

The aggregate queue length of a system of parallel queues under a work-conserving scheduling policy evolves like a discrete time M/G/1 queue, from which we infer that $\sum_{i=1}^{n} \mathbb{E}[Q_i] = \Theta\Big(\frac{1}{1-\rho}\Big)$. So, in the context of parallel queues, the scaling provided by Corollary 3 is tight with respect to the traffic intensity, but not necessarily tight with respect to the size of the network.

Example 2: (**$n \times n$ Input-Queued Switch**) Consider a $n \times n$ input-queued switch. The arriving traffic is assumed to be Bernoulli, with traffic intensity $\rho < 1$. In this case $k^* = n$ and $S_{\max} = n$. Corollary 3 implies that under the Max-Weight scheduling policy, the sum of the steady state queue lengths is bounded from above as follows:

$$\sum_{i=1}^{n}\sum_{j=1}^{n} \mathbb{E}[Q_{ij}] \leq \frac{4n^2}{1-\rho}.$$

In the context of an input-queued switch, the joint scaling provided by Corollary 3, in terms of both the traffic intensity and the size of the network, is the tightest currently known. However, it should be noted that the correct scaling under Max-Weight as ρ goes to one, and as n becomes large, is an open problem [26]. On a related note, a different scheduling policy has been recently shown to achieve the optimal joint scaling [29].

VI. Discussion

The main conclusion of this paper is that the Max-Weight scheduling policy performs poorly in the presence of heavy-tailed traffic. More specifically, we show that the phenomenon of delay instability not only arises but can propagate to a significant part of the network, possibly depending on the arrival rates. However, from the sketches of the proofs of Propositions 1 and 2, it becomes obvious that analyzing rate-dependent delay (in)stability phenomena is somewhat involved, even in simple queueing systems like the one in Figure 3. The analysis of more complex networks is the subject of ongoing research, and relies on fluid approximations, renewal theory, and stochastic Lyapunov theory [19].

Another important conclusion is that the Max-Weight-α scheduling policy can be used to alleviate the effects of heavy-tailed traffic, and is order optimal if its α-parameters are chosen suitably. However, for Max-Weight-α to perform well, some knowledge of higher order statistics of the traffic flows is required. If the α-parameters are not chosen appropriately, then in light of Theorem 4, this policy may also perform poorly.

An important direction for future research is to consider queueing networks with correlated traffic, as opposed to the IID arrivals that are assumed in this paper. As alluded to earlier, evidence suggests that traffic in data communication networks exhibits strong correlations in time, e.g., long-range dependence. We believe that the shortcomings of Max-Weight persist in the presence of correlated traffic. In particular, queues that receive LRD traffic are likely to dominate the service for long periods of time, leading to large delays at conflicting queues. For a related work in a more abstract context, see also [23].

Appendix 1 - Background Material

1.1 BASTA, Little's Law, and Delay Stability

In this section we state the steady state versions of two important results in queueing theory, the "Bernoulli Arrivals See Time Averages" property and Little's Law, which we subsequently use to prove Lemma 1.

Consider the switched network described in Section II. Let $\tau_{f,k}$ be the random time slot of the arrival of the k^{th} file to queue f, $k \in \mathbb{N}$, $f \in \{1, \ldots, F\}$. We assign two marks to this file: (i) the vector of queue lengths upon its arrival $Q^{c(f)}(k) = (Q_g(\tau_{f,k});\ g = 1, \ldots, F)$; and (ii) its end-to-end delay $D_f(k)$.

Under a regenerative scheduling policy, and for a given $f \in \{1, \ldots, F\}$, the vector-valued sequences $\{Q^{c(f)}(k);\ k \in \mathbb{N}\}$ and $\{Q(t);\ t \in \mathbb{Z}_+\}$ are (possibly delayed) aperiodic and positive recurrent regenerative processes. Therefore, they converge in distribution, and their limiting distributions do not depend on $Q(0)$; see [30]. We denote by $Q^{c(f)} = (Q_g^{c(f)};\ g = 1, \ldots, F)$ and $Q = (Q_g;\ g = 1, \ldots, F)$ generic random vectors distributed according to these limiting distributions.

The arrival of files at queue f constitutes a Bernoulli process with parameter $p_f = \mathbb{P}(A_f(0) > 0)$, since all arrival processes are IID. The Bernoulli Arrivals See Time Averages (BASTA) property relates the limiting distributions $Q^{c(f)}$ and Q.

Theorem 5: (BASTA) Consider the switched network described in Section II under a regenerative scheduling policy. The random vectors Q and $Q^{c(f)}$ are identically distributed, for all $f \in \{1, \ldots, F\}$.

Now let $L_f(t)$ be the number of files in queue f at time slot t, either in queue or in service. Under a regenerative scheduling policy, the sequences $\{L_f(t);\ t \in \mathbb{Z}_+\}$ and $\{D_f(k);\ k \in \mathbb{N}\}$ are (possibly delayed) aperiodic and positive recurrent regenerative processes. Hence, they converge in distribution, and their limiting distributions do not depend on $Q(0)$; see [30]. We denote by L_f and D_f generic random variables distributed according to these limiting distributions. Little's Law relates their expected values.

Theorem 6: (Little's Law) Consider the switched network described in Section II under a regenerative scheduling policy. Then,

$$\mathbb{E}[L_f] = p_f \cdot \mathbb{E}[D_f], \qquad \forall f \in \{1, \ldots, F\}.$$

This holds even if the above expectations are infinite.

Theorems 5 and 6 can be proved by combining the rather elementary time-average versions of BASTA and Little's Law (which can be found in [17] and [35], respectively), with well-known ergodicity properties of regenerative stochastic systems.

We now use these results to prove **Lemma 1**. Let us start with the implication

$$\mathbb{E}[Q_f] < \infty \implies \mathbb{E}[D_f] < \infty, \qquad \forall f \in \{1, \ldots, F\}.$$

Fix a traffic flow $f \in \{1, \ldots, F\}$, and assume that $\mathbb{E}[Q_f]$ is finite. Since every file has at least one packet, then for all $t \in \mathbb{Z}_+$, and for all $b \in \mathbb{Z}_+$,

$$\mathbb{P}(Q_f(t) > b) \geq \mathbb{P}(L_f(t) > b).$$

We have argued that under a regenerative scheduling policy, the sequences $\{Q_f(t);\ t \in \mathbb{Z}_+\}$ and $\{L_f(t);\ t \in \mathbb{Z}_+\}$ converge in distribution. So, taking the limit as t goes to infinity, we have

$$\mathbb{P}(Q_f > b) \geq \mathbb{P}(L_f > b), \qquad \forall b \in \mathbb{Z}_+,$$

which implies that

$$\mathbb{E}[Q_f] \geq \mathbb{E}[L_f].$$

Combining this inequality with Little's Law and the assumption that $\mathbb{E}[Q_f]$ is finite, we conclude that $\mathbb{E}[D_f]$ is finite.

Let us now prove the implication

$$\mathbb{E}[Q_f] = \infty \implies \mathbb{E}[D_f] = \infty, \qquad \forall f \in \{1, \ldots, F\}.$$

Fix a traffic flow $f \in \{1, \ldots, F\}$, and assume that $\mathbb{E}[Q_f]$ is infinite. The end-to-end delay of a file is bounded from below by the length of the respective queue upon its arrival, since the service discipline within each queue is "First Come, First Served." So, for all $k \in \mathbb{N}$, and for all $b \in \mathbb{Z}_+$,

$$\mathbb{P}(D_f(k) > b) \geq \mathbb{P}(Q_f(\tau_{f,k}) > b).$$

We have argued that under a regenerative scheduling policy, the sequences $\{D_f(k);\ k \in \mathbb{N}\}$ and $\{Q_f(\tau_{f,k});\ k \in \mathbb{N}\}$ converge in distribution. So, taking the limit as k goes to infinity, we have

$$\mathbb{P}(D_f > b) \geq \mathbb{P}(Q_f^{c(f)} > b), \qquad \forall b \in \mathbb{Z}_+.$$

Combining this with the BASTA property, we get

$$\mathbb{P}(D_f > b) \geq \mathbb{P}(Q_f > b), \qquad \forall b \in \mathbb{Z}_+,$$

which implies that

$$\mathbb{E}[D_f] \geq \mathbb{E}[Q_f].$$

Finally, the assumption that $\mathbb{E}[Q_f]$ is infinite implies that $\mathbb{E}[D_f]$ is infinite as well.

1.2 Truncated Rewards

Consider the switched network described in Section II under a regenerative scheduling policy. By definition, there exists a sequence of stopping times $\{\tau_n;\ n \in \mathbb{Z}_+\}$, which constitutes a (possibly delayed) renewal process, i.e., the sequence $\{\tau_{n+1} - \tau_n;\ n \in \mathbb{Z}_+\}$ is IID. Moreover, the lattice distribution of cycle lengths has span equal to one and finite expectation.

For $t \in \mathbb{Z}_+$, let $R(t)$ be an instantaneous reward, which is assumed to be an arbitrary scalar-valued function of $Q(t)$. We define the truncated reward as $R^M(t) = \min\{R(t), M\}$, where M is a positive integer. Under a regenerative scheduling policy, the sequences $\{R(t);\ t \in \mathbb{Z}_+\}$ and $\{R^M(t);\ t \in \mathbb{Z}_+\}$ are (possibly delayed) aperiodic and positive recurrent regenerative processes. Consequently, they converge in distribution, and their limiting distributions do not depend on $Q(0)$; see [30]. Let R and R^M be generic random variables distributed according to these limiting distributions. We denote by R_{agg} the aggregate reward, i.e., the reward accumulated over a regeneration cycle. Similarly, R_{agg}^M represents the aggregate truncated reward.

Lemma 3: Consider the switched network described in Section II under a regenerative scheduling policy. Suppose that there exists a random variable Y with infinite expectation, and a nondecreasing function $f(\cdot)$, such that

$$\lim_{M\to\infty} f(M) = \infty,$$

and

$$\mathbb{E}[\min\{Y, f(M)\}] \le \mathbb{E}[R^M_{agg}]. \tag{1}$$

Then, $\mathbb{E}[R]$ is infinite.

Proof: By definition, regeneration cycle lengths have finite expectation, and $\mathbb{E}[R^M_{agg}]$ is bounded from above by $M \cdot \mathbb{E}[\tau_1 - \tau_0]$. Then, the Renewal Reward theorem implies that

$$\frac{\mathbb{E}[R^M_{agg}]}{\mathbb{E}[\tau_1 - \tau_0]} = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} R^M(t), \tag{2}$$

almost surely; see Section 3.4 of [7]. The sequence $\{R^M(t);\ t \in \mathbb{Z}_+\}$ is a (possibly delayed) positive recurrent regenerative process, which is also uniformly bounded by M. Then, the Ergodic theorem for regenerative processes implies that

$$\begin{aligned}\lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} R^M(t) &= \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \min\{R(t), M\} \\ &= \mathbb{E}[\min\{R, M\}], \end{aligned} \tag{3}$$

almost surely; see [30]. Eqs. (1)-(3) give

$$\frac{\mathbb{E}[\min\{Y, f(M)\}]}{\mathbb{E}[\tau_1 - \tau_0]} \le \mathbb{E}[\min\{R, M\}].$$

By taking the limit as M goes to infinity on both sides, and using the Monotone Convergence theorem, we obtain

$$\frac{\mathbb{E}[Y]}{\mathbb{E}[\tau_1 - \tau_0]} \le \mathbb{E}[R];$$

see Section 5.3 of [34]. Finally, the fact that $\mathbb{E}[Y]$ is infinite implies that $\mathbb{E}[R]$ is infinite as well. ■

1.3 "Average Behavior" of IID Processes

The following result is a well-known corollary of the Strong Law of Large Numbers. We provide a proof for completeness.

Lemma 4: Consider a sequence of IID random variables $\{B(\tau);\ \tau \in \mathbb{N}\}$, taking values in $\mathbb{Z}_+$, with finite rate $\lambda = \mathbb{E}[B(1)] > 0$. For any given $\epsilon > 0$, there exists a constant $\delta > 0$, such that

$$\mathbb{P}\Big(\Big\{(\lambda-\epsilon)t - \delta \le \sum_{\tau=1}^{t} B(\tau) \le (\lambda+\epsilon)t + \delta\Big\},\ \forall t \in \mathbb{N}\Big) > 0.$$

Proof: We define an event C_m by

$$C_m = \Big\{ \Big|\frac{1}{t}\sum_{\tau=1}^{t} B(\tau) - \lambda\Big| \le \epsilon,\quad \forall\ t \ge m\Big\}.$$

By the Strong Law of Large Numbers,

$$\mathbb{P}(\cup_{m\ge1} C_m) = 1.$$

Because the sequence of events C_m is nondecreasing, the continuity property of probabilities implies that

$$\lim_{m\to\infty} \mathbb{P}(C_m) = 1.$$

So, let us therefore fix some T, such that $\mathbb{P}(C_T) > 1/2$. Now consider the event

$$D = \Big\{ 0 \le \sum_{\tau=1}^{T} B(\tau) \le \delta \Big\}.$$

We choose δ large enough so that $\mathbb{P}(D) > 1/2$ and $\delta \ge \lambda T$. Notice that

$$\mathbb{P}(C_T \cap D) \ge \mathbb{P}(C_T) + \mathbb{P}(D) - 1 > \frac{1}{2} + \frac{1}{2} - 1 = 0.$$

Finally, note that when both C_T and D occur, then

$$\Big|\sum_{\tau=1}^{t} B(\tau) - \lambda t\Big| \le \epsilon t + \delta, \qquad \forall\ t \in \mathbb{N},$$

so that the latter event has positive probability, which is the desired result. ■

Appendix 2 - Proof of Theorem 1

Consider a heavy-tailed traffic flow $h \in \{1, \ldots, F\}$. We will show that $\mathbb{E}[Q_h]$ is infinite under any regenerative scheduling policy. Combined with Lemma 1, this will imply that traffic flow h is delay unstable.

Consider a fictitious queue, denoted by $\tilde{h}$, which has exactly the same arrivals and initial length as queue h, but is served at unit rate whenever nonempty. We denote by $Q_{\tilde{h}}(t)$ the length of queue $\tilde{h}$ at time slot t. Since the arriving traffic is assumed admissible, the queue length process $\{Q_{\tilde{h}}(t);\ t \in \mathbb{Z}_+\}$ converges to a limiting distribution $Q_{\tilde{h}}$.

An easy inductive argument can show that the length of queue h dominates the length of queue $\tilde{h}$ at all time slots, under any regenerative scheduling policy. This implies that for all $t \in \mathbb{Z}_+$, and for all $b \in \mathbb{Z}_+$,

$$\mathbb{P}(Q_h(t) > b) \ge \mathbb{P}(Q_{\tilde{h}}(t) > b).$$

Taking the limit as t goes to infinity, and using the fact that both queue length processes converge in distribution, we have

$$\mathbb{P}(Q_h > b) \ge \mathbb{P}(Q_{\tilde{h}} > b), \qquad \forall b \in \mathbb{Z}_+.$$

So, in order to prove the desired result, it suffices to show that $\mathbb{E}[Q_{\tilde{h}}]$ is infinite.

Notice that the length of queue $\tilde{h}$ evolves as a positive recurrent Markov chain, and the empty state is recurrent. Hence, the time slots that initiate busy periods of queue $\tilde{h}$ constitute a (possibly delayed) renewal process. We define an instantaneous reward on this renewal process:

$$R^M(t) = \min\{Q_{\tilde{h}}(t), M\}, \qquad \forall t \in \mathbb{Z}_+,$$

where M is some finite integer.

Without loss of generality, assume that a busy period starts at time slot 0, and let b be the size of the file that initiates it. Since queue $\tilde{h}$ is served at unit rate, its length is at least $b/2$ packets over a time period of length at least $b/2$ time slots.

This implies that the aggregate reward R^M_{agg}, i.e., the reward accumulated over a renewal period, is bounded from below as follows:

$$R^M_{agg} \geq \frac{b}{2} \cdot \min\left\{\frac{b}{2}, M\right\} \geq \min\left\{\frac{b^2}{4}, M^2\right\}.$$

Consequently, the expected aggregate reward is bounded from below as follows:

$$\begin{aligned}\mathbb{E}[R^M_{agg}] &\geq \sum_{b=0}^{\infty} \min\left\{\frac{b^2}{4}, M^2\right\} \cdot \mathbb{P}(A_h(0) = b) \\ &= \mathbb{E}\Big[\min\left\{\frac{A_h^2(0)}{4}, M^2\right\}\Big].\end{aligned}$$

Then, Lemma 3 (see Appendix 1.2) applied to $Y = (1/4)A_h^2(0)$, implies that $\mathbb{E}[Q_{\tilde{h}}]$ is infinite. This, in turn, implies that $\mathbb{E}[Q_h]$ is infinite, which, combined with Lemma 1, gives the desired result.

Appendix 3 - Proof of Theorem 2

Consider a heavy-tailed traffic flow h, and a light-tailed flow l that conflicts with h. We will show that, under the Max-Weight scheduling policy, $\mathbb{E}[Q_l]$ is infinite. Combined with Lemma 1, this will imply that traffic flow l is delay unstable.

Notice that the vector of queue lengths evolves as a positive recurrent Markov chain, and the empty state is recurrent. Hence, the time slots that initiate busy periods of the system constitute a (possibly delayed) renewal process. We define an instantaneous reward on this renewal process:

$$R^M(t) = \min\{Q_l(t), M\}, \qquad \forall t \in \mathbb{Z}_+,$$

where M is a positive integer.

Without loss of generality, assume that a busy period of the network starts at time slot 0. Consider the set of sample paths where at time slot 0, queue h receives a file of size b packets, and all other queues receive no traffic; we denote this set of sample paths by $H(b)$. Since the arrival processes of different traffic flows are mutually independent,

$$\mathbb{P}(H(b)) = \mathbb{P}(A_h(0) = b) \cdot \prod_{g \neq h} \mathbb{P}(A_g(0) = 0).$$

This quantity is positive as long as b is in the support of $A_h(0)$, because the rate vector is admissible (hence $\lambda_g < 1$) and $\mathbb{P}(A_g(0) = 0) \geq 1 - \lambda_g > 0$. For sample paths in $H(b)$, denote by T_b the first time slot when the length of queue h becomes less than or equal to the sum of the lengths of all other queues:

$$T_b = \min\Big\{t > 0 \;\Big|\; \sum_{g \neq h} Q_g(t) \geq Q_h(t)\Big\} \cdot 1_{H(b)}.$$

Under the Max-Weight scheduling policy, queue l receives no service until time slot T_b. Moreover, queue h is served at unit rate. So, for sample paths in $H(b)$,

$$b - (T_b - 1) \leq Q_h(T_b) \leq \sum_{g \neq h} Q_g(T_b) = \sum_{g \neq h} \sum_{t=1}^{T_b - 1} A_g(t).$$

Lemma 4 implies that for every $\epsilon > 0$ there exists $\delta > 0$, such that the set of sample paths

$$\Delta = \Big\{\Big|\sum_{\tau=1}^{t} A_g(\tau) - \lambda_g\Big| \leq \epsilon \cdot t + \delta, \ \forall t \in \mathbb{N}, \ \forall g \neq h\Big\},$$

has positive probability (see Appendix 1.3 for a proof of this result.) We denote by $\tilde{H}(b)$ the set of sample paths $\Delta \cap H(b)$. Due to the IID nature of the arriving traffic,

$$\mathbb{P}(\tilde{H}(b)) = \mathbb{P}(\Delta) \cdot \mathbb{P}(H(b)).$$

For sample paths in $\tilde{H}(b)$, we have

$$T_b - 1 \geq \frac{b - (F-1) \cdot \delta}{\sum_{g \neq h}(\lambda_g + \epsilon) + 1}.$$

Moreover,

$$Q_l(T_b) = \sum_{t=1}^{T_b - 1} A_l(t) \geq (\lambda_l - \epsilon) \cdot (T_b - 1) - \delta.$$

Consequently, there exist positive constants c and b_0 such that for every $b \geq b_0$, and any sample path in $\tilde{H}(b)$, we have

$$Q_l(T_b) \geq cb.$$

Since at most one packet from queue l can be served at each time slot, the length of queue l is at least $cb/2$ over a time period of length at least $cb/2$ time slots. This implies that the aggregate reward R^M_{agg}, i.e., the reward accumulated over a renewal period, satisfies the lower bound

$$R^M_{agg} \cdot 1_{\{b \geq b_0\}} \cdot 1_{\tilde{H}(b)} \geq \min\left\{\left(\frac{cb}{2}\right)^2 \cdot 1_{\{b \geq b_0\}}, M^2\right\} \cdot 1_{\tilde{H}(b)}.$$

Then, the expected aggregate reward satisfies

$$\begin{aligned}\mathbb{E}[R^M_{agg}] &\geq \mathbb{P}(\Delta) \cdot \prod_{g \neq h} \mathbb{P}(A_g(0) = 0) \\ &\quad \cdot \sum_{b=1}^{\infty} \min\left\{\left(\frac{cb}{2}\right)^2 \cdot 1_{\{b \geq b_0\}}, M^2\right\} \cdot \mathbb{P}(A_h(0) = B).\end{aligned}$$

So, there exists a positive constant c', such that

$$\mathbb{E}[R^M_{agg}] \geq c' \cdot \mathbb{E}\Big[\min\left\{\left(\frac{cA_h(0)}{2}\right)^2 \cdot 1_{\{A_h(0) \geq b_0\}}, M^2\right\}\Big].$$

Lemma 3 (see Appendix 1.2) applied to $Y = (1/4)c^2 A_h^2(0) \cdot 1_{\{A_h(0) \geq b_0\}}$, implies that $\mathbb{E}[Q_l]$ is infinite. Then, Lemma 1 gives the desired result.

Appendix 4 - Proof of Theorem 3

The admissibility of the arriving traffic implies that we can find a set of feasible schedules $\{\sigma^k;\ k = 1, \ldots, k^*\}$ that satisfies

$$\bigcup_{k=1}^{k^*} \sigma^k = \{1, \ldots, F\}.$$

By the definition of the intensity parameter $\rho \in (0, 1)$, there exist nonnegative numbers $\zeta_i,\ i = 1, \ldots, I$, adding up to 1, and feasible schedules $\tilde{s}^i,\ i=1,\dots,I$, such that

$$\lambda \leq \rho \cdot \sum_{i=1}^{I} \zeta_i \cdot \tilde{s}^i.$$

Notice that

$$\Big((1-\rho)\cdot \sum_{k=1}^{k^*} \frac{1}{k^*}\cdot \sigma^k + \rho \cdot \sum_{i=1}^{I} \zeta_i \cdot \tilde{s}^i\Big) \in \overline{\Lambda},$$

where $\overline{\Lambda}$ denotes the closure of the set Λ. This is because we have a convex combination of $(I+k^*)$ feasible schedules, and the stability region is known to be a convex set; see Section 3.2 of [10]. Moreover,

$$(1-\rho)\cdot \sum_{k=1}^{k^*} \frac{1}{k^*}\cdot \sigma^k = \frac{1-\rho}{k^*}\cdot \sum_{k=1}^{k^*} \sigma^k \geq \frac{1-\rho}{k^*}\cdot 1_F,$$

where 1_F denotes the F-dimensional vector of ones.

A well-known monotonicity property of the stability region is the following: if $0 \leq \lambda' \leq \lambda''$ componentwise, and $\lambda'' \in \Lambda$, then $\lambda' \in \Lambda$. Using this property, we have that

$$\Big(\frac{1-\rho}{k^*}\cdot 1_F + \lambda\Big) \in \overline{\Lambda}.$$

This, in turn, implies the existence of nonnegative numbers $\theta_j,\ j=1,\cdots,J$, adding up to 1, and of feasible schedules $s^j=(s^j_f),\ j=1,\cdots,J$, such that

$$\lambda_f \leq \sum_{j=1}^{J} \theta_j \cdot s^j_f - \frac{1-\rho}{k^*}, \qquad \forall f \in \{1,\dots,F\}. \tag{4}$$

Under the Max-Weight-α scheduling policy, the sequence $\{Q(t);\ t\in\mathbb{Z}_+\}$ is a time-homogeneous, irreducible, and aperiodic Markov chain on the countable state-space $\mathbb{Z}_+^F$. In order to establish positive recurrence, we use the convex Lyapunov function

$$V(t) = \sum_{f=1}^{F} \frac{1}{\alpha_f+1} Q_f^{\alpha_f+1}(t).$$

We have

$$\begin{aligned}&\mathbb{E}[V(t+1) \mid \mathcal{F}_t] \\ &\qquad = \sum_{f=1}^{F} \mathbb{E}\Big[\frac{1}{\alpha_f+1}(Q_f(t)+\Delta_f(t))^{\alpha_f+1}\ \Big|\ \mathcal{F}_t\Big],\end{aligned}$$

where

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

and $\mathcal{F}_t$ is the σ-algebra generated by $Q(0), A(0), \dots, Q(t-1), A(t-1), Q(t)$. Throughout the proof we use the shorthand notation

$$V_f(t) = \frac{1}{\alpha_f+1} Q_f^{\alpha_f+1}(t).$$

We consider the conditional expectation of the terms $V_f(t+1)$, distinguishing between two cases.

(i) $\alpha_f \leq 1$: Consider the zeroth order Taylor expansion around $Q_f(t)$ (i.e., the mean value theorem):

$$\begin{aligned}&\frac{1}{\alpha_f+1}\Big(Q_f(t)+\Delta_f(t)\Big)^{\alpha_f+1} \\ &\qquad = \frac{1}{\alpha_f+1} Q_f^{\alpha_f+1}(t) + \Delta_f(t)\cdot \xi(t)^{\alpha_f},\end{aligned}$$

for some $\xi(t) \in \Big[Q_f(t) - S_f(t)\cdot 1_{\{Q_f(t)>0\}}, Q_f(t)+A_f(t)\Big]$. Thus,

$$V_f(t+1) = V_f(t) + \Delta_f(t)\cdot \xi(t)^{\alpha_f},$$

and

$$\mathbb{E}[V_f(t+1) \mid \mathcal{F}_t] = V_f(t) + \mathbb{E}[\Delta_f(t)\cdot \xi(t)^{\alpha_f} \mid \mathcal{F}_t].$$

Consider the event $\Gamma_f(t) = \{\Delta_f(t) < 0\}$ and its complement. We have

$$\begin{aligned}&\mathbb{E}[V_f(t+1) \mid \mathcal{F}_t] \\ &\leq V_f(t) + \mathbb{E}[\Delta_f(t)\cdot (Q_f(t)+A_f(t))^{\alpha_f}\cdot 1_{\{\Gamma_f^c(t)\}} \mid \mathcal{F}_t] \\ &+ \mathbb{E}[\Delta_f(t)\cdot (Q_f(t) - S_f(t)\cdot 1_{\{Q_f(t)>0\}})^{\alpha_f}\cdot 1_{\{\Gamma_f(t)\}} \mid \mathcal{F}_t].\end{aligned} \tag{5}$$

Since $Q_f(t),\ Q_f(t) - S_f(t)\cdot 1_{\{Q_f(t)>0\}}$, and $A_f(t)$ are nonnegative numbers and $\alpha_f \in (0,1]$, it can be verified that

$$\Big(Q_f(t)+A_f(t)\Big)^{\alpha_f} \leq Q_f^{\alpha_f}(t) + A_f^{\alpha_f}(t). \tag{6}$$

Moreover, because they are also integers, it can be verified that

$$\Big(Q_f(t) - S_f(t)\cdot 1_{\{Q_f(t)>0\}}\Big)^{\alpha_f} \geq Q_f^{\alpha_f}(t) - S_f(t)\cdot 1_{\{Q_f(t)>0\}}. \tag{7}$$

Eqs. (5)-(7) imply that

$$\begin{aligned}&\mathbb{E}[V_f(t+1) \mid \mathcal{F}_t] \\ &\quad \leq V_f(t) + \mathbb{E}[\Delta_f(t) \mid \mathcal{F}_t]\cdot Q_f^{\alpha_f}(t) \\ &\qquad + \mathbb{E}[\Delta_f(t)\cdot A_f^{\alpha_f}(t)\cdot 1_{\{\Gamma_f^c(t)\}} \mid \mathcal{F}_t] \\ &\qquad + \mathbb{E}[-\Delta_f(t)\cdot S_f(t)\cdot 1_{\{Q_f(t)>0\}}\cdot 1_{\{\Gamma_f(t)\}} \mid \mathcal{F}_t].\end{aligned}$$

If $\Delta_f(t) < 0$, which is the event $\Gamma_f(t)$, then $-\Delta_f(t) \leq 1$. Also, if $\Delta_f(t) \geq 0$, which is the event $\Gamma_f^c(t)$, then $\Delta_f(t) \leq A_f(t)$, so that $\Delta_f(t)\cdot A_f^{\alpha_f}(t) \leq A_f^{\alpha_f+1}(t)$. Consequently,

$$\begin{aligned}\mathbb{E}[V_f(t+1) \mid \mathcal{F}_t] \leq & V_f(t) + \mathbb{E}[\Delta_f(t) \mid \mathcal{F}_t]\cdot Q_f^{\alpha_f}(t) \\ & + \mathbb{E}[A_f^{\alpha_f+1}(t)\cdot 1_{\{\Gamma_f^c(t)\}} \mid \mathcal{F}_t] + 1.\end{aligned}$$

Finally, the fact that the random variables $\{A_f(t);\ t\in\mathbb{Z}_+\}$ are IID gives

$$\begin{aligned}&\mathbb{E}[V_f(t+1) \mid \mathcal{F}_t] \\ &\quad \leq V_f(t) + \mathbb{E}[\Delta_f(t) \mid \mathcal{F}_t]\cdot Q_f^{\alpha_f}(t) + \mathbb{E}[A_f^{\alpha_f+1}(0)] + 1.\end{aligned}$$

The inequality above implies trivially that

$$\begin{aligned}\mathbb{E}[V_f(t+1) \mid \mathcal{F}_t] \leq & V_f(t) + \mathbb{E}[\Delta_f(t) \mid \mathcal{F}_t]\cdot Q_f^{\alpha_f}(t) \\ & + \frac{1-\rho}{2k^*}\cdot Q_f^{\alpha_f}(t) + \mathbb{E}[A_f^{\alpha_f+1}(0)] + 1.\end{aligned} \tag{8}$$

(ii) $\alpha_f > 1$: Consider the first order Taylor expansion around $Q_f(t)$:

$$\begin{aligned}&\frac{1}{\alpha_f+1}\Big(Q_f(t)+\Delta_f(t)\Big)^{\alpha_f+1}\\=&\frac{1}{\alpha_f+1}Q_f(t)^{\alpha_f+1}+\Delta_f(t)\cdot Q_f^{\alpha_f}(t)+\frac{\Delta_f^2(t)}{2}\cdot\alpha_f\cdot\xi(t)^{\alpha_f-1}\end{aligned}$$

for some $\xi(t)\in\Big[Q_f(t)-S_f(t)\cdot 1_{\{Q_f(t)>0\}},Q_f(t)+A_f(t)\Big]$. Then,

$$\begin{aligned}\mathbb{E}[V_f(t+1)\mid\mathcal{F}_t]=&V_f(t)+\mathbb{E}[\Delta_f(t)\mid\mathcal{F}_t]\cdot Q_f^{\alpha_f}(t)\\&+\mathbb{E}\Big[\frac{\Delta_f^2(t)}{2}\cdot\alpha_f\cdot\xi(t)^{\alpha_f-1}\ \Big|\ \mathcal{F}_t\Big].\end{aligned}\tag{9}$$

Since $\Delta_f^2(t)\cdot\alpha_f\geq 0$ and $\alpha_f-1\geq 0$, the last term can be bounded from above as follows:

$$\begin{aligned}&\mathbb{E}\Big[\frac{\Delta_f^2(t)}{2}\cdot\alpha_f\cdot\xi(t)^{\alpha_f-1}\ \Big|\ \mathcal{F}_t\Big]\\&\qquad\leq\mathbb{E}\Big[\frac{\Delta_f^2(t)}{2}\cdot\alpha_f\cdot(Q_f(t)+A_f(t))^{\alpha_f-1}\ \Big|\ \mathcal{F}_t\Big].\end{aligned}\tag{10}$$

Moreover, it is easy to verify that for $\alpha_f\geq 1$,

$$\Big(Q_f(t)+A_f(t)\Big)^{\alpha_f-1}\leq 2^{\alpha_f-1}\cdot\Big(Q_f^{\alpha_f-1}(t)+A_f^{\alpha_f-1}(t)\Big),\tag{11}$$

and also that

$$\Delta_f^2(t)\leq A_f^2(t)+1.\tag{12}$$

Eqs. (10)-(12) imply that

$$\begin{aligned}&\mathbb{E}\Big[\frac{\Delta_f^2(t)}{2}\cdot\alpha_f\cdot\xi^{\alpha_f-1}\ \Big|\ \mathcal{F}_t\Big]\\&\quad\leq 2^{\alpha_f-2}\cdot\alpha_f\cdot\Big(\mathbb{E}[A_f^2(t)]+1\Big)\cdot Q_f^{\alpha_f-1}(t)\\&\qquad+2^{\alpha_f-2}\cdot\alpha_f\cdot\Big(\mathbb{E}[A_f^{\alpha_f+1}(t)]+\mathbb{E}[A_f^{\alpha_f-1}(t)]\Big)\\&\quad\leq K\cdot Q_f^{\alpha_f-1}(t)+K,\end{aligned}\tag{13}$$

where $K=2^{\alpha_f-1}\cdot\alpha_f\cdot\Big(\mathbb{E}[A_f^{\alpha_f+1}(0)]+1\Big)$. Then, Eqs. (9) and (13) imply that

$$\begin{aligned}&\mathbb{E}[V_f(t+1)\mid\mathcal{F}_t]\\&\leq V_f(t)+\mathbb{E}[\Delta_f(t)\mid\mathcal{F}_t]\cdot Q_f^{\alpha_f}(t)+K\cdot Q_f^{\alpha_f-1}(t)+K\\&=V_f(t)+\mathbb{E}[\Delta_f(t)\mid\mathcal{F}_t]\cdot Q_f^{\alpha_f}(t)+\frac{1-\rho}{2k^*}\cdot Q_f^{\alpha_f}(t)\\&\quad+\Big(K\cdot Q_f^{\alpha_f-1}(t)-\frac{1-\rho}{2k^*}\cdot Q_f^{\alpha_f}(t)+K\Big).\end{aligned}\tag{14}$$

Our goal is to bound from above the last term on the right-hand side of Eq. (14). Relaxing the constraint that $Q_f(t)$ has to be an integer, we have

$$\begin{aligned}&K\cdot Q_f^{\alpha_f-1}(t)-\frac{1-\rho}{2k^*}\cdot Q_f^{\alpha_f}(t)+K\\&\leq\max_{x\in\mathbb{R}_+}\Big\{K\cdot x^{\alpha_f-1}-\frac{1-\rho}{2k^*}\cdot x^{\alpha_f}+K\Big\},\qquad\forall t\in\mathbb{Z}_+.\end{aligned}\tag{15}$$

It can be verified that the optimization problem on the right-hand side has the unique solution $x^*=\frac{2k^*K}{1-\rho}\cdot\frac{\alpha_f-1}{\alpha_f}$. The corresponding optimal value is

$$\begin{aligned}K^{\alpha_f}\cdot\Big(\frac{2k^*}{1-\rho}\Big)^{\alpha_f-1}\cdot\frac{(\alpha_f-1)^{\alpha_f-1}}{\alpha_f^{\alpha_f}}+K&\\\leq K^{\alpha_f}\cdot\Big(\frac{2k^*}{1-\rho}\Big)^{\alpha_f-1}+K.&\end{aligned}\tag{16}$$

Eqs. (15) and (16) imply that for all $t\in\mathbb{Z}_+$,

$$K\cdot Q_f^{\alpha_f-1}(t)-\frac{1-\rho}{2k^*}\cdot Q_f^{\alpha_f}(t)+K\leq K^{\alpha_f}\cdot\Big(\frac{2k^*}{1-\rho}\Big)^{\alpha_f-1}+K.\tag{17}$$

Finally, Eqs. (14) and (17) give

$$\begin{aligned}\mathbb{E}[V_f(t+1)\mid\mathcal{F}_t]=&V_f(t)+\mathbb{E}[\Delta_f(t)\mid\mathcal{F}_t]\cdot Q_f^{\alpha_f}(t)\\&+\frac{1-\rho}{2k^*}\cdot Q_f^{\alpha_f}(t)\\&+K^{\alpha_f}\cdot\Big(\frac{2k^*}{1-\rho}\Big)^{\alpha_f-1}+K.\end{aligned}\tag{18}$$

Summarizing our findings from cases (i) and (ii), Eqs. (8) and (18) imply that

$$\begin{aligned}\mathbb{E}[V_f(t+1)\mid\mathcal{F}_t]\leq&V_f(t)+\mathbb{E}[\Delta_f(t)\mid\mathcal{F}_t]\cdot Q_f^{\alpha_f}(t)\\&+\frac{1-\rho}{2k^*}\cdot Q_f^{\alpha_f}(t)\\&+H\Big(\rho,k^*,\alpha_f,\mathbb{E}[A_f^{\alpha_f+1}(0)]\Big),\end{aligned}$$

for all $f\in\{1,\ldots,F\}$, where

$$\begin{aligned}&H\Big(\rho,k^*,\alpha_f,\mathbb{E}[A_f^{\alpha_f+1}(0)]\Big)\\&\qquad=\begin{cases}\frac{2k^*}{1-\rho}\cdot\Big(\mathbb{E}[A_f^{\alpha_f+1}(0)]+1\Big), & \alpha_f\leq 1,\\ \Big(\frac{2k^*}{1-\rho}\Big)^{\alpha_f}\cdot K^{\alpha_f}+\frac{2k^*}{1-\rho}\cdot K, & \alpha_f>1,\end{cases}\end{aligned}$$

and $K=2^{\alpha_f-1}\cdot\alpha_f\cdot\Big(\mathbb{E}[A_f^{\alpha_f+1}(0)]+1\Big)$. Summing over all $f\in\{1,\ldots,F\}$, gives

$$\begin{aligned}&\mathbb{E}[V(t+1)\mid\mathcal{F}_t]\\&\quad\leq V(t)+\sum_{f=1}^{F}(\lambda_f-S_f(t)\cdot 1_{\{Q_f(t)>0\}})\cdot Q_f^{\alpha_f}(t)\\&\quad+\frac{1-\rho}{2k^*}\cdot\sum_{f=1}^{F}Q_f^{\alpha_f}(t)+\sum_{f=1}^{F}H\Big(\rho,k^*,\alpha_f,\mathbb{E}[A_f^{\alpha_f+1}(0)]\Big).\end{aligned}$$

Taking into account Eq. (4), we have

$$\begin{aligned}\mathbb{E}[V(t+1)\mid\mathcal{F}_t]\leq&V(t)-\frac{1-\rho}{2k^*}\cdot\sum_{f=1}^{F}Q_f^{\alpha_f}(t)\\&+\sum_{f=1}^{F}H\Big(\rho,k^*,\alpha_f,\mathbb{E}[A_f^{\alpha_f+1}(0)]\Big)\\&+\sum_{f=1}^{F}\Big(\sum_{j=1}^{J}\theta_j\cdot s_f^j-S_f(t)\Big)\cdot Q_f^{\alpha_f}(t).\end{aligned}$$

By definition of the Max-Weight-α policy, the last term is nonpositive. So,

$$\mathbb{E}[V(t+1) - V(t) \mid \mathcal{F}_t] \leq -\frac{1-\rho}{2k^*} \cdot \sum_{f=1}^{F} Q_f^{\alpha_f}(t) + \sum_{f=1}^{F} H\Big(\rho, k^*, \alpha_f, \mathbb{E}[A_f^{\alpha_f+1}(0)]\Big).$$

Then, the Foster-Lyapunov stability criterion and moment bound (e.g., see Corollary 2.1.5 of [12]) implies that the sequence $\{Q(t);\ t \in \mathbb{Z}_+\}$ converges in distribution. Moreover, its limiting distribution $(Q_f;\ f = 1, \ldots, F)$ does not depend on $Q(0)$, and satisfies

$$\sum_{f=1}^{F} \mathbb{E}[Q_f^{\alpha_f}] \leq \frac{2k^*}{1-\rho} \cdot \sum_{f=1}^{F} H\Big(\rho, k^*, \alpha_f, \mathbb{E}[A_f^{\alpha_f+1}(0)]\Big).$$

Based on this, it can be verified that the sequence $\{D(k);\ k \in \mathbb{N}\}$ is a (possibly delayed) aperiodic and positive recurrent regenerative process. Hence, it also converges in distribution, and its limiting distribution does not depend on $Q(0)$; see [30].

References

[1] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar, P. Whiting (2004). Scheduling in a queueing system with asynchronously varying service rates. Probability in the Engineering and Informational Sciences, 18, 191-217.

[2] S. Borst, O. Boxma, R. Nunez-Queija, B. Zwart (2003). The impact of the service discipline on delay asymptotics. Performance Evaluation, 54, 175-206.

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

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

[5] L. Bui, R. Srikant, A. Stolyar (2009). Novel architectures and algorithms for delay reduction in back-pressure scheduling and routing. In: Proc. Infocom 2009.

[6] A. Eryilmaz, R. Srikant, J. Perkins (2005). Stable scheduling policies for fading wireless channels. IEEE/ACM Transactions on Networking, 13, 411-424.

[7] R. Gallager (1996). Discrete stochastic processes. Kluwer Academic.

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

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

[10] L. Georgiadis, M. Neely, L. Tassiulas (2006). Resource allocation and cross-layer control in wireless nertworks. Foundations and Trends in Networking, 1, 1-144.

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

[12] B. Hajek (2006). Notes on communication network analysis. Available online at: http://www.ifp.illinois.edu/∼hajek/Papers/networkanalysis Dec06.pdf.

[13] 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, 1096-1111.

[14] 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, 289-298.

[15] W. Leland, M. Taqqu, W. Willinger, D. Wilson (1994). On the self-similar nature of ethernet traffic. IEEE/ACM Transactions on Networking, 2, 1-15.

[16] S.T. Maguluri, R. Srikant, L. Ying (2012). Stochastic models of load balancing and scheduling in cloud computing clusters. In: Proc. Infocom 2012.

[17] A. Makowski, B. Melamed, W. Whitt (1989). On averages seen by arrivals in discrete time. In: Proc. CDC 1989.

[18] M.G. Markakis, E. Modiano, J.N. Tsitsiklis (2012). Delay stability regions of the Max-Weight policy under heavy-tailed traffic. arXiv e-prints, July 2012. Available at http://arxiv.org/pdf/1207.5746v1.pdf

[19] M.G. Markakis, E. Modiano, J.N. Tsitsiklis. Dynamic scheduling in queueing systems with heavy-tailed traffic: fluid approximations and the "single big event" principle; in preparation.

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

[21] J. Nair, K. Jagannathan, A. Wierman (2013). When heavy-tailed and light-tailed flows compete: the response time tail under generalized Max-Weight scheduling. In: Proc. Infocom 2013.

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

[23] B. Oguz, V. Anantharam (2012). Hurst index of functions of long-range-dependent Markov chains. Journal of Applied Probability, 49, 451-471.

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

[25] A. Rybko, A. Stolyar (1992). Ergodicity of stochastic processes describing the operation of open queueing networks. Probl. Peredachi Inf., 3, 3-26.

[26] D. Shah, J. N. Tsitsiklis, Y. Zhong (2011). Optimal scaling of average queue sizes in an input-queued switch: an open problem. Queueing Systems, 68, 375-384.

[27] D. Shah, D. Wischik (2006). Optimal scheduling algorithms for input-queued switches. In: Proc. Infocom 2006.

[28] D. Shah, D. Wischik (2008). Lower bound and optimality in switched networks. In: Proc. Allerton 2008.

[29] D. Shah, N. Walton, Y. Zhong (2012). Optimal queue-size scaling in switched networks. In: Proc. Sigmetrics 2012.

[30] K. Sigman, R. Wolff (1993). A review of regenerative processes. SIAM Review, 35, 269-288.

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

[32] 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, 1936-1948.

[33] L. Tassiulas, A. Ephremides (1993). Dynamic server allocation to parallel queues with randomly varying connectivity. IEEE Transactions on Information Theory, 39, 466-478.

[34] D. Williams (1991). Probability with Martingales. Cambridge University Press.

[35] W. Whitt (1991). A review fo $L = \lambda W$ and extensions. Queueing Systems, 9, 235-268.