On the Convergence to Fairness in Overloaded FIFO Systems

Florin Ciucu, Oliver Hohlfeld
T-Labs / TU Berlin
{florin,oliver}@net.t-labs.tu-berlin.de

Lydia Y. Chen
IBM Zurich Research Laboratory
yic@zurich.ibm.com

***Abstract*—Many of today's computing and communication systems are based on FIFO queues whose performance, e.g., in terms of throughput and fairness, is highly impacted by load fluctuations, especially in the case of short-term overload. This paper analytically proves that overloaded FIFO queues are fair in the sense that each flow or aggregate of flows receives a proportional fair share of the service rate. The convergence rate is evaluated with respect to flow sizes and intensity of overutilization for two broad and distinctive arrival classes: Markovian and heavy-tailed/self-similar. For the former class the paper shows smaller convergence times at higher utilizations, which is exactly the opposite behavior characteristic to underloaded queueing systems.**

I. Introduction

FIFO queues is a commonly used system design choice in various aspects of today's computing and communication systems, e.g., multi-core CPUs, computing clouds, or network routers and switches. Such systems are often subject to short-term fluctuation of their workloads which leads to overload and consequently to performance degradation. Examples include web servers experiencing severe transient overload [33], or network links getting unavoidably congested. To optimize the run-time/real-time performance of FIFO queue based systems, it is thus important to understand the transient behavior of an overloaded system with respect to its workload.

The workload typically consists of multiple users/clients/applications/flows competing for resources (e.g., CPU clock, threads, or bandwidth), and is modelled in terms of arrival processes. In modern systems, these arrival processes can be roughly categorized into two classes: (1) light-tailed Markovian, here represented by the class of Exponentially Bounded Burstiness (EBB) processes [34], and (2) heavy-tailed and self-similar [23]. The key characteristic of EBB processes is that deviations from the mean decay in probability exponentially fast. In turn, heavy-tailed and self-similar processes fundamentally differ from EBB processes in that deviations from the mean increase in time and decay in probability as a power law, i.e., more slower than the exponential. While the EBB class is used to model voice/video traffic [30], the class of heavy-tailed and self-similar processes is particularly relevant because it fits measurements of aggregate Internet traffic [22], [7], web server, and modern CPU execution times [15].

FIFO queues with EBB arrival processes are well understood mainly in terms of steady-state performance metrics and in underloaded regimes; relevant tools include the classical queueing theory [21] and the stochastic network calculus [3]. In turn, transient metrics are analyzed using various ordinary differential equations based numerical methods [2] which do however not scale in the size of the system, having thus limited applicability. Transient metrics but in overloaded regimes are studied in [16], [14] under the processor sharing scheduling, and in [18] for a multi-class queueing system with one class experiencing transient surge.

Compared to the analysis of the EBB class, the analysis of heavy-tailed or self-similar arrivals is fundamentally more difficult and has produced limited results mainly in large-buffer and many-sources asymptotic regimes [27]. Since such asymptotic results may be inaccurate in finite regimes [1], approximative non-asymptotic results were derived using numerical inversion transforms [28], or fitting heavy-tailed distributions with sums of exponentials [31]. The first non-asymptotic results on end-to-end delays in networks with heavy-tailed and self-similar traffic were recently obtained in [23]. To the best of our knowledge, there has been no analytical transient analysis for this class of traffic in overloaded regimes.

In this paper we provide a transient analysis of a FIFO queue in overloaded regimes for both EBB and heavy-tailed and self-similar processes. We follow the approach of the stochastic network calculus [3], [17], which is a probabilistic extension of the deterministic network calculus conceived by Cruz [8]. An attractive feature of the stochastic network calculus is that it can account for a broad class of arrivals, including the notoriously difficult heavy-tailed and self-similar processes [23]. The central analytical tool used in this paper is a probabilistic extension of the FIFO service curve [9], with a suitably chosen time parameter, which models probabilistic lower bounds on the service at FIFO schedulers, and in particular becomes appropriate for analyzing overloaded queueing systems. In this way we can show how each flow converges to the fair share of the service rate, with respect to arrivals' class, flow size, and intensity of overload.

The main contribution consists in the derivation of the convergence rates to fairness, i.e., the amount of time it takes for a flow, or an aggregate of flows, to attain the fair share in a FIFO queue when multiplexed with other flows and resulting in overload. Concretely, for the two considered arrival classes, we provide closed-form probabilistic lower bounds on the departure processes at an overloaded queue. These results determine the rates of convergence with respect to factors such as flow sizes and overutilization factor, i.e., by how much the queue is overloaded. An interesting behavior is that in the case of the EBB class the convergence is faster when increasing the utilization factor (e.g., from 1.01 to 1.25), which is exactly the opposite behavior characteristic to underloaded regimes (e.g., it takes longer for output rates to converge to input rates by increasing the utilization, e.g., from 0.75 to 0.99).

The rest of this paper is structured as follows. In Section II we formally state the problem. In Section III we analyze the case of EBB arrivals. The extension to heavy-tailed and self-similar arrivals is then considered in Section IV. Several numerical results are illustrated in Section V, and brief conclusions are given in Section VI.

II. Problem Statement

Consider the scenario from Figure 1 depicting a FIFO queue serving two arrival (or aggregate of) flows $A(t)$ and $A'(t)$, with corresponding departure flows $D(t)$ and $D'(t)$, respectively. The service rate of the scheduler is $R > 0$, and the arrival flows $A(t)$ and $A'(t)$ have the (long-term) average rates r and r', respectively. Denote the utilization factor $u = \frac{r+r'}{R}$. It was conjectured in [24] that, depending on utilization, the departure flow $D(t)$ has the average rate r_{out} given by

$$r_{\text{out}} = \begin{cases} r & , \quad \text{if } u < 1 \\ \frac{r}{r+r'}R & , \quad \text{otherwise (i.e., overload)} \end{cases} .$$

In other words, overloaded FIFO schedulers offer a proportional fair share to each of the arrival flows.

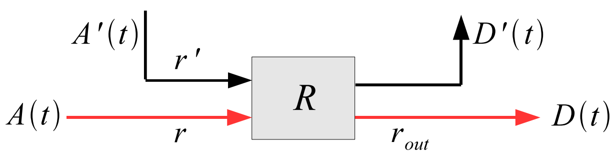

Fig. 1. A scheduler with rate R serving two arrival (or aggregate of) flows $A(t)$ and $A'(t)$ with the (long-term) average rates r and r', respectively; the departure flow $D(t)$ has the average rate r_{out}.

In this paper we focus on the overload regime and prove that if $u > 1$, i.e., the overload condition $r + r' > R$, then the departure flow $D(t)$ satisfies for all $t, \sigma \geq 0$

$$Pr\left(D(t) < \left(\frac{r}{r+r'} - \delta\right) Rt - \sigma\right) \leq \varepsilon(\sigma) \ , \qquad (1)$$

for two broad classes of arrival processes: EBB, and heavy-tailed and self-similar[1]. Here, $\delta > 0$ is a correction factor and $\varepsilon(\sigma)$ is an error function, non-increasing in σ and δ, and satisfying $\lim_{\sigma\to\infty} \varepsilon(\sigma) = 0$. The limit $\delta \to 0$ can be taken at the expense of a slower convergence of the error function. When Eq. (1) holds we say that $A(t)$ gets a probabilistic fair share from the scheduler. In this paper, we evaluate Eq. (1) by providing/deriving the specific error functions corresponding to arrival processes in the next two sections. Note that as flows here are broadly defined, i.e., single or aggregate of flows, the analysis derived in this paper are applicable to any number of flows.

[1]A proof in the special case of constant-rate arrivals, i.e., $A(t) = rt$ and $A'(t) = r't$, appeared recently in [12].

III. Scenario 1: EBB Arrivals

We start out this section by introducing the model preliminary used in the paper. We then review the class of EBB arrivals, in particular the MGF upper and lower envelope models. When $A(t)$ and $A'(t)$ from Figure 1 belong to the EBB class we demonstrate that overloaded FIFO schedulers guarantee a probabilistic fair share. To this end, we consider scenarios depending on (1) the statistical independence of $A(t)$ and $A'(t)$, and also (2) the independence of increments of $A(t)$ and $A'(t)$.

The model considered is as follows. The time model is continuous. The arrivals and departures at the FIFO scheduler from Figure 1 are given by non-decreasing and left-continuous processes, defined on some joint probability space. The scheduler is workconserving, has an infinite sized buffer, and serves the arrivals in a fluid manner, i.e., the data service units of $A(t)$ and $A'(t)$ are infinitesimally small. This fluid representation of service can be extended to account for packetization, i.e., the departure processes change once whole packets are processed rather than once infinitesimally parts of packets are processed. For the purpose of this paper we prefer a fluid service representation because it simplifies notation and the loss of generality is minimal, i.e., roughly in the order of one packet size. We also assume the existence of (long-term) average rates for the arrival processes, e.g., $r = \lim_{t\to\infty} \frac{A(t)}{t}$.

A. MGF Upper and Lower Envelopes

The class of EBB arrival processes can be described in terms of bounds on their Moment Generating Functions (MGFs) [3], [10]. Let us first introduce for convenience the bivariate process $A(s,t) = A(t) - A(s)$ for all $s \leq t$.

An arrival process $A(t)$ has an *MGF upper envelope* with rate r_u depending on some $\theta > 0$, if for all $s \leq t$

$$E\left[e^{\theta A(s,t)}\right] \leq e^{\theta r_u (t-s)} \ . \qquad (2)$$

A distinguishing feature of EBB arrivals is that the rate r_u is invariant to time parameters; this means that the characterization from Eq. (2) excludes fractional Brownian motion and also heavy-tailed arrivals which have infinite MGFs.

To demonstrate the existence of probabilistic lower bounds on a departure process as in Eq. (1), not only we need probabilistic upper bounds on the arrivals as in Eq. (2), but we also need probabilistic lower bounds on the arrivals. More precisely, we additionally need an arrival model which enforces lower bounds on the average arrival rates such that sufficient conditions for the scheduler to be in an overloaded regime can be given.

We enforce lower bounds on the average arrival rates by assuming that the arrival processes have exponential bounds on their Laplace transforms. Formally, we say that an arrival process $A(t)$ has an *MGF lower envelope* with rate r_l depending on some $\theta > 0$, if for all $s \leq t$

$$E\left[e^{-\theta A(s,t)}\right] \leq e^{-\theta r_l (t-s)} \ . \tag{3}$$

Similar to the upper envelope, the rate r_l in the lower envelope is invariant to time parameters. Slightly more general bounds can be considered in both Eqs. (2) and (3), e.g., $Me^{\theta r_u (t-s)}$ in (2), where $M \geq 1$ is the prefactor of the exponential; such EBB characterization is equivalent to the original EBB description from [34]. To keep notation simple, we consider the models from Eqs. (2) and (3).

Given an arrival process $A(t)$ with MGF upper and lower envelopes with rates r_u and r_l, respectively, for some $\theta > 0$, we have from Jensen inequality for all $t \geq 0$

$$r_l \leq \frac{E\left[A(t)\right]}{t} \leq r_u \ , \tag{4}$$

which implies that the long-term average rate of the process $r = \lim_{t\to\infty} \frac{A(t)}{t}$ satisfies $r_l \leq r \leq r_u$. Therefore, the models from Eqs. (2) and (3) can be used in conjunction to enforce overloaded regimes for the scheduler, and also prevent trivial situations of zero or infinite arrivals.

B. Examples: Compound Poisson and Markov-Modulated On-Off (MMOO) Processes

We now briefly show how to construct MGF upper and lower envelopes as in Eqs. (2) and (3). Consider first the compound Poisson process, having independent increments,

$$A(t) = \sum_{i=1}^{N(t)} X_i \ , \tag{5}$$

where $N(t)$ is a Poisson process with rate λ, and X_i's are independent exponential random variables with mean $1/\mu$. Using a conditioning argument, one can immediately derive the MGF upper and lower envelopes of $A(t)$ with the rates

$$r_u = \frac{\lambda}{\mu - \theta}, \quad r_l = \frac{\lambda}{\mu + \theta} \ ,$$

for all $\theta > 0$, with an additional constraint $\theta < \mu$ for the upper envelope.

Consider now the MMOO process which is characterized by burstiness and lack of independent increments. Let a continuous homogeneous Markov chain $S(u)$ with two states 'On' and 'Off', and transition rates μ and λ between the 'On' and 'Off' states, and vice-versa, respectively. If a source produces at some constant rate $P > 0$ while the chain is in the 'On' state, then the MMOO process is

$$A(t) = \int_0^t P I_{\{S(u)=\text{`}On'\}} du \ , \tag{6}$$

where $I_{\{\cdot\}}$ is the indicator function. The corresponding MGF upper and lower envelopes have the rates [6]

$$\begin{aligned} r_u &= \frac{1}{2\theta}\left(P\theta - \mu - \lambda + \sqrt{(P\theta - \mu + \lambda)^2 + 4\lambda\mu}\right) \\ r_l &= \frac{1}{2\theta}\left(P\theta + \mu + \lambda - \sqrt{(P\theta + \mu - \lambda)^2 + 4\lambda\mu}\right) \ . \end{aligned}$$

The case of Markov-modulated processes with more than two states can be treated using results from [19].

C. Output Bounds in Eq. (1)

In addition to the two probabilistic arrival models described earlier, we also need probabilistic models to describe the service received by arrival flows at a scheduler. The network calculus uses the concept of service curves, which model lower bounds on the amount of service received by a flow, or an aggregate of flows, at a server. The service curves are essentially functions, or random processes, whose expressions depend on the service rate, the arrival representation of the competing flow, the scheduling algorithm, and even the packetization model [11]. In this paper we use a probabilistic service curve model from [3].

Formally, a bi-variate random process $S(s,t)$ is a (probabilistic) service curve for an arrival process $A(t)$ if the corresponding departure process $D(t)$ satisfies for all $s \leq t$

$$D(t) \geq A * S(t) \ , \tag{7}$$

where $A * S(t) := \inf_{0 \leq s \leq t} \{A(s) + S(s,t)\}$ denotes the $(min, +)$ convolution of $A(t)$ and $S(s,t)$. In other words, the service curve $S(s,t)$ models probabilistic lower bounds on the service received by the arrival process $A(t)$.

We are now ready to show that overloaded FIFO schedulers yield probabilistic fair shares for the incoming EBB arrival processes. The next theorem considers both the cases of statistically independent arrivals, and also of possibly statistically correlated arrivals, without the assumption that the increments of flows are independent.

Theorem 1: (OUTPUT BOUNDS: EBB ARRIVALS) Consider a FIFO node with fixed capacity $R > 0$, and serving two arrival flows $A(t)$ and $A'(t)$. Fix $\theta > 0$ and assume that $A(t)$ is bounded by an MGF lower envelope with rate r_l, and $A'(t)$ is bounded by MGF upper and lower envelopes with rates r'_u and r'_l, respectively. If $r_l + r'_l > R$ then the FIFO scheduler offers the probabilistic fair share

$$Pr\Big(D(t) < \frac{r_l}{r_l + r'_u} Rt - \sigma\Big) \leq \varepsilon(\sigma) \ , \tag{8}$$

with the error function $\varepsilon(\sigma)$. When $A(t)$ and $A'(t)$ are statistically independent,

$$\varepsilon(\sigma) = e^{-\theta\sigma} + 2\left(\frac{e(r_l + r'_l)}{r_l + r'_l - R}\right)^{\frac{1}{2}} e^{-\frac{\theta}{2}\sigma} \ . \tag{9}$$

Without the independence assumption the error function is

$$\varepsilon(\sigma) = e^{-\theta\sigma} + 3\left(\frac{e(r_l + r'_l)}{r_l + r'_l - R}\right)^{\frac{2}{3}} e^{-\frac{\theta}{3}\sigma} \ , \tag{10}$$

where e is Euler's constant.

Before deriving the proof, let us make some remarks. The overload condition $r_l + r'_l > R$ is critical for the error function $\varepsilon(\sigma)$ to be a real function. We also note that the theorem does not need an MGF upper envelope for $A(t)$; this is only needed if one further seeks a probabilistic fair share for $A'(t)$. By denoting the average rates of $A(t)$ and $A'(t)$ by r and r', respectively, we can derive Eq. (1) from Eq. (8) by letting

$$\delta = \frac{r_l}{r_l + r'_u} - \frac{r}{r + r'} \ ,$$

and preserving the error function from Eq. (8). Note that $\delta > 0$ according to the relationship between lower, average, and upper rates from Eq. (4).

We also point out that the error function increases by dispensing with the independence assumption, which means that FIFO's convergence in σ to the proportional share is slower than in the case of independent arrivals. This is due to the fact that large bursts in both $A(t)$ and $A'(t)$ are likely to occur simultaneously, unless the assumption of independence between $A(t)$ and $A'(t)$ is enforced.

PROOF. Denote by $[a]_+ = \max\{a, 0\}$ the positive part of a number a. From [9] we have that for any $x > 0$ the random process

$$S(s,t) = [R(t-s) - A'(s, t-x)]_+ 1_{\{t-s>x\}}$$

is a probabilistic service curve for $A(t)$, i.e., $D(t) \geq A * S(t)$. Let $\theta > 0$, and r_l, r'_l, and r'_u as in the theorem. Let us choose

$$x = \left(1 - \frac{R}{r_l + r'_u}\right) t \ . \qquad (11)$$

We point out that this choice of the time parameter x, which appears in the expression of the service curve $S(s,t)$ above, is the *key* to prove the relationship stated in Eq. (1) for overloaded regimes.

Using the definition of the service curve from Eq. (7), we can now bound the probability in Eq. (8) as follows

$$\begin{aligned} Pr&\left(D(t) < \frac{r_l}{r_l + r'_u} Rt - \sigma\right) \\ &\leq Pr\left(A * S(t) < \frac{r_l}{r_l + r'_u} Rt - \sigma\right) \\ &\leq Pr\Bigg(\sup_{0 \leq s \leq t} \Big\{\frac{r_l}{r_l + r'_u} Rt - A(s) \\ &\qquad - [R(t-s) - A'(s, t-x)]_+ 1_{\{t-s>x\}}\Big\} > \sigma\Bigg) \end{aligned}$$

By separating the supremum in two parts, i.e., $s \geq t - x$ and $s < t - x$, we can further bound the probability by

$$\begin{aligned} &Pr\left(\frac{r_l}{r_l + r'_u} Rt - A(t-x) > \sigma\right) \\ &\quad + Pr\Bigg(A'(t-x) - \frac{r'_u}{r_l + r'_u} Rt \\ &\qquad + \sup_{0 \leq s < t-x} \{Rs - A(s) - A'(s)\} > \sigma\Bigg) \\ &\leq e^{-\theta\sigma} + Pr\left(A'(t-x) - \frac{r'_u}{r_l + r'_u} Rt > \sigma_1\right) \\ &\quad + Pr\left(\sup_{0 \leq s < t-x} \{Rs - A(s) - A'(s)\} > \sigma_2\right) \ , \end{aligned} \qquad (12)$$

where $\sigma_1 + \sigma_2 = \sigma$. the last inequality follows by applying the Chernoff and Boole bounds.[2]

With the choice of x from Eq. (11) and the MGF upper envelope for $A'(t-x)$ from Eq. (2), the Chernoff bound yields that the second term in the sum is bounded by

$$Pr\left(A'(t-x) - \frac{r'_u}{r_l + r'_u} Rt > \sigma_1\right) \leq e^{-\theta\sigma_1} \ .$$

To bound the third term in the sum we introduce a discretization parameter $\tau_0 > 0$. Furthermore, for some $s \geq 0$, we denote $j = \lfloor \frac{s}{\tau_0} \rfloor$ the integer part of $\frac{s}{\tau_0}$. Assume now that $A(t)$ and $A'(t)$ are statistically independent. Then the third is bounded by

$$\begin{aligned} Pr&\left(\sup_{0 \leq s < t-x} \{Rs - A(s) - A'(s)\} > \sigma_2\right) \\ &\leq \sum_{j=0}^{\infty} Pr\left(R(j+1)\tau_0 - A(j\tau_0) - A'(j\tau_0) > \sigma_2\right) \\ &\leq e^{\theta R \tau_0} \sum_{j=0}^{\infty} e^{\theta(R - r_l - r'_l) j \tau_0} e^{-\theta\sigma_2} \\ &\leq e^{\theta R \tau_0} \frac{e^{-\theta(R - r_l - r'_l)\tau_0}}{\theta(r_l + r'_l - R)\tau_0} e^{-\theta\sigma_2} \leq \frac{e(r_l + r'_l)}{r_l + r'_l - R} e^{-\theta\sigma_2} \ . \end{aligned} \qquad (13)$$

In the first inequality we used Boole inequality. Then we used the independence between $A(t)$ and $A'(t)$, and the MGF lower envelopes for $A(j\tau_0)$ and $A'(j\tau_0)$ from Eq. (3). Finally we applied the inequality $\sum_{j \geq 0} e^{-aj} \leq e^a / a$ for any $a > 0$, and optimized the resulted convex function with $\tau_0 = \frac{1}{\theta(r_l + r'_l)}$.

Collecting terms yields

$$\begin{aligned} Pr&\left(D(t) < \frac{r_l}{r_l + r'_u} Rt - \sigma\right) \\ &\qquad \leq e^{-\theta\sigma} + e^{-\theta\sigma_1} + \frac{e(r_l + r'_l)}{r_l + r'_l - R} e^{-\theta\sigma_2} \end{aligned}$$

Next we use the convex optimization for $M > 0$ [5]

$$\inf_{\sigma_1 + \sigma_2 = \sigma} \left\{e^{-\theta\sigma_1} + M e^{-\theta\sigma_2}\right\} \leq 2 M^{\frac{1}{2}} e^{-\frac{\theta}{2}\sigma} \ .$$

The proof for statistically independent $A(t)$ and $A'(t)$ is thus complete.

In turn, when $A(t)$ and $A'(t)$ are not necessarily statistically independent, we can bound the probabilities in the sum from the second line of Eq. (13) by a sum of two probabilities, using Boole inequality, i.e.,

$$\begin{aligned} Pr&\left(\frac{R}{2}\tau_0 + \frac{R - r_l - r'_l}{2} j\tau_0 - \left(A(j\tau_0) - r_l j\tau_0\right) > \sigma_3\right) \\ &+ Pr\left(\frac{R}{2}\tau_0 + \frac{R - r_l - r'_l}{2} j\tau_0 - \left(A'(j\tau_0) - r'_l j\tau_0\right) > \sigma_4\right) \ , \end{aligned}$$

where $\sigma_3 + \sigma_4 = \sigma_2$. The bound from Eq. (13) then becomes

$$M e^{-\theta\sigma_3} + M e^{-\theta\sigma_4} \ ,$$

[2] For a random variable X, the Chernoff bound states that $Pr\left(e^{\theta X} \geq x\right) \leq E\left[e^{\theta X}\right] / x$ for all $\theta, x > 0$. For two events A and B, the Boole bound states that $Pr(A \cup B) \leq Pr(A) + Pr(B)$.

where the value

$$M = \frac{e(r_l + r_l')}{r_l + r_l' - R}$$

is obtained exactly as in the last lines of Eq. (13). Finally, by applying the inequality [5]

$$\inf_{\sigma_1+\sigma_3+\sigma_4=\sigma} \left\{ e^{-\theta\sigma_1} + Me^{-\theta\sigma_3} + Me^{-\theta\sigma_4} \right\} \leq 3M^{\frac{2}{3}} e^{-\frac{\theta}{3}\sigma} ,$$

completes the proof. □

Now, we turn to the specialized case of EBB arrivals with statistically independent increments, i.e., $A(s_1, s_2)$ and $A(s_3, s_4)$ are statistically independent for all $0 \leq s_1 \leq s_2 \leq s_3 \leq s_4$. Note that the Poisson or the compound Poisson process considered earlier belong to this category. By accounting for the independent increments property, we can improve the bounds from Theorem 1 as follows.

Theorem 2: (OUTPUT BOUNDS: EBB ARRIVALS WITH INDEPENDENT INCREMENTS) Consider the hypothesis from Theorem 1. In addition, assume that both $A(t)$ and $A'(t)$ have independent increments. If $r_l + r_l' > R$, then the FIFO scheduler offers the probabilistic fair share

$$Pr\Big(D(t) < \frac{r_l}{r_l + r_u'} Rt - \sigma\Big) \leq \varepsilon(\sigma) , \qquad (14)$$

with the error function, when $A(t)$ and $A'(t)$ are statistically independent,

$$\varepsilon(\sigma) = e^{-\theta\sigma} + 2e^{-\frac{\theta}{2}\sigma} . \qquad (15)$$

Without the independence assumption between $A(t)$ and $A'(t)$, the error function is

$$\varepsilon(\sigma) = e^{-\theta\sigma} + 3e^{-\frac{\theta}{3}\sigma} . \qquad (16)$$

We point out that the (over)utilization factor $u = \frac{r_l + r_l'}{R}$ plays a critical role in the convergence. In marginally overloaded regimes, i.e., when $u \approx 1$, the bound from the theorem considerably improves upon the bound from Theorem 1: the second exponential in the error function from Theorem 1 is reduced by a factor of as much as $\Omega\left(\frac{1}{u-1}\right)$. In turn, in extreme overloaded regimes, i.e., when $u \gg 1$, the gain is almost negligible: the second exponential is only reduced by a factor of roughly $\sqrt{e}$.

PROOF. The proof of the theorem proceeds along the same lines as the proof of Theorem 1, with the exception that the sample path probability from Eq. (12), more exactly the third term in the sum, is now evaluated using a supermartingale argument [20], [4].

Consider first the case when $A(t)$ and $A'(t)$ are statistically independent, and let us introduce the process

$$T(s) = e^{\theta\left(Rs - A(s) - A'(s)\right)}$$

for all $s \geq 0$. Consider also the filtration of σ-algebras

$$\mathcal{F}_s = \sigma\left\{(A(u), A'(u)) : 0 \leq u \leq s\right\} ,$$

i.e., $\mathcal{F}_s \subseteq \mathcal{F}_t$ for all $s \leq t$. Note that $T(s)$ is $\mathcal{F}_s$-measurable for all $s \geq 0$. Then we can write for the conditional expectations for all $s, u \geq 0$

$$\begin{aligned} E\left[T(s+u) \parallel \mathcal{F}_s\right] &= E\left[T(s)e^{\theta(Ru - A(s,s+u) - A'(s,s+u))} \parallel \mathcal{F}_s\right] \\ &= T(s)E\left[e^{\theta(Ru - A(s,s+u) - A'(s,s+u))} \parallel \mathcal{F}_s\right] \\ &= T(s)E\left[e^{\theta(Ru - A(s,s+u) - A'(s,s+u))}\right] \\ &= T(s)e^{\theta(R - r_l - r_l')u} \leq T(s) . \end{aligned}$$

In the second line we applied that $T(s)$ is $\mathcal{F}_s$-measurable, and then we used the independent increments property of $A(t)$ and $A'(t)$, e.g., $A(s, s+u)$ is independent of $\mathcal{F}_s$. Finally we used the MGF lower envelopes for $A(s, s+u)$ and $A'(s, s+u)$, and the overload condition from the theorem. Therefore, the process $T(s)$ is a supermartingale, i.e., [13]

$$E\left[T(s+u) \parallel \mathcal{F}_s\right] \leq T(s)$$

for all $s, u \geq 0$. We can thus evaluate the third term in the sum from Eq. (12) as

$$\begin{aligned} &Pr\left(\sup_{0\leq s<t-x} \left\{Rs - A(s) - A'(s)\right\} > \sigma_2\right) \\ &\leq Pr\left(\sup_{s\geq 0} T(s) > e^{\theta\sigma_2}\right) \\ &\leq e^{-\theta\sigma_2} . \end{aligned}$$

In the last line we applied Doob inequality for the supermartingale $T(s)$ [13]. The rest of the first part of the proof follows immediately by collecting terms.

When $A(t)$ and $A'(t)$ are not necessarily independent we proceed as in the second part of the proof of Theorem 1. Concretely, we bound the sample-path probability from Eq. (12) with

$$\begin{aligned} &Pr\left(\sup_{s\geq 0}\left\{\frac{R - r_l - r_l'}{2} - (A(s) - r_l s)\right\} > \sigma_3\right) \\ &+ Pr\left(\sup_{s\geq 0}\left\{\frac{R - r_l - r_l'}{2} - (A'(s) - r_l' s)\right\} > \sigma_4\right) , \end{aligned}$$

where $\sigma_3 + \sigma_4 = \sigma_2$. The rest of the proof then follows using the same supermartingale argument as above. □

IV. SCENARIO 2: HEAVY-TAILED AND SELF-SIMILAR ARRIVALS

In this section we extend the analysis from Section III to the class of heavy-tailed and self-similar arrival processes. First we describe two statistical envelope models, including a novel one for enforcing overloaded regimes, and then present bounds on the departure processes as in Eq. (1).

A. Upper and Lower Statistical Envelopes

To capture the properties of heavy-tailed and self-similar arrivals $A(t)$, we use the following (upper) statistical envelope model for all $0 \leq s \leq t$ and $\sigma \geq 0$ [23]

$$Pr\Big(A(s,t) > r(t-s) + \sigma(t-s)^H\Big) \leq K\sigma^{-\alpha} . \qquad (17)$$

Here, r is the long-term arrival rate. The tail index $-\alpha$ describes the shape of the tail and smaller values of α indicate heavier tails; we are interested in the case when $\alpha \in (1, 2)$, i.e., arrivals with finite mean but infinite variance. The self-similarity index H (or Hurst parameter) satisfies $H \in [1/\alpha, 1)$ and describes the arrivals behavior when rescaling time. The burst parameter σ captures the tail behavior, and $K > 0$ is a constant. The function $G(t) = rt + \sigma t^H$ is called a *htss*-envelope, i.e., heavy-tailed and self-similar (upper) envelope [23].

The key characteristic of the *htss*-envelope is that the error function $\varepsilon(\sigma) = K\sigma^{-\alpha}$ is given by a power law. This means that the arrivals $A(s, t)$ may deviate from the mean $r(t-s)$ by some very large values with non-negligible probabilities; moreover, these deviations increase as a function of time due to self-similarity.

Similar to the case of EBB arrivals, we also need to introduce a lower envelope model for heavy-tailed and self-similar arrivals in order to give sufficient conditions for overloaded regimes. To this end we propose the following (lower) statistical envelope model for an arrival process $A(t)$, such that for all $0 \le s \le t$ and $\sigma \ge 0$

$$Pr\Big(A(s,t) < r(t-s) - \sigma(t-s)^H\Big) \le e^{-L\sigma^{\frac{\alpha}{\alpha-1}}} \quad . \qquad (18)$$

Here, r is the long-term arrival rate. The parameter $\alpha \in (1, 2]$ defines the Weibull distribution for the error function, with smaller values indicating heavier tails. The self-similarity index H is as in Eq. (17), and describes the arrivals behavior when rescaling time. Finally, $L > 0$ is a constant and the function $G(t) = rt + \sigma t^H$ is called a *wss*-envelope, i.e., Weibull and self-similar (lower) envelope.

In contrast to the *htss*-envelope, the *wss*-envelope has an error function with a Weibull tail. An explanation for this choice is that large deviations of the arrivals to the left of the mean are much less likely to happen than to the right of the mean, especially at small time scales, because $A(t)$ is non-negative. Moreover, the decay of the deviations is faster than the exponential (note that $\frac{\alpha}{\alpha-1} > 1$), yet the deviations increase as a function of time due to self-similarity.

B. Example: Multiplexed Heavy-Tailed On-Off

Here we construct the *htss* and *wss*-envelopes from Eqs. (17) and (18), respectively, for an aggregate of N independent On-Off flows consisting of alternating and statistically independent 'On' and 'Off' periods [25]. Each flow transmits at a constant rate P during 'On' periods and is silent during 'Off' periods. Unlike the exponential distributions in the states for the MMOO process from Section III, the right-tail distribution of the 'On' periods is now given by

$$Pr\Big(X > \sigma\Big) = \sigma^{-\alpha} \ ,$$

where $\alpha \in (1, 2)$. The distribution of the 'Off' periods is given similarly with tail index $\alpha_{\text{off}} > \alpha$. The means of the 'On' and 'Off' periods are μ_{on} and μ_{off}, respectively. The per-flow arrival rate is $r = \frac{\mu_{\text{on}} P}{\mu_{\text{on}} + \mu_{\text{off}}}$.

The cumulative process $A(t)$ generated by all flows can be written for all $t \ge 0$

$$A(t) = \int_0^t PM(s)ds \ , \qquad (19)$$

where $M(s)$ denotes the number of active flows at time s.

To construct the *htss*-envelope for $A(t)$ we first need to introduce α-stable processes [26], which typically appear as the limit of normalized sums of i.i.d. random variables, e.g., by the generalized central limit theorem (GCLT) [32]. Explicit expressions for the densities of stable distributions are known only in very special cases, e.g., the Gaussian. However, all stable distributions have explicit descriptions for their characteristic functions (see [32], pp. xvi) in terms of four parameters: the stability index $\alpha \in (0, 2]$, the skewness $\beta \in [-1, 1]$, the scale $a > 0$, and the location $\mu \in \mathbb{R}$. According to these parameters a stable random variable is denoted by $S_\alpha(\beta, a, \mu)$; its normalized version is denoted by

$$S_\alpha(\beta) := S_\alpha(\beta, 1, 0) = \frac{S_\alpha(\beta, a, \mu) - \mu}{a} \ .$$

Now we look at the scaled process $A(Tt)$ for large time scale T, and show how to construct *htss*-envelopes for it. By taking a suitable simultaneous limit in T and N, then $A(Tt)$ converges in distribution to an α-stable Lévy motion. This is an α-stable process with stationary and independent increments, càdlàg sample paths (i.e. right-continuous and limits from the left exist), and self-similar with index $1/\alpha$ [25]. Formally

$$Pr\left(A(Tt) > NTtr + (NT)^{\frac{1}{\alpha}}\sigma\right) \to_{T,N} Pr\left(\sigma_0 S_\alpha\left(1, c_\alpha t^{\frac{1}{\alpha}}, 0\right) > \sigma\right) \ , \qquad (20)$$

where $\sigma_0 = \frac{\mu_{\text{off}}}{(\mu_{\text{on}} + \mu_{\text{off}})^{1+1/\alpha}}$, and $c_\alpha = \left(\frac{2\Gamma(\alpha)\sin\frac{\pi\alpha}{2}}{\pi}\right)^{-\frac{1}{\alpha}}$, and $\Gamma(\cdot)$ is the Gamma function; we point out that here we rewrote the result from [25] using Euler's equality $\Gamma(\alpha)\Gamma(1-\alpha) = \frac{\pi}{2\sin(\frac{\pi\alpha}{2})\cos(\frac{\pi\alpha}{2})}$.

Using the tail approximation of α-stable distributions [26]

$$Pr\Big(S_\alpha(1) > \sigma\Big) \sim (c_\alpha\sigma)^{-\alpha} \quad , \ \sigma \to \infty \ , \qquad (21)$$

and arranging terms in Eq. (20), we obtain the *htss*-envelope with parameters

$$Nr, \quad \alpha, \quad \eta = \frac{1}{\alpha}, \quad K \approx N\sigma_0^\alpha \ . \qquad (22)$$

Next, to construct the *wss*-(lower)-envelope, we use the Laplace transform of the α-stable process $S_\alpha(1, a, 0)$, i.e., for all $\theta > 0$ [29]

$$E\left[e^{-\theta S_\alpha(1,a,0)}\right] = e^{-\frac{a^\alpha}{\cos\frac{\pi\alpha}{2}}\theta^\alpha} \ , \qquad (23)$$

which in particular holds for our choice of $\alpha \in (1, 2]$. From this and the Chernoff bound we obtain for some $\theta > 0$

$$Pr\left(A(Tt) < NTtr - (Tt)^{\frac{1}{\alpha}}\sigma\right) \le e^{-\frac{(\sigma_0 c_\alpha)^\alpha t}{\cos\frac{\pi\alpha}{2}}\theta^\alpha} e^{-\theta\left(\frac{t}{N}\right)^{\frac{1}{\alpha}}\sigma}$$
$$\le e^{-(\alpha-1)\left(-\frac{\cos\frac{\pi\alpha}{2}}{N}\right)^{\frac{1}{\alpha-1}}\left(\frac{1}{\alpha c_\alpha \sigma_0}\right)^{\frac{\alpha}{\alpha-1}}\sigma^{\frac{\alpha}{\alpha-1}}} \ .$$

In the second line we optimized $\theta = \left(-\frac{\sigma \cos\frac{\pi\alpha}{2}}{\alpha(\sigma_0 c_\alpha)^\alpha \left(\frac{N}{t}\right)^{\frac{1}{\alpha}}}\right)^{\frac{1}{\alpha-1}}$, which eliminates the time parameter t in the exponent. Therefore we get the *wss*-envelope with parameters

$$Nr,\ \alpha,\ H=\frac{1}{\alpha},\ L=(\alpha-1)\left(-\frac{\cos\frac{\pi\alpha}{2}}{N}\right)^{\frac{1}{\alpha-1}}\left(\frac{1}{\alpha c_\alpha \sigma_0}\right)^{\frac{\alpha}{\alpha-1}} \quad (24)$$

Finally, we note that for the construction from (24), if $A(t)$ and $A'(t)$ are statistically independent and have the same parameters except for the average rates, i.e., r and r', respectively, then the sum $A(t)+A'(t)$ has the *wss*-envelope with parameters

$$r+r',\quad \alpha,\quad H,\quad L\ , \quad (25)$$

with L from Eq. (24).

C. Output Bounds in Eq. (1)

The next theorem extends Theorem 1 to heavy-tailed and self-similar arrivals. We now consider only the case of statistically independent arrivals; the case of not necessarily independent arrivals can be also considered similarly by constructing a *wss*-envelope for the sum, as in Eq. (25), but after using the splitting argument from the proof of Theorem 1. On the other hand, an extension of Theorem 2 is not possible, since self-similar processes do not have statistically independent increments.

Theorem 3: (OUTPUT BOUNDS: HEAVY-TAILED AND SELF-SIMILAR ARRIVALS) Consider a FIFO node with fixed capacity $R>0$, and serving two statistically independent arrival flows $A(t)$ and $A'(t)$. Assume that $A(t)$ is bounded by a *htss*-envelopes with parameters r,α,H,K, and a *wss*-envelope with parameters r,α,H,L. Assume also that $A'(t)$ and $A(t)+A'(t)$ are bounded by *wss*-envelopes with parameters r',α,H,L, and $r+r',\alpha,H,L$, respectively, according to Eq. (25). If $r+r'>R$, then the FIFO scheduler offers the probabilistic fair share for all $t,\sigma\geq 0$

$$Pr\Big(D(t) < \frac{r}{r+r'}Rt - \sigma t^H\Big) \leq \varepsilon(\sigma)\ , \quad (26)$$

with the error function

$$\varepsilon(\sigma) = e^{-L\sigma^{\frac{\alpha}{\alpha-1}}} + \inf_{\sigma_1+\sigma_2=\sigma}\left\{K_1\sigma_1^{-\alpha} + K_2 e^{-L\sigma_2^{\frac{\alpha}{\alpha-1}}}\right\}\ ,$$

where K_1 and K_2 are given in the proof.

The error function from the theorem is dominated by the power function with tail index $-\alpha$, which indicates that FIFO's convergence in σ to the proportional share is very slow. This is due to the fact that the FIFO scheduler may have to continuously serve extremely large bursts of $A'(t)$, which can occur with non-negligible probabilities. Let us also remark that the theorem can be extended to the case when $A(t)$ is an EBB process, whereas $A'(t)$ is a heavy-tailed and similar process. In this case, the error function would be still dominated by the power function with tail index $-\alpha$, as in Eq. (26).

PROOF. Fix $t,\sigma\geq 0$, and choose

$$x = \left(1-\frac{R}{r+r'}\right)t\ .$$

Similarly as in the proof of Theorem 1, this choice of the time parameter x is critical for the proof. Following the proof of Theorem 1, we can bound the probability in Eq. (26) by

$$\begin{aligned} &Pr\left(\frac{r}{r+r'}Rt - A(t-x) > \sigma t^H\right) \\ &+ Pr\left(A'(t-x) - \frac{r'}{r+r'}Rt > \sigma_1 t^H\right) \\ &+ Pr\left(\sup_{0\leq s<t-x}\{Rs - A(s) - A'(s)\} > \sigma_2 t^H\right)\ , \end{aligned}$$

where $\sigma_1+\sigma_2=\sigma$.

According to the definition of the *wss*-envelope from Eq. (18), the first term in the sum is bounded by

$$Pr\left(\frac{r}{r+r'}Rt - A(t-x) > \sigma t^H\right) \leq e^{-L\sigma^{\frac{\alpha}{\alpha-1}}}\ .$$

Next, from the definition of the *htss*-envelope from Eq. (17), the second term in the sum is bounded by

$$Pr\left(A'(t-x) - \frac{r'}{r+r'}Rt > \sigma_1 t^H\right) \leq K'\left(\frac{\sigma_1}{\left(\frac{R}{r+r'}\right)^H}\right)^{-\alpha}\ ,$$

such that we can set

$$K_1 = K\left(\frac{r+r'}{R}\right)^{H\alpha}\ . \quad (27)$$

Finally, to bound the third probability in the sum, we introduce a discretization parameter $\tau_0>0$. For some $s\geq 0$ we denote $j=\lfloor\frac{s}{\tau_0}\rfloor$ the integer part of $\frac{s}{\tau_0}$. Then the probability is bounded, using Boole inequality, by

$$\begin{aligned} &\sum_{j\geq 0} Pr\Big(\left(A+A'\right)(j\tau_0) < (r+r')\,j\tau_0 \\ &\quad -(j\tau_0)^H\left(\sigma_2 + (r+r'-R)(j\tau_0)^{1-H} - R\tau_0^{1-H}j^{-H}\right)\Big)\ . \end{aligned}$$

Setting $\tau_0 = \left(\frac{1}{R}\right)^{\frac{1}{1-H}}$, we can ignore the term $R\tau_0^{1-H}j^{-H}$ which becomes negligible relative to the second term in the sum, for sufficiently large j. Let us now introduce the positive parameter $\gamma=(r+r'-R)/R$. Applying the definition of the *wss*-envelope from Eq. (18), the sum is further bounded by

$$\sum_{j\geq 0} e^{-L\left(\sigma_2+\gamma j^{1-H}\right)^{\frac{\alpha}{\alpha-1}}} \leq e^{-L\sigma_2^{\frac{\alpha}{\alpha-1}}}\sum_{j\geq 0} e^{-L\gamma^{\frac{\alpha}{\alpha-1}} j^{\frac{\alpha(1-H)}{\alpha-1}}}$$

Since the string in the sum is non-increasing, the sum is bounded by [**?**]

$$1 + \int_0^\infty e^{-ax^b}dx = 1 + \frac{\Gamma\left(\frac{1}{b}\right)}{ba^{\frac{1}{b}}}\ ,$$

with the parameters $a = L\gamma^{\frac{\alpha}{\alpha-1}}$ and $b=\frac{\alpha(1-H)}{\alpha-1}\leq 1$. Therefore, the parameter K_2 from the error function is $K_2 =$

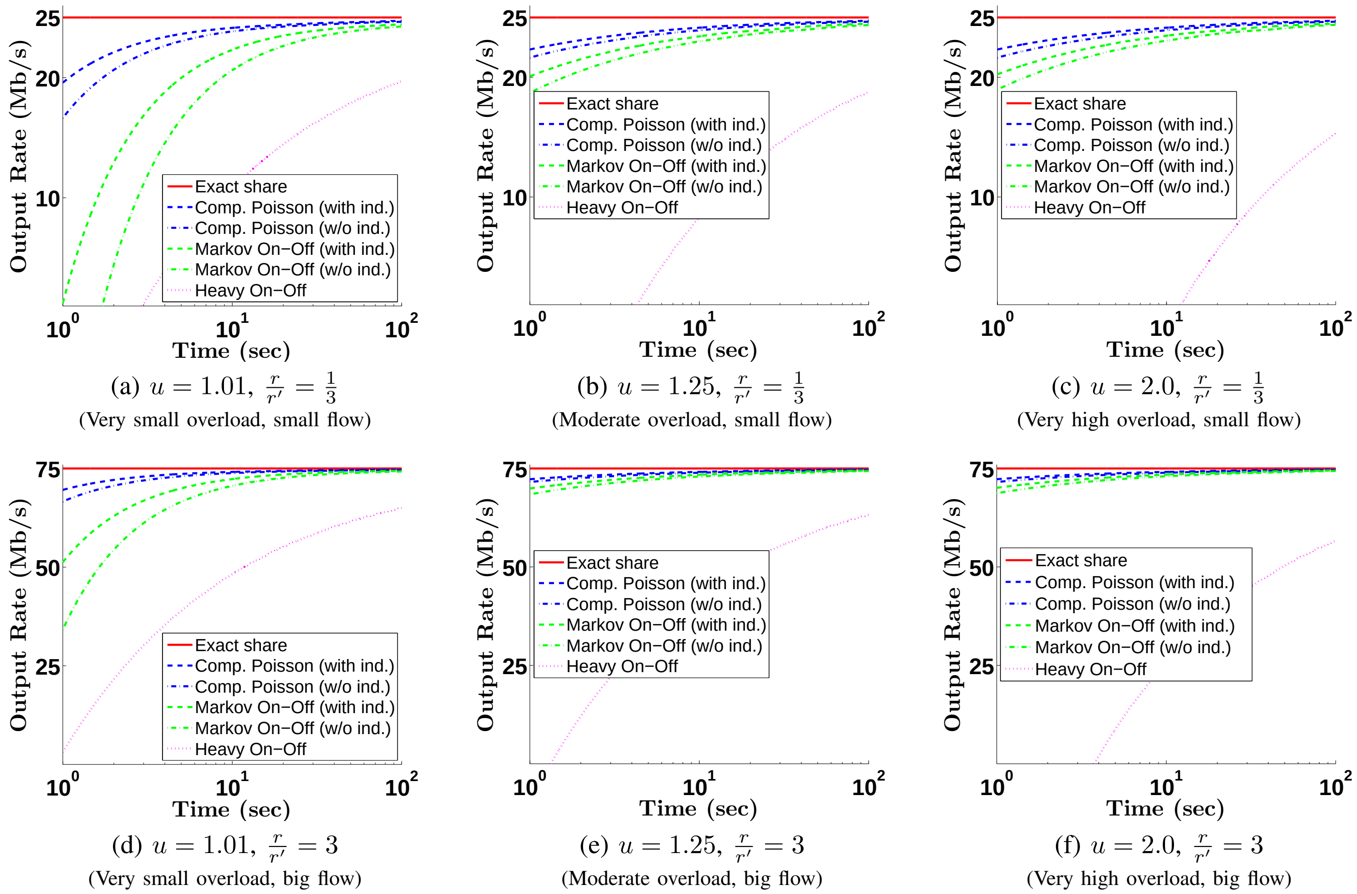

(a) $u = 1.01,\ \frac{r}{r'} = \frac{1}{3}$
(Very small overload, small flow)

(b) $u = 1.25,\ \frac{r}{r'} = \frac{1}{3}$
(Moderate overload, small flow)

(c) $u = 2.0,\ \frac{r}{r'} = \frac{1}{3}$
(Very high overload, small flow)

(d) $u = 1.01,\ \frac{r}{r'} = 3$
(Very small overload, big flow)

(e) $u = 1.25,\ \frac{r}{r'} = 3$
(Moderate overload, big flow)

(f) $u = 2.0,\ \frac{r}{r'} = 3$
(Very high overload, big flow)

Fig. 2. Probabilistic lower bounds on the output rate of $D(t)$ in Eq. (1), for a compound Poisson process from Eq. (5) ($E[X_i] = 400\ (Bytes)$, independent (Eq. (15)) and not independent (Eq. (16)) arrivals), MMOO process from Eq. (6) ($\frac{1}{\mu} = .4$ s, $\frac{1}{\lambda} = .6$ s, $P = 64$ Kb/s [30], independent (Eq. (9)) and not independent (Eq. (10)) arrivals), and Heavy-tailed On-Off process from Eq. (19) ($\alpha = 1.75,\ \alpha_{\text{off}} = 1.95,\ P = 64$ Kb/s) (server rate $R = 100$ Mb/s, (over)utilization factors ($u = 1.01$, $u = 1.25$ and $u = 2.0$), violation probability $\varepsilon = 10^{-3}$)

$1 + \frac{\Gamma(\frac{1}{b})}{ba^{\frac{1}{b}}}$. In the special case when $H = \frac{1}{\alpha}$ (see Eq. (24)), we have that $b = 1$ and the expression of K_2 becomes

$$K_2 = 1 + \frac{1}{L}\left(\frac{R}{r + r' - R}\right)^{\frac{\alpha}{\alpha - 1}} \quad . \tag{28}$$

The proof is thus complete. □

V. Numerical Results

In this section we illustrate the convergence of the output rates from Eq. (1) to the fair shares with respect to different arrival processes, flow sizes and intensity of overload. We consider both EBB and heavy-tailed/self-similar arrival processes. For the EBB case, we show both cases of independent and not necessarily independent arrivals.

We present numerical examples mainly characteristic to Internet traffic, so that the rates of arrival processes and output processes are measured in Mb/s. Note that the numerical units can be adjusted according to the application of queueing systems, e.g., web services or CPU cores. Fig. 2 depicts the lower bounds on the output rates for arrivals of (1) compound Poisson as in Eq. (14), with and without the independence assumption of flows, (2) MMOO as in Eq. (8), with and without the independence assumption of flows, and (3) Multiplexed Heavy-Tailed On-Off as in Eq. (26). Specific numerical values used to compute r_l, r'_l , and r_u in the lower bounds for all three arrival processes are given in the caption. Two flow sizes of $A(t)$, i.e., 25% and 75%, corresponding to $\frac{r}{r'} = \frac{1}{3}$ and $\frac{r}{r'} = 3$, are considered in Fig. 2 (a,b,c) and (d,e,f), respectively. Also, three intensities of overutilization, $u = 1.01$, $u = 1.25$ and $u = 2.0$ are considered in Fig. 2 (a,d), (b,e), and (c,f), respectively.

Besides the expected behavior that burstier sources take longer to converge, one can observe that increasing the load of $A(t)$ (in (d,e,f), as opposed to (a,b,c)) significantly increases the convergence speed to the fair share, which indicates that FIFO biases proportional fairness towards larger flows. Moreover, the independence assumption for EBB flows has a significant impact on the convergence rates, especially at barely overutilizations (i.e., $u = 1.01$). The immediate implication is that statistical multiplexing strongly manifests itself in overloaded queues, as it is known for underloaded queues.

An interesting observation from Fig. 2 is on how the utilization factor influences the convergence speed. Namely, for the EBB case, higher utilizations lead to faster convergence speeds to the fair shares, and also a stabilization effect after $u = 1.25$ (further increasing the utilization does not change the convergence speed). Note that this behavior is exactly the opposite as in underloaded queues, where higher utilizations lead to slower convergence rates (of the output rates to the input rates). This observation can be analytically explained by the error functions (e.g., in Eq. (9)) which roughly increase as $\mathcal{O}\left(\frac{1}{u-1}\right)$; in contrast, in underloaded queues, the error functions roughly increase as $\mathcal{O}\left(\frac{1}{1-u}\right)$ [5]. However, heavy-tailed self-similar traffic preserves the nondecreasing convergence time behavior from underloaded queues, which can be explained by the interplay between $K_1 = \mathcal{O}(u)$ from Eq. (27) and $K_2 = \mathcal{O}\left(\frac{1}{u-1}\right)^{\frac{\alpha}{\alpha-1}}$ in Eq. (28), both appearing in the error function from Eq. (26). Note that K_1 and K_2 have opposite monotonicity, i.e., the former increasing in u and the latter decreasing in u.

VI. Conclusions

In this paper we have analyzed a FIFO queue in overloaded regimes, i.e., when the total arrival rate exceeds the service rate. We have considered both the class of EBB arrivals (including Markovian) and also the class of heavy-tailed and self-similar arrivals, which is adequate to model the characteristics of modern workloads. For both classes we have shown that FIFO guarantees probabilistic proportional fair shares in overloaded regimes, proving thus a conjecture from the literature in great generality. Specifically, we showed that the convergence rate to the fair share is nondecreasing in flow sizes. Most interestingly, we also showed that for EBB arrivals, the convergence rate is nonincreasing in the utilization factor for the overloaded regime; in turn, for heavy-tailed and self-similar arrivals, the corresponding convergence rate is nondecreasing, as it is generally the case for the underloaded regime. To conclude, our results indicate the need for sophisticated system oriented solutions for overload management, especially in the likely situation of highly varying load fluctuations in short term.

References

[1] J. Abate, G. L. Choudhury, and W. Whitt. Waiting-time tail probabilities in queues with long-tail service-time distributions. *Queueing Systems*, 16(3-4):311–338, Sept. 1994.

[2] G. Bolch, S. Greiner, H. Meer, and K. Trivedi. *Queueing Networks and Markov Chains: Modeling and Performance Evaluation With Computer Science Applications*. Wiley, 2006.

[3] C.-S. Chang. *Performance Guarantees in Communication Networks*. Springer Verlag, 2000.

[4] F. Ciucu. Network calculus delay bounds in queueing networks with exact solutions. In *International Teletraffic Congress (ITC)*, pages 495–506, 2007.

[5] F. Ciucu, A. Burchard, and J. Liebeherr. A network service curve approach for the stochastic analysis of networks. In *ACM Sigmetrics*, volume 33, pages 279–290, 2005.

[6] C. Courcoubetis and R. Weber. Buffer overflow asymptotics for a buffer handling many traffic sources. *Journal of Applied Probability*, 33(3):886–903, Sept. 1996.

[7] M. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic: evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, Dec. 1997.

[8] R. Cruz. A calculus for network delay, parts I and II. *IEEE Transactions on Information Theory*, 37(1):114–141, Jan. 1991.

[9] R. L. Cruz. SCED+: Efficient management of quality of service guarantees. In *IEEE Infocom*, pages 625–634, 1998.

[10] M. Fidler. An end-to-end probabilistic network calculus with moment generating functions. In *IEEE International Workshop on Quality of Service (IWQoS)*, pages 261–270, 2006.

[11] M. Fidler. A survey of deterministic and stochastic service curve models in the network calculus. *To appear in IEEE Communications Surveys and Tutorials*, 12(1), 2010.

[12] Y. Ghiassi-Farrokhfal and J. Liebeherr. Output characterization of constant bit rate traffic in FIFO networks. *IEEE Communications Letters*, 13(8):618–620, Aug. 2009.

[13] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, 2001.

[14] R. C. Hampshire, M. Harchol-Balter, and W. A. Massey. Fluid and diffusion limits for transient sojourn times of processor sharing queues with time varying rates. *Queueing Systems: Theory and Applications*, 53(1-2):19–30, June 2006.

[15] M. Harchol-Balter and A. Downey. Exploiting process lifetime distributions for dynamic load balancing. *ACM Transactions on Computer Systems*, 15(3):253–285, 1997.

[16] A. Jean-Marie and P. Robert. On the transient behavior of the processor sharing queue. *Queueing Systems*, 17(1-2):129–136, Mar. 1994.

[17] Y. Jiang and Y. Liu. *Stochastic Network Calculus*. Springer, 2008.

[18] M. Jonckheere, R. Nunez-Queija, and B. Prabhu. Performance analysis of traffic surges in multi-class communication networks. In *22nd International Teletraffic Congress (to appear)*, 2010.

[19] G. Kesidis, J. Walrand, and C. Chang. Effective bandwidths for multiclass Markov fluids and other ATM sources. *IEEE/ACM Transactions on Networking*, 1(4):424–428, Aug. 1993.

[20] J. F. C. Kingman. A martingale inequality in the theory of queues. *Cambridge Philosophical Society*, 60(2):359–361, Oct. 1964.

[21] L. Kleinrock. *Queueing Systems*, volume 1. John Wiley and Sons, 1975.

[22] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic. *IEEE/ACM Transactions on Networking*, 2(1):1–15, Feb. 1994.

[23] J. Liebeherr, A. Burchard, and F. Ciucu. Non-asymptotic delay bounds for networks with heavy-tailed traffic. In *IEEE Infocom*, 2010.

[24] B. Melander, M. Bjorkman, and P. Gunningberg. First-come-first-served packet dispersion and implications for TCP. In *IEEE Globecom*, volume 1-3, pages 2170–2174, 2002.

[25] T. Mikosch, S. Resnick, H. Rootzn, and A. Stegeman. Is network traffic approximated by stable Lévy motion or fractional Brownian motion? *Annals of Applied Probability*, 12(1):23–68, Feb. 2002.

[26] J. P. Nolan. *Stable Distributions - Models for Heavy Tailed Data*. Birkhäuser, 2010. In progress, Chapter 1 online at academic2.american.edu/∼jpnolan.

[27] K. Park and W. Willinger, editors. *Self-Similar Network Traffic and Performance Evaluation*. John Wiley & Sons, 2000.

[28] M. Roughan, D. Veitch, and M. P. Rumsewicz. Computing queue-length distributions for power-law queues. In *IEEE Infocom*, pages 356–363, 1998.

[29] G. Samorodnitsky and M. Taqqu. *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*. Chapman and Hall, 1994.

[30] M. Schwartz. *Broadband Integrated Networks*. Prentice Hall, 1996.

[31] D. Starobinski and M. Sidi. Modeling and analysis of power-tail distributions via classical teletraffic methods. *Queueing Systems*, 36(1-3):243–267, Nov. 2000.

[32] V. V. Uchaikin and V. M. Zolotarev. *Chance and Stability: Stable Distributions and Their Applications*. VSP Science, 1999.

[33] M. Welsh and D. Culler. Adaptive overload control for busy internet servers. In *USENIX Symposium on Internet Technologies and Systems*, pages 43–56, 2003.

[34] O. Yaron and M. Sidi. Performance and stability of communication networks via robust exponential bounds. *IEEE/ACM Transactions on Networking*, 1(3):372–385, June 1993.