Stochastic Analysis of Some Expedited Forwarding Networks

Milan Vojnović and Jean-Yves Le Boudec

Abstract—We consider stochastic guarantees for networks with aggregate scheduling, in particular, Expedited Forwarding (EF). Our approach is based on the assumption that a node can be abstracted by a service curve, and the input flows are regulated individually at the network ingress. Both of these assumptions are inline with EF [1], [2]. For a service curve node, we derive bounds on the complementary distributions of the steady-state backlog and backlog as seen by packet arrivals. We also give a bound on the long-run loss ratio for a service curve node where the buffer size is too small to guarantee loss-free operation. For a Packet Scale Rate Guarantee node [3], [1], we use the delay from backlog bound to obtain a probabilistic bound on the delay. Our analysis is exact under the given assumptions. Our results should help us to understand the performance of networks with ag gregate scheduling, and provide the basis for dimensioning such networks.

Keywords— Expedited Forwarding, Differentiated Services, Aggregate Scheduling, Statistical Multiplexing, Stochastic QoS, Service Curve, Packet Scale Rate Guarantee, Queueing, Loss Ratio, Network Calculus

I. INTRODUCTION

EXPEDITED FORWARDING (EF) is a per-hop behavior (PHB) of Differentiated Services (DiffServ) [1], [2]. With EF, individual flows (called "micro-flows", or "inputs" in this paper) are shaped separately at network access; from there on, they are served in an aggregate manner. Our objective is to derive probabilistic guarantees for EF networks.

The definition of EF PHB [1], [2] gives an abstract model of a node called as Packet Scale Rate Guarantee (PSRG) [1], [3]. A node is said to offer a PSRG with a rate c and a latency e to EF aggregate if the departure d_n of the n-th packet, in the order of arrivals, satisfies

$$d_n \le f_n + e$$

where f_n is given recursively as $f_0 = 0$ and

$$f_n = \max\{a_n, \min\{d_{n-1}, f_{n-1}\} + \frac{l_n}{c}\}, \ n \ge 1,$$

for the *n*-th packet arrival at time a_n of length l_n bits.

The intention of PSRG is to give a formally correct definition of the intuitive concept that the rate of the node guaranteed to the EF aggregate is at least c, with a tolerance e (the tolerance depends on the specific details of the node). A special case of PSRG is a scheduler that gives static non-preemptive priority to EF traffic over non-EF traffic; here the rate is the server rate and the latency is service time of a maximum-length non-EF packet. In general, though, it cannot be assumed that an Internet router is a simple scheduler; in contrast, PSRG is intended to model complex nodes, such as Internet routers, that consist of many components; viewed as black-boxes, such nodes are generally not work conserving. In this paper we use two properties of PSRG. First, PSRG with rate c and latency e implies the service curve $\beta(t) = [c(t - e) - L_{\max}]^+$ where L_{\max} is the maximum size of an EF packet (this is called a rate-latency service curve; we use the notation $x^+ := \max[x, 0]$). Formally:

(A1) We suppose a node offers to the aggregate of all EF traffic a service curve β , i.e. for all t there exists $s \leq t$ such that

$$A^*(t) \ge A(s) + \beta(t-s), \tag{1}$$

where $A^*(t)$ is the output data from the node on [0, t] and A(t) is the data which is accepted for service (i.e. not lost) at the input of the node during [0, t] [4], [5], [6].

Second, in Section IV, we use another property of PSRG, namely, the fact that delay can be bounded from backlog. In addition, we do the following assumptions.

(A2) We suppose EF traffic inputs (micro-flows) at the network ingress points are mutually independent.

This assumption is also made in other work [7]. Note that we make no independence assumption for flows inside the network. (A3) We suppose each EF input (micro-flow) at the network ingress point is regulated, that is to say, for a given input *i* there exists a wide-sense increasing function α_i such that

$$A_i^0(t) - A_i^0(s) \le \alpha_i(t-s)$$
, for any $s \le t$,

where $A_i^0(t)$ is the data observed on [0, t] of the input *i* at the network ingress point.

In general, we derive our results for arbitrary arrival curves, and, in particular, we study leaky-bucket regulated inputs; $\alpha_i(t) = \rho_i t + \sigma_i$. When we consider individual EF flows with identical arrival curve constraints, we say the input flows are homogeneously regulated (resp. heterogeneously regulated for non-identical arrival curve constraints).

(A4) We suppose $\mathbb{E}[A_i^0(t) - A_i^0(s)] \le \rho_i(t-s)$, for any $s \le t$, where

$$\rho_i = \lim_{t \to \infty} \frac{\alpha_i(t)}{t}.$$

Indeed, the assumption (A4) is implied for the input flows with stationary and ergodic increments [8], [9], but not *vice versa*. Thus, (A4) is a weaker assumption. Note that we allow for the input flows with non-stationary increments as long as (A4) is verified. However, for some of our results we need stationary ergodic increments of the inputs to ensure certain limits exist; we explicitly indicate when such an assumption is needed.

Our results are obtained by combining some queueing results based on stochastic comparisons (see [9] and references therein) with some concepts of network calculus (see [10] and references therein).

Author Affiliations: EPFL, CH-1015 Lausanne, Switzerland. E-mail: <mi lan.vojnovic@epfl.ch> and <jean-yves.leboudec@epfl.ch>

We now explain the organization of the paper and highlight our main findings. We discuss the state of the art in Section II. In Section III we give theoretical foundations of our work; the results given in this section are of general interest for statistical multiplexing of regulated inputs to a multiplexer that offers a service curve to the aggregate input. Our prior work [9] gives us a catalog of probabilistic bounds on the backlog of the latter system. In Section III-A we go a step further and give a bound on the backlog that accommodates heterogeneously regulated inputs (Theorem 1), and which performs better than the bounds of Theorems 4 and 5 in [9]. Moreover, in Theorem 2, for a special case of leaky-bucket regulated inputs, we give three bounds on the complementary distribution of the backlog. A remarkable feature of the three bounds is that they are functions of some aggregate parameters of the leaky-bucket regulators.

In order to derive a probabilistic bound on the delay of a packet through a node we need an upper bound on the complementary distribution of the backlog as seen at packet arrival epochs. In Section III-B we find an inequality between the complementary distribution of the backlog seen at arrival epochs and the steady-state complementary distribution of the backlog (as seen at a randomly chosen point). In fact, we prove a more general result in Theorem 3, and then specialize the result to the complementary distribution in Corollary 1.

When evaluating the performance of statistical multiplexing, a common performance metric considered is the probability that the backlog exceeds a given level. However, a performance metric of practical relevance is the loss ratio, and in particular, longrun loss ratio (fraction of the data lost over a long time interval). In Theorem 4 (Section III-C) we give an exact upper bound on long-run loss rate and loss ratio for a service curve node. The bound is in terms of the complementary distribution of the backlog, a deterministic bound on the loss ratio [11], [10], and the aggregate arrival curve. We study the many sources limit in Section III-D; we identify the many sources limit and Bahadur-Rao improvement of the backlog bound in Theorem 1; we also discuss typical time-scale to overflow with leaky-bucket regulated inputs.

In Section IV we apply our findings to EF. We show how to obtain a probabilistic delay bound, based on the delay from backlog bound of PSRG nodes (Proposition 1). Then we apply a majorization by fresh traffic in order to find bounds at any node inside a network. Finally, we briefly address computation of an upper bound on the complementary distribution of the end-toend delay through a sequence of nodes.

Section V shows some numerical computations. We conclude the paper in Section VI. Proofs of the theorems are deferred to Appendix.

II. RELATED WORK

One approach to study EF is to derive deterministic bounds; this is pursued by Charny and Le Boudec in [12] and Bennett, Benson, Charny, Courtney, and Le Boudec in [3]. A worst-case bound on delay jitter for leaky-bucket regulated EF input flows [12] is sup-linear in the maximum hop count, and it explodes at certain utilization that may be rather low. Thus, the deterministic approach gives us hard QoS guarantees that may be quite pessimistic estimate of the performance. This leads us to seek for probabilistic guarantees.

An alternative probabilistic approach is proposed by Bonald, Proutière, and Roberts [7]. Their approach relies on two main assumptions. First, EF traffic at the network ingress is Better than Poisson meaning that the virtual waiting time distribution for EF input traffic to a single node is stochastically smaller than if the input is replaced with a Poisson process with the same intensity as the original input. Second, it uses a conjecture that delay jitter remains negligible, which would ensure, if EF traffic is Better than Poisson at the network ingress, it remains so as the EF traffic passes through a sequence of nodes in the network. A remarkable property of Better than Poisson approach is that it is parsimonious in the parameters needed to characterize the input traffic; it requires only the intensity of the aggregate input. However, it is necessary to shape individual flows in some way to ensure they are Better than Poisson at the network ingress. It must also be noted that [7] only presents plausibility arguments for the negligible jitter conjecture which has not yet been proved. Our approach does not make such assumptions, and our analysis is *exact* under the given set of assumptions. In addition, [7] assumes that a node offers a static non-preemptive priority for EF traffic over non-EF traffic. Our results are valid for a node that offers a service curve, and thus apply to a PSRG node as discussed earlier.

In our prior work [9] we derive probabilistic upper bounds on the backlog for a node that offers a service curve to the aggregate of independent individually regulated flows. The catalog of bounds given there consists of two sets of the bounds. The first set of bounds is derived upon the virtual segregation of the backlog to individual input flows, and then observing that such virtual backlogs are with bounded support we applied Hoeffding's inequalities [13] to obtain closed-form bounds for both homogeneously and heterogeneously regulated inputs. It turns out that the bound for homogeneously regulated inputs generalizes a result by Kesidis and Konstantopoulos [14], [15], which is for a work-conserving constant service rate server. The second set of bounds is derived upon an approach originally due to Chang, Song, and Chiu [8] for a work-conserving constant service rate server. Our extension is to a super-additive service curve. Moreover, we derive bounds that hold exactly in continuous time and improve upon the bound in [8]. In the present paper, we use the second set of bounds since the bounds of the second set exhibit superior tightness than the bounds of the first set [9], [8]. Unlike the related work [14], [8], [9] gives closed-form bounds for heterogeneously regulated inputs.

III. THEORETICAL FOUNDATIONS

We introduce some further notation. For the input aggregate A(t), consisting of *I* flows, we write $A(t) = \sum_{i=1}^{I} A_i(t)$. Also, the aggregate arrival curve is denoted as $\alpha(t) = \sum_{i=1}^{I} \alpha_i(t)$, and the upper bound on the aggregate sustainable rate as $\rho =$

 $\sum_{i=1}^{I} \rho_i. \text{ Define } \lambda_{\rho}(t) = \rho t \mathbb{1}\{t \ge 0\}. \text{ Define also}$ $\hat{\beta}(t+u) - \beta(t)$

$$\hat{\beta} = \sup_{t,u \ge 0} \frac{\beta(t+u) - \beta(t)}{u}$$

Notice $\hat{\beta}$ is maximum slope of β ; for a rate-latency service curve $\beta(t) = c(t-e)^+, \hat{\beta} = c.$

For two functions f, g, define $v(f, g) = \sup_{u \ge 0} \{f(u) - g(u)\}$. For example, $v(\alpha, \beta)$ is the required buffer size to ensure loss-free operation. We use the operators $x \lor y = \max(x, y)$ and $x \land y = \min(x, y)$.

Let Q(t) be the backlog at time t of a node that offers a service curve β . For a super-additive β , we know $Q(t) \leq \tilde{Q}(t)$, for any t (Lemma 1 [9]), where

$$\tilde{Q}(t) = \sup_{t-\tau \le s \le t} \{A(t) - A(s) - \beta(t-s)\}.$$
 (2)

and $\tau = \inf\{u \ge 0 | \alpha(u) \le \beta(u)\}.$

We say that β is super-additive if $\beta(t + s) \ge \beta(t) + \beta(s)$, for all $t, s \ge 0$. Many service curves are super-additive, but not all. A sufficient condition is that β is convex; in particular, the rate-latency service curve is super-additive, thus this additional assumption is not restrictive for our application to EF.

Note that τ is the intersection of the aggregate arrival curve α and the service curve β . Intuitively, think of τ as an upper bound on the busy period. This is formally correct if β is a *strict* service curve β (the service curve β is strict if in addition to (1) the backlog Q(s) = 0, for s given in (1)).

We recall two bounds from [9] that are for heterogeneously regulated flows. First, the bound of Theorem 4 [9]

$$\mathbb{P}(Q(t) > b) \le \sum_{k=0}^{K-1} \exp\left(-\frac{2[(b+\beta(s_k)-\rho s_{k+1})^+]^2}{\sum_{i=1}^{I} \alpha_i(s_{k+1})^2}\right),\tag{3}$$

for any $K \in \mathbb{N}$, and any $0 = s_0 \le s_1 \le \cdots \le s_K = \tau$. Second, the bound of Theorem 5 [9]

$$\mathbb{P}(Q(t) > b) \le \sum_{k=0}^{K-1} \exp\left(-\frac{[(b+\beta(s_k)-\rho s_{k+1})^+]^2}{2\sum_{i=1}^{I} v(\alpha_i, \lambda_{\rho_i})^2}\right),$$
(4)

for any $K \in \mathbb{N}$, and any $0 = s_0 \leq s_1 \leq \cdots \leq s_K = \tau$.

In the next section we give a backlog bound that improves upon both (3) and (4).

A. An Improved Backlog Bound

Theorem 1: (A Bound on Backlog) Consider a node that offers a super-additive service curve β . Then, under (A1)-(A4) and $\rho < \hat{\beta}$, for any t,

$$\mathbb{P}(Q(t) > b) \leq \\
\leq \sum_{k=0}^{K-1} \exp\left(-\frac{2[(b+\beta(s_k)-\rho s_{k+1})^+]^2}{[\sum_{i=1}^{I} \alpha_i(s_{k+1})^2] \wedge [4\sum_{i=1}^{I} v(\alpha_i,\lambda_{\rho_i})^2]}\right),$$
(5)

for any $K \in \mathbb{N}$, and any $0 = s_0 \leq s_1 \leq \cdots \leq s_K = \tau$. *Proof:* Appendix I.

Note that (5) and other bounds in [9] satisfy the *economy of* scale, a notion originally introduced by Botvich and Duffield

[16]. It means that if we scale b and β as O(I), then the probability to overflow decays exponentially with I. We also note that with fixed aggregate arrival curve, the bound in (5) is tightest for all the inputs having identical arrival curves $\alpha_i(t) = \frac{\alpha(t)}{I}$. We call this *the economy of equality*; it tells us that the best performance is achieved for all the input flows having the same arrival curves.

Next we give three bounds on the probability to overflow for leaky-bucket regulated inputs. The bounds require some *aggregate knowledge* about the leaky-buckets. As such, they merit is when knowledge about individual leaky-bucket parameters is not available, but some knowledge about the aggregate is given; for instance, an upper bound on the aggregate load or aggregate burstiness parameters. This is inline with DiffServ philosophy.

Theorem 2: (Three Backlog Bounds for Leaky-Bucket Regulated Inputs) Consider a node that offers a super-additive service curve β , fed with leaky-bucket regulated inputs; $\alpha_i(t) = \rho_i t + \sigma_i$. Then, under (A1)-(A4) and $\rho < \hat{\beta}$, for any t,

$$\mathbb{P}(Q(t) > b) \leq \\
\stackrel{(a)}{\leq} \sum_{k=0}^{K-1} \exp\left(-\frac{2[(b+\beta(s_{k})-\rho s_{k+1})^{+}]^{2}}{[\sum_{i=1}^{I}(\rho_{i}s_{k+1}+\sigma_{i})^{2}]\wedge(4\sum_{i=1}^{I}\sigma_{i}^{2})}\right) \\
\stackrel{(b)}{\leq} \sum_{k=0}^{K-1} \exp\left(-\frac{2[(b+\beta(s_{k})-\rho s_{k+1})^{+}]^{2}}{(\sqrt{\sum_{i=1}^{I}\rho_{i}^{2}s_{k+1}}+\sqrt{\sum_{i=1}^{I}\sigma_{i}^{2}})^{2}\wedge(4\sum_{i=1}^{I}\sigma_{i}^{2})}\right) \\
\stackrel{(c)}{\leq} \sum_{k=0}^{K-1} \exp\left(-\frac{2[(b+\beta(s_{k})-\rho s_{k+1})^{+}]^{2}}{(\rho s_{k+1}+\sqrt{\sum_{i=1}^{I}\sigma_{i}^{2}})^{2}\wedge(4\sum_{i=1}^{I}\sigma_{i}^{2})}\right), \\
\text{For any } K \in \mathbb{N} \text{ and any } 0 = s_{0} \leq s_{1} \leq \dots \leq s_{K} = T$$

for any $K \in \mathbb{N}$, and any $0 = s_0 \leq s_1 \leq \cdots \leq s_K = \tau$. *Proof:* Appendix II.

Let $\underline{\rho} = [\rho_1, \dots, \rho_I]$ and $\underline{\sigma} = [\sigma_1, \dots, \sigma_I]$ be vectors of the sustainable rates and burstiness parameters, respectively. Consider the following aggregate parameters: (P1) $\sum_{i=1}^{I} \rho_i$, (P2) $\sum_{i=1}^{I} \sigma_i^2$, (P3) $\sum_{i=1}^{I} \rho_i^2$, and (P4) $\sum_{i=1}^{I} \rho_i \sigma_i$. The parameters can be interpreted, respectively, as the aggregate input load, the variability of $\underline{\sigma}$, the variability of $\underline{\rho}$, and the correlation of $\underline{\rho}$ and $\underline{\sigma}$.

Notice that (6) require to know upper bounds on: (a) (P1)-(P4), (b) (P1)-(P3), and (c) (P1)-(P2). It is a remarkable property that (c) in (6) requires only two aggregate parameters, namely, (P1) and (P2). An issue of interest is how much we loose in terms of tightness as we know fewer aggregate parameters. We explore this numerically in Section V.

B. Bound on Backlog at Arrival Epochs

In the previous section we consider the steady-state complementary distribution of the backlog. This may be empirically interpreted as a fraction of time the backlog is above a given level (time average). Here we consider the complementary distribution of the backlog as seen by the packet arrivals, which may be empirically interpreted as a fraction of the arrival data that encounter the backlog above a given level (Palm average [17]). We denote this as \mathbb{P}_A for the arrival process A (\mathbb{E}_A is the expectation with respect to \mathbb{P}_A).

Theorem 3: (Bound on Backlog Seen By Arrivals) Consider a node that offers a service curve β . Suppose the input A is with stationary increments and intensity $\bar{\rho} < \hat{\beta}$. Then, for any run loss ratio is measurable function $\psi : \mathbb{R}^+ \to \mathbb{R}^+$,

$$\mathbb{E}_{A}[\psi(\tilde{Q}(0))] \leq \frac{\hat{\beta}}{\bar{\rho}} \mathbb{E}[\psi(\tilde{Q}(0))].$$
(7)

Proof: Appendix III.

Use in (7) $\psi(x) = \mathbf{1}_{(b,\infty)}(x)$, then we directly obtain the following corollary.

Corollary 1: Under the assumptions of the theorem, it holds

$$\mathbb{P}_A(\tilde{Q}(0) > b) \le \frac{\hat{\beta}}{\bar{\rho}} \mathbb{P}(\tilde{Q}(0) > b).$$

Konstantopoulos and Last [18] study a work-conserving constant service rate server. They prove equality in (7) for any measurable function $\psi : \mathbb{R}^+ \to \mathbb{R}$. This equality is known as distributional version of Little's law [18], [19].

The corollary tells us that for the majorizing process \hat{Q} , the complementary distribution of the backlog at arrival epochs is less than or equal to the steady-state complementary distribution of the backlog, times the ratio of the maximum slope of β on $[0, \tau]$ and the intensity of the arrival process $\bar{\rho}$. For the ratelatency service curve $\beta(s) = c(s - e)^+$, it reads as $\mathbb{P}_A(\tilde{Q}(0) > b) \leq \frac{c}{\bar{\rho}}\mathbb{P}(\tilde{Q}(0) > b)$.

Note that the result in Corollary 1 is established for a majorizing process $\tilde{Q}(\cdot)$ to the backlog $Q(\cdot)$. As such, it enables us to state:

$$\mathbb{P}_A(Q(0) > b) \le \frac{\beta}{\bar{\rho}} \mathbb{P}(\tilde{Q}(0) > b).$$

C. Bound on Long-Run Loss Ratio

Insofar, we consider upper bounds on the complementary distribution of the backlog. However, in practice, one would be more interested in an upper bound on the loss ratio. Consider a lossy node with arrival process A(t). Let L(t) be the data lost in [0, t]. Then the loss ratio is defined as $l(t) = \frac{L(t)}{A(t)}$. Rather than looking at l(t), we consider the long-run loss ratio $\bar{l} = \lim_{t \to \infty} l(t)$; thus, the fraction of packets lost over a long time interval.

The next theorem gives us an exact bound on the long-run loss ratio \bar{l} . The bound is in terms of the complementary distribution of \tilde{Q} (2), which is a majorization on the backlog of a virtual system identical to the one we consider here, but with sufficient buffer size to insure no losses. Note that having identified an upper bound on the backlog of a loss-free service curve node, the theorem directly gives us an upper bound on the long-run loss ratio of a lossy service curve node.

Theorem 4: (Bound on Long-run Loss Ratio) Consider a node that offers a service curve β and finite buffer of size B. Time is discrete. Then, an upper bound on the long-run loss rate is

$$\mathbb{E}[L(t) - L(t-1)] \leq \int_{B}^{v(\alpha,\beta)} \mathbb{P}(\tilde{Q}(t) > z) dz$$

$$\leq (v(\alpha,\beta) - B) \mathbb{P}(\tilde{Q}(t) > B).$$
(8)

Moreover, for ergodic inputs with stationary increments, and the intensity of the aggregate input $\bar{\rho}$, an upper bound on the long-

$$\bar{l} \leq \frac{1}{\bar{\rho}} \int_{B}^{v(\alpha,\beta)} \mathbb{P}(\tilde{Q}(0) > z) dz
\leq \frac{v(\alpha,\beta) - B}{\bar{\rho}} \mathbb{P}(\tilde{Q}(0) > B).$$
(9)

Proof: Appendix IV.

A similar expression to (8) was obtained by Likhanov and Mazumdar [20] for a work-conserving constant rate server. Their result is different in that it is for the limit of many independent identically distributed input flows. We show the bounds on long-run loss rate and loss ratio hold exactly (not only for the many sources limit). Our bounds are for a service curve node, which encompasses a work-conserving constant rate server.

We comment the bound on the long-run loss rate in (8). The first bound in (8) reads as: $\mathbb{E}[L(t) - L(t-1)] \leq \mathbb{E}[(\tilde{Q}(t) - B)^+]$ (see Appendix IV). Notice that \tilde{Q} takes its values on $[0, v(\alpha, \beta)]$. This allows us to use $v(\alpha, \beta)$ as the upper boundary of the integral in (8). The second bound in (8) can be easily derived directly by observing, for any t such that Q(t) > 0, $L(t) - L(t-1) \leq \tilde{Q}(t) - Q(t)$ (Appendix IV). Now note that $L(t) - L(t-1) \geq 0$ implies Q(t) = B, and thus $\mathbb{E}[L(t) - L(t-1)] \leq \mathbb{E}[(\tilde{Q}(t) - B)\mathbf{1}_{Q(t)=B}]$. Finally, use $\tilde{Q}(t) - B \leq v(\alpha, \beta) - B$, and easily proven fact that $\{Q(t) = B\} \subseteq \{\tilde{Q}(t) \geq B\}$.

Lastly, we compare with a known deterministic upper bound on the loss ratio [11], [10] over a time interval [0, t], given that Q(0) = 0,

$$\hat{l}(t) = \left[1 - \inf_{0 < s \le t} \frac{B + \beta(s)}{\alpha(s)}\right]^{-1}$$

This gives us an upper bound on the long-run loss ratio $\hat{l} = \lim_{t\to\infty} \hat{l}(t)$; we expect this to be a conservative bound as we exemplify in the following example.

Example 1: Consider a lossy node that offers the rate-latency service curve $\beta(t) = c(t - e)^+$. The node is fed with leaky-bucket regulated inputs $\alpha_i(t) = \rho_i t + \sigma_i$. Let $\sum_{i=1}^{I} \rho_i = \rho$ and $\sum_{i=1}^{I} \sigma_i = \sigma$. Then,

$$\hat{l}(t) = \begin{cases} \left[1 - \frac{B}{\rho t + \sigma}\right]^+, & t \le e \\ 1 - \frac{B}{v(\alpha, \beta)}, & t > e, \end{cases}$$
(10)

where $v(\alpha, \beta) = \rho e + \sigma$. Notice \hat{l} is linear in *B*.

D. Many Sources Limit

In the preceding sections, and [9], the bounds derived are *exact*. The bounds hold exactly for any setting of the parameters, and, in particular, the bounds are valid for any number of the input flows. In this section we consider asymptotic counterparts to the backlog bounds given earlier. In particular, we study the many sources limit – the buffer size and capacity scale as O(I) as the number of the input flows I tends to infinity.

We will see that in the asymptotic regime our bounds admit a simpler form. Recall that our backlog bounds are obtained by



Fig. 1. Typical time-scale to overflow versus the buffer size *b*. The node offers the rate-latency service curve with rate c = 150 Mbps and latency e=MTU/c; MTU=1500 Bytes. Individual input flows are regulated with identical leaky-bucket regulators to rate $\alpha c/I$ and burstiness 5 MTU; there are I = 100 input flows. In the brackets we give values of $\tau_c \wedge \tau$ as a fraction of τ .

using the union bound on

$$\mathbb{P}(\bigcup_{k \in \{0, \dots, K-1\}} \{A(0) - A(s_{k+1}) > b + \beta(s_k)\}).$$

One reason to consider the many sources limit is to show that by using the union bound we have asymptotically a tight bound. Another reason is to gain insight how the bounds behave, in particular, what is the most likely way the backlog build up.

From [9] and a simple majorization by using (A3) we have

$$\mathbb{P}(Q(0) > b) \leq \sum_{k=0}^{K-1} \mathbb{P}(A(0) - A(-s_{k+1}) > b + \beta(s_k)) \\
\leq \sum_{k=0}^{K-1} \mathbb{P}(A(0) - A(-s_k) > b + \beta(s_k) - \alpha(s_{k+1} - s_k)) \\
\leq \sum_{k=0}^{K-1} \mathbb{P}(A(0) - A(-s_k) > b + \beta(s_k) - \alpha(\delta)),$$
(11)

for any $K \in \mathbb{N}$, and any $0 = s_0 \le s_1 \le \cdots \le s_K = \tau$, where $\delta = \max_{k \in \{0, 1, \dots, K-1\}} \{s_{k+1} - s_k\}.$

Let $e^{-g(s_k)}$ be upper bound on the k-th summation term in the last inequality of (11); g is some positive-valued function (see [9] for a catalog of functions g). Suppose $0 = s_0 \leq s_1 \leq \cdots \leq s_K = \tau$ is such that there exists a unique $k^* \in \{0, 1, \ldots, K-1\}$ and $\epsilon > 0$ such that $g(s_{k^*}) + I\epsilon \leq g(s_k)$, for all $k \in \{0, 1, \ldots, K-1\} \setminus k^*$.

For the many sources scaling, the known functions g [9] satisfy $g(\cdot) \sim O(I)$. Thus

$$\mathbb{P}(Q(0) > b) \le e^{-g(s^*)} [1 + O(e^{-I\epsilon})], \tag{12}$$

as $I \to \infty$, where $s^* \in [0, \tau]$ such that $g(s^*) = \inf_{s \in [0, \tau]} g(s)$. Note that (12) does not require s^* to be unique on $[0, \tau]$; however, the partition $0 = s_0 \le s_1 \le \cdots \le s_K = \tau$ needs to ensure a unique minimum of $\{g(s_0), g(s_1), \ldots, g(s_{K-1})\}$. If s^* is unique on $[0, \tau]$, then it may be interpreted as the typical time-scale to overflow a given level of the buffer. This is a known concept; see e.g. [20], [8].

Note that (12) holds for any $K \in \mathbb{N}$. We can take a uniform partition of $[0, \tau]$ such that $\delta = \tau/K$, and then let $K \to \infty$. This allows us to replace $\alpha(\delta)$ in (13) with $\lim_{\delta \downarrow 0} \alpha(\delta)$. For a right-hand continuous α at 0, we replace $\alpha(\delta)$ with $\alpha(0)$; if in addition $\alpha(0) = 0$, the term $\alpha(\delta)$ in (13) vanishes. In practice, there always exists a peak rate constraint, and thus indeed $\lim_{\delta \downarrow 0} \alpha(\delta) = 0$. In the further expressions we carry on $\alpha(\delta)$, but we keep in mind the observation made here.

Next we consider (5) and identify the function g as

$$g(s) = \frac{2[(b+\beta(s)-\rho s-\alpha(\delta))^+]^2}{\left[\sum_{i=1}^{I} \alpha_i(s)^2\right] \wedge \left[4\sum_{i=1}^{I} v(\alpha_i, \lambda_{\rho_i})^2\right]}.$$
 (13)

In particular, for a node that offers the rate-latency service curve $\beta(t) = c(t-e)^+$, fed with leaky-bucket regulated inputs, we can show that s^* is unique and is equal to:

$$s^* = \begin{cases} (u \wedge \tau_c \wedge \tau) \lor e, & b < b^* \\ \tau, & b \ge b^* \end{cases} .$$
(14)

Where

$$u = \frac{\frac{b-ce-\alpha(\delta)}{c-\rho} \sum_{i=1}^{I} \rho_i \sigma_i - \sum_{i=1}^{I} \sigma_i^2}{\sum_{i=1}^{I} \rho_i \sigma_i - \frac{b-ce-\alpha(\delta)}{c-\rho} \sum_{i=1}^{I} \rho_i^2},$$

$$\tau_c = \frac{\sum_{i=1}^{I} \rho_i \sigma_i}{\sum_{i=1}^{I} \rho_i^2} \left(\sqrt{1 + 3 \frac{\sum_{i=1}^{I} \rho_i^2 \sum_{i=1}^{I} \sigma_i^2}{(\sum_{i=1}^{I} \rho_i \sigma_i)^2}} - 1 \right),$$

and

$$b^* = (c - \rho) \frac{\sum_{i=1}^{I} \rho_i \sigma_i}{\sum_{i=1}^{I} \rho_i^2} + ce + \alpha(\delta)$$

Note that τ_c is a value of s in (13) at which the two terms acting in the minimum operator are equal. Also, note that b^* is a cut-off buffer level at which the typical time-scale to overflow turns from one value to another.

In Fig. 1 we show some numerical values of s^* versus the buffer level *b*. For the utilization α larger than 0.5, we observe that s^* is equal to *e* for $b < b^*$ and otherwise to $\tau_c \wedge \tau$ for $b \ge b^*$. For the utilization α smaller than 0.5, s^* is equal to *e* for all the values of interest of the buffer level *b*. At this point, we compare with the typical time-scale to overflow of the bounds (3) and (4). It can be shown that for (3), $s^* = e$ for $b < b^*$, and else $s^* = \tau$. For (4), $s^* = e$. It is noteworthy that for (3) with buffer levels larger than b^* , $s^* = \tau$, which may be quite large, in particular, for high utilization. This explains some numerical results given later in Section V. Our improved bound (5) remedies the latter effect by having the typical time-scale to overflow $s^* = \tau_c \wedge \tau$, for $b > b^*$. In Fig. 1 we show $\tau_c \wedge \tau$ as a fraction of τ .

In the sequel, we consider Bahadur-Rao [21] improvement of the many sources limit (12)

$$\mathbb{P}(Q(0) > b) \approx \frac{1}{\sqrt{4\pi g(s^*)}} e^{-g(s^*)},$$
(15)

where

$$g(s^*) = -\frac{2[(b + \beta(s^*) - \rho s^* - \alpha(\delta))^+]^2}{\sum_{i=1}^{I} \alpha_i(s^*)^2}$$

In [20], Likhanov and Mazumdar show that for independent identically distributed input flows (15) is an exact asymptotics up to a multiplicative constant 1 + O(1/I). Their result is under two assumptions: (1) s^* is unique, and (2) $\liminf_{t\to\infty} g(t)/\ln t > 0$. It can be shown that in our case (1) holds, and (2) is not needed given that we have a finite summation in (11). On the other hand, for heterogeneously regulated inputs, one may use the central limit approximation as discussed in [22] (Section 5.4). Notice the pre-factor in (15) scales as $1/\sqrt{I}$, which was already observed elsewhere, e.g. by Montgomery and de Veciana [23]. We come back to the Bahadur-Rao improvement in Section V with some numerical computations.

IV. APPLICATION TO EF

A. Delay from Backlog for a PSRG Node

In general, for an arbitrary node, one cannot directly deduce a bound on the complementary distribution of the delay from the complementary distribution of the backlog. However, this is possible for a PSRG node. It is shown that the delay from backlog bound holds for PSRG FIFO nodes [3], and also for non-FIFO PSRG in [24].

Proposition 1: (Bound on Delay) For a PSRG node with rate *c* and latency *e*, it holds

$$\mathbb{P}(d(0) > u) \le \mathbb{P}_A(Q(0) > c(u-e)), \text{ for } u \ge e, \qquad (16)$$

where d(0) is a delay incurred by an arbitrary packet that arrives at time 0.

Proof: By Theorem 1 in [3] and Theorem III.1 in [24], the delay for a packet arriving at time t is bounded by Q(t)/c + e; simply use this majorization to obtain (16).

Notice, combining (16) with Corollary 1, and any upper bound on the steady-state complementary distribution of the backlog, we obtain an upper bound on the complementary distribution of the delay.

B. Majorization by the Fresh Traffic

Our bounds in [9] and (5), and typically the bounds found elsewhere, are based on the assumption that the input flows are mutually independent. Thus, we cannot apply the bounds directly, because it is not realistic to suppose the input flows to any node in the network are mutually independent; the flows may get correlated as they share common upstream nodes. However, it is reasonable to suppose at the network ingress the flows are mutually independent; (A2) in Section I. We suppose the delay jitter incurred at the upstream nodes to a given node is bounded by Δ . Such a bound indeed exists with finite buffer sizes; use the delay from backlog bound of a PSRG node [3], [24] to obtain $\Delta = (h - 1) \max\{B_n/c_n + e_n\}$, where h is the maximum hop count, B_n is the buffer size, c_n the service rate, and e_n the latency of the node n. Then, we majorize increments of the input flows to a given node by the fresh traffic at the network ingress

$$A_i(t) - A_i(s) \le A_i^0(t) - A_i^0(s - \Delta).$$

Such a majorization was suggested by Chang, Chiu, and Song [8]. In particular, for our bounds in (5) and (6) this amounts to replace s_{k+1} with $s_{k+1} + \Delta$.

We note that one can easily generalize our bounds to nonindependent input flows by using Hölder's inequality in our application of Chernoff-Hoeffding's inequalities. Then, it can be shown that all our bounds on the backlog remain the same, but with the exponent divided by I; this would preclude the statistical multiplexing gain.

C. Delay Through a Sequence of Nodes

Let d_n be the delay of an arbitrary packet through a node n. Suppose the packet traverses h nodes. Then, the end-to-end delay incurred by a packet is

$$d = d_1 + d_2 + \dots + d_h.$$
(17)

In Section IV-A we show how to obtain per-node probabilistic bound on the delay, i.e. $\mathbb{P}(d_n > u) \leq F_n(u)$. Here we consider how to obtain $\mathbb{P}(d > u) \leq G(u)$, where G(u) is an upper bound on the complementary distribution of the end-to-end delay.

One approach gives us

$$G(u) = \sum_{n=1}^{h} F_n(\frac{u}{h}).$$
(18)

This is readily shown by noting $d \leq h \max\{d_1, \ldots, d_h\}$, and then

$$\mathbb{P}(d > u) \leq \mathbb{P}(\max\{d_1, \dots, d_h\} > \frac{h}{u}) \\ \leq \mathbb{P}(\bigcup_{n \in \{1, \dots, h\}} \{d_n > \frac{u}{h}\}) \\ \leq \sum_{n=1}^{h} \mathbb{P}(d_n > \frac{u}{h}) \leq \sum_{n=1}^{h} F_n(\frac{u}{h}).$$

Note (18) is obtained by summing up h times the maximum delay along the path, as such it may be a conservative bound on the end-to-end delay. Notice, also, (18) is sup-linear in the number of hops h.

One may say more on the end-to-end delay by assuming independence of the delays incurred at different hops; this would be an approximation. We defer a more elaborate study of the end-to-end delay for future work.

V. NUMERICAL RESULTS

In this section we give a numerical comparison of the backlog bounds. We do not show numerical results for our bound on the backlog as seen by packet arrivals (Theorem 3) and longrun loss ratio (Theorem 4). The latter two bounds are in fact expressed in terms of the bound on the complementary distribution of the steady-state backlog. We consider some aspects of our analytical results through numerical computations; we do not compare with simulations results.

We first numerically compare the bounds (a), (b), and (c) in (6) and the bounds in (3) and (4). We refer to the bounds in (6)



HETEROGENEOUSLY REGULATED INPUT FLOWS



Fig. 2. Backlog bounds HET(a), HET(b), HET(c), (3) and (4). The input aggregate consists of two traffic classes; class-1 individual flows are (ρ_1, σ_1) leaky-bucket regulated (resp. class-2 with (ρ_2, σ_2)). Two upper graphs are for homogeneous case $\rho_1/\rho_2 = 1$ and $\sigma_1/\sigma_2 = 1$. Two lower graphs are for heterogeneous case $\rho_1\rho_2 = 1/4$, $\sigma_1 = 2$ MTU, and $\sigma_2 = 8$ MTU. The node offers the rate-latency service curve with rate c = 150 Mbps and latency e = MTU/c; MTU=1500 Bytes,.



Fig. 3. A comparison of backlog bounds HET(a) (-x- line), (3) (solid light line), and HOM (Theorem 3 in [6]) (solid bold line). All individual flows have identical leaky-bucket regulators; homogeneous case.

respectively as HET(a), HET(b), and HET(c). We suppose two traffic classes each consisting of 50 flows; thus I = 100. Class-1 flows are (ρ_1, σ_1) leaky-bucket regulated; respectively, class-2 flows are (ρ_2, σ_2) leaky-bucket regulated. The node offers the rate-latency service curve with rate c = 150 Mbps and latency e = MTU/c; MTU is set to 1500 Bytes. The results are given for the node utilizations $\alpha = 0.2$ and 0.8, which are representative of light and heavy loaded node, respectively. The bounds are computed as the infimum over uniform partition of $[0, \tau]$ ($s_k = k\tau/K$, for $k = 0, \ldots, K$).

In Fig. 2 (two upper graphs), we show the backlog bounds for the homogeneous case; $\rho_1 = \rho_2 = \alpha c/I$ and $\sigma_1 = \sigma_2 = 5$ MTU (ρ and σ non-correlated). In Fig. 2 (two lower graphs) we show the backlog bounds for the heterogeneous case; $\rho_1 = 0.4\alpha c/I$, $\rho_2 = 1.6\alpha c/I$, $\sigma_1 = 2$ MTU, $\sigma_2 = 8$ MTU (ρ and σ positively correlated). We make the following observations:

- in many cases (3) is better than (4); exception is a heavy loaded node;
- HET(a) is very close to (3) for light utilization;
- for heavy utilization, HET(a) remedies the deviation of (3); this is expected from our discussion in Section III-D;
- in all the cases, HET(a) and HET(b) almost coincide;

• HET(c) is close to HET(a) for light utilization; as the utilization increases, HET(c) gradually moves from (3) to (4); thus HET(c) may have some merit for a lightly loaded node (recall that HET(c) is parsimonious with respect to the aggregate parameters needed);

• the deviation of (4) for heavy utilization is indeed stronger for positively correlated ρ and $\underline{\sigma}$ and larger $\sum_{i=1}^{I} \sigma_i^2$.

• for light and moderate utilization, (3), HET(a), and HET(b) are insensitive to the node utilization (i.e. $\underline{\rho}$); this can be explained by considering the many sources limit – for light to moderate utilization the typical time-scale to overflow s^* is equal to e, thus very small, so all the terms that act in the bound multiplied with s^* do not have a significant impact (including $\underline{\rho}$); a dominant effect have the burstiness parameters, i.e. $\sum_{i=1}^{I} \sigma_i^2$.

Our next objective is to compare the backlog bound (3) and the backlog bound in Theorem 3 [9]; the latter bound is derived under assumption that input flows are regulated with identical regulators; we call this bound HOM. We know the latter bound is tighter than (3) [9]. We demonstrate when discrepancy between the two bounds is particularly emphasized; we give numerical results for the homogeneous case, i.e. identically regulated input flows. In Fig. 3 we show the backlog bound (3), HET(a), and HOM versus normalized buffer level. Our first observation is that, for light utilization, (3) and HET(a) are conservative with respect to HOM. However, for heavier load, discrepancy between HET(a) and HOM becomes less pronounced; the bounds get fairly close to each other, and even, for heavy load, HET(a) outperforms HOM. It is also noteworthy that, unlike (3), HOM is *not* insensitive with respect to ρ under light to moderate utilization.

By numerical results in Fig. 4 we compare our exact bound HET(a) with its many sources limit, and Bahadur-Rao improvement. The results are given for the homogeneous case with $\rho_i = \alpha c/I$, $\sigma_i = 5$ MTU, I = 100. As earlier, c = 150 Mbps, e = MTU/c, and MTU=1500 Bytes. One observation is that the Bahadur-Rao improvement is not dramatic; it is about an order of magnitude uniformly over the buffer size.

Last but not least, we compare our backlog bounds with bound of Better than Poisson [7] (discussed in Section II); see Fig. 5. We fix the aggregate arrival curve to $\alpha(t) = \rho t + \sigma$ with $\rho = \alpha c$ and $\sigma = 500$ MTU; MTU=1500 Bytes. We show the results for three different utilizations $\alpha = 0.2, 0.5, 0.8$. Number of the input flows is I = 100 and 500. We consider a node with c = 150 Mbps and e = 0. Notice that for I = 500, $\sigma_i = 1$ MTU, which in fact gives the burstiness constraint for a constant bit rate input flow. Observe also that by setting the latency of the node e = 0, the backlog bound of Better than Poisson corresponds to that of M/D/1 queue; we use the asymptotic expansion for M/D/1 [7] in our numerical computations.

As pointed out in Section 5, our bounds admit the economy of scale. Thus as we scale the buffer size and the service rate as O(I), the backlog bound decays exponentially with I. On the other hand, the backlog bound of Better than Poisson is invariant to the number of input flows; it depends solely on the load of the aggregate input ρ . The results in Fig. 5 confirm the economy of scale effect. We observe that for I = 500, the backlog



Fig. 4. Backlog bounds HET(a) (solid line), its many sources limit (thin solid line), and Bahadur-Rao improvement (dashed line); the homogeneous case with $\rho_i = \alpha c/I$, $\sigma_i = 5$ MTU, I = 100, c = 150 Mbps, e = MTU/c, MTU=1500 Bytes.

bound HOM (identically regulated input flows), is very close to the backlog bound of Better than Poisson. Thus relying on the economy of scale effect, we expect our bounds to favorably compete with Better than Poisson in terms of tightness as the number of input flows increases (high multiplexing). Observe that for I = 100 (moderate multiplexing of bursty flows) our bounds are more pessimistic than Better than Poisson. It has to be noted that Better than Poisson does not apply in this case; we consider bursty individual flows, while Better than Poisson assumes non bursty flows.

VI. CONCLUSION

We propose a framework to derive probabilistic guarantees for networks with individually regulated input flows and aggregate scheduling, in particular, Expedited Forwarding DiffServ networks. Our approach is based on assumption that a node can be abstracted as a service curve node; this is verified for the definition of EF PHB, namely, PSRG [1], [2]. A remarkable feature



Fig. 5. A comparison with Better than Poisson. The graphs show backlog bounds for $\beta(t) = ct$ node with utilization $\alpha = 0.2$ (upper graph) and 0.8 (lower graph); HET(a) is shown as solid line, HOM as dashed line, and Better than Poisson (M/D/1) [7] as dotted line.

of our approach is that the bounds we obtain are exact, they are valid for any setting of the parameters, and in particular for any number of the input flows.

ACKNOWLEDGMENT

The authors thank Thomas Bonald, Alexandre Proutière, and Jim Roberts for discussions on [7]. A discussion with Don Towsley is kindly acknowledged. The first author gratefully thank George Kesidis for hospitality and discussion during his visit to him.

REFERENCES

- A. Charny, F. Baker, J. Bennett, K. Benson, J.-Y. Le Boudec, A. Chiu, [1] W. Courtney, B. Davie, S. Davari, V. Firoiu, C. Kalmanek, K. K. Ramakrishnan, and D. Stiliadis, "Ef phb redefined," November 2000.
- B. Davie (Editor), A. Charny, F. Baker, J. Bennett, K. Benson, J.-Y. Le [2] Boudec, A. Chiu, W. Courtney, S. Davari, V. Firoiu, C. Kalmanek, K. K. Ramakrishnan, and D. Stiliadis, "An expedited forwarding phb," April 2001
- [3] Jon C. R. Bennett, Kent Benson, Anna Charny, William F. Courtney, and

Jean-Yves Le Boudec, "Delay jitter bounds and packet scale rate guarantee for expedited forwarding," in Proc. of IEEE INFOCOM'2001, March 2001

- [4] R. Agrawal, R. L. Cruz, C. Okino, and R. Rajan, "Performance bounds for flow control protocols," IEEE/ACM Transactions on Networking (7) 3, pp. 310-323, June 1999.
- [5] C.S. Chang, "On deterministic traffic regulation and service guarantee: A systematic approach by filtering," IEEE Transactions on Information Theory, vol. 44, pp. 1096-1107, August 1998.
- J.-Y. Le Boudec, "Application of network calculus to guaranteed service [6] networks," IEEE Transactions on Information Theory, vol. 44, pp. 1087-1096, May 1998.
- [7] T. Bonald, A. Proutiére, and J. Roberts, "Statistical performance guarantees for streaming flows using expedited forwarding," in Proc. of IEEE INFOCOM'2001, March 2001.
- [8] Cheng-Shang Chang, Wheyming Song, and Yuh ming Chiu, "On the performance of multiplexing independent regulated inputs," in Proc. of Sigmetrics 2001, Massachusetts, USA, May 2001.
- [9] Milan Vojnović and Jean-Yves Le Boudec, "Bounds for independent regulated inputs multiplexed in a service curve network element," in Proc. of Globecom 2001 (Internet Performance Symposium), San Antonio, Texas, USA, November 2001.
- [10] J.-Y. Le Boudec and P. Thiran, Network Calculus, Springer-Verlag, 2001.
- [11] Chen-Ming Chuang and Jin-Fu Chang, "Deterministic loss ratio quality of service guarantees for high speed networks," IEEE Communication Letters, vol. 4, no. 7, pp. 236-238, July 2000.
- [12] Anna Charny and Jean-Yves Le Boudec, "Delay bounds in a network with aggregate scheduling," in First International Workshop on Quality of future Internet Services, September 2000.
- [13] Wassily Hoeffding, "Probability inequalities for sums of bounded random variables," American Statistical Association Journal, pp. 13-30, March 1963
- [14] George Kesidis and Takis Konstantopoulos, "Extremal traffic and worstcase performance for queues with shaped arrivals," Analysis of Communication Networks: Call Centers, Traffic and Performance (edt. D. R. Mc-Donald and S. R. E. Turner), Fields Institute Communications/AMS, ISBN 0-8218-1991-7, 2000.
- [15] George Kesidis and Takis Konstantopoulos, "Worst-case performance of a buffer with independent shaped arrival processes," IEEE Communication Letters, 2000.
- [16] D. Botvich and N. Duffield, "Large deviations, the shape of the loss curve, and economies of scale in large multiplexers," Queueing Systems, vol. 20, pp. 293-320, 1995.
- [17] Francois Baccelli and Pierre Brémaud, Elements of Queueing Theory, vol. 26, Applications of Mathematics, Springer-Verlag, 1991.
- T. Konstantopoulos and G. Last, "On the dynamics and performance of [18] stochastic fluid systems," Journal of Applied Probability, vol. 37, pp. 652-667, 2000.
- [19] T. Konstantopoulos, M. Zazanis, and G. de Veciana, "Conservation laws and reflection mappings with an application to multiclass mean value analysis for stochastic fluid queues," Stoch. Proc. Appl., vol. 65, no. 1, pp. 139-146, 1997.
- [20] Nikolay Likhanov and Ravi R. Mazumdar, "Cell loss asymptotics in buffers fed with a large number of independent stationary sources," Journal of Applied Probability, vol. 36, pp. 86-96, 1999.
- [21] R. R. Bahadur and R. Ranga Rao, "On deviations of the sample mean," Ann. Math. Statist., vol. 31, pp. 1015–1027, 1960.J. W. Roberts (Editor), COST 224: Performance evaluation and design of
- [22] multiservice networks, Commision of the European Communities, 1991.
- [23] M. Montgomery and G. de Veciana, "On the relevance of time scales in performance oriented traffic characterization," in Infocom'96, 1996, pp. 4d.3.1-4.d.3.8.
- Jean-Yves Le Boudec and Anna Charny, "Packet scale rate guarantee for [24] non-fifo nodes," in Proc. of IEEE Infocom 2002, New York, USA, June 2002.

APPENDIX

I. PROOF OF THEOREM 1

From Lemma 2 [9], it holds

$$\mathbb{P}(Q(t) > b) \le \sum_{k=0}^{K-1} \mathbb{P}(A(t) - A(t - s_{k+1}) > b + \beta(s_k)),$$
(19)

for any $K \in \mathbb{N}$ and any $0 = s_0 \leq s_1 \leq \cdots \leq s_K = \tau$.

Next, by Hoeffding's inequality for non-uniformly bounded random variables $0 \le A_i(t) - A_i(s) \le \alpha_i(t-s)$, for any $s \le t$ (A3), we obtain $\mathbb{P}(A(t) - A(t-s_{k+1}) > b + \beta(s_k)) \le$

$$\leq \exp\left(-\frac{2[(b+\beta(s_k)-\rho s_{k+1})^+]^2}{\sum_{i=1}^{I} \alpha_i(s_{k+1})^2}\right),$$
 (20)

where in the nominator of the exponent we use (A4).

On the other hand, in [14] [9] we show $\mathbb{P}(A(t) - A(t - s_{k+1}) > b + \beta(s_k)) \le$

$$\leq \exp\left(-\frac{[(b+\beta(s_k)-\rho s_{k+1})^+]^2}{2\sum_{i=1}^{I} v(\alpha_i, \lambda_{\rho_i})^2}\right).$$
 (21)

Finally, indeed, the minimum of (20) and (21) is an upper bound on $\mathbb{P}(A(t) - A(t - s_{k+1}) > b + \beta(s_k))$. Using this minimum in (19) completes the proof.

II. PROOF OF THEOREM 2

The first inequality in (6) is a corollary of Theorem 1 for leaky-bucket regulated inputs. The second inequality is obtained by upper-bounding the first term in the minimum operation in (6) (a) as follows

$$\begin{split} &\sum_{i=1}^{I} (\rho_i s_{k+1} + \sigma_i)^2 = \\ &= \sum_{i=1}^{I} \rho_i^2 s_{k+1}^2 + 2 \sum_{i=1}^{I} \rho_i \sigma_i s_{k+1} + \sum_{i=1}^{I} \sigma_i^2 \\ &\leq \sum_{i=1}^{I} \rho_i^2 s_{k+1}^2 + 2 \sqrt{\sum_{i=1}^{I} \rho_i^2} \sqrt{\sum_{i=1}^{I} \sigma_i^2} s_{k+1} + \sum_{i=1}^{I} \sigma_i^2 \\ &= (\sqrt{\sum_{i=1}^{I} \rho_i^2} s_{k+1} + \sqrt{\sum_{i=1}^{I} \sigma_i^2})^2, \end{split}$$

where the inequation is by Cauchy-Schwartz's inequality.

The last inequality in (6) is by a trivial bound $\sum_{i=1}^{I} \rho_i^2 \leq \rho^2$. This completes the proof of the theorem.

III. PROOF OF THEOREM 3

Let $\tilde{A}^* = A \otimes \beta$. $A \otimes \beta$ is called the min-plus convolution of A and β , defined by $(A \otimes \beta)(t) = \inf_{u \in [0,t]} \{A(t-u) + \beta(u)\}$. By [8], [9], the infimum is obtained for $u \in [0, \tau]$, thus the majorizing process $\tilde{Q}(t)$ defined in Equation (2) satisfies $\tilde{Q}(t) = A(t) - \tilde{A}^*(t)$.

We now state and prove a preparatory lemma, and then continue with the proof of the theorem.

Lemma 1: We have

$$\tilde{A}^*(t+u) - \tilde{A}^*(t) \le u\hat{\beta}.$$

Proof: Define $\gamma(u) = u\hat{\beta}\mathbf{1}\{u \ge 0\}$. It follows from the definition of $\hat{\beta}$ that, for all $0 \le s \le t$:

$$\beta(t-s) + \gamma(s) \ge \beta(t),$$

thus

$$\beta \otimes \gamma \geq \beta$$
.

It follows that

$$\tilde{A}^* = A \otimes \beta < A \otimes (\beta \otimes \gamma) = (A \otimes \beta) \otimes \gamma = \tilde{A}^* \otimes \gamma.$$

Coming back to the definition of \otimes we find that

$$\tilde{A}^*(t+u) \le \tilde{A}^*(t) + u\hat{\beta}$$

For a wide-sense increasing measurable function φ such that $\varphi'=\psi$

$$\varphi(\tilde{Q}(t)) - \varphi(\tilde{Q}(0)) = \int_0^t \varphi'(\tilde{Q}(s))\tilde{Q}(ds), \qquad (22)$$

where $\tilde{Q}(ds) = A(ds) - \tilde{A}^*(ds)$. It follows from the lemma that

$$\int_0^t \psi(\tilde{Q}(s))\tilde{A}^*(ds) \le \hat{\beta} \int_0^t \psi(\tilde{Q}(s))ds.$$
(23)

Combining (22) with (23) we obtain

$$\varphi(\tilde{Q}(t)) - \varphi(\tilde{Q}(0)) \le \int_0^t \psi(\tilde{Q}(s)) A(ds) - \hat{\beta} \int_0^t \psi(\tilde{Q}(s)) ds$$

Take the expectation at both sides to obtain

$$0 \leq \bar{\rho} t \mathbb{E}_A[\varphi'(\tilde{Q}(0))] - \hat{\beta} t \mathbb{E}[\varphi'(\tilde{Q}(0))]$$

where the Palm expectation is by Campbell's formula [17]. Replacing φ' with ψ we prove (7).

IV. PROOF OF THEOREM 4

By the service curve property, for all t, there exists $s \le t$ such that $A^*(t) \ge A'(s) + \beta(t-s)$, where A'(s) = A(s) - L(s). Note $L(t) - L(t-1) = (L(t) - L(t-1))\mathbf{1}_{Q(t)=B}$. For s = t, $A'(s) + \beta(t-s) \ge A'(t)$, and thus Q(t) = 0. Since we are interested in the events $\{Q(t) = B\}$, we are allowed to only consider s < t. Then, for s < t,

$$L(t) - L(t - 1) =$$

$$= A(t) - A^{*}(t) - Q(t) - L(t - 1)$$

$$\leq A(t) - A(s) - \beta(t - s) - Q(t) + L(s) - L(t - 1)$$

$$\leq A(t) - A(s) - \beta(t - s) - Q(t).$$
(24)

And, thus

$$\begin{split} & (L(t) - L(t-1))\mathbf{1}_{Q(t)=B} \leq \\ & \leq (A(t) - A(s) - \beta(t-s) - Q(t))\mathbf{1}_{Q(t)=B} \\ & = (A(t) - A(s) - \beta(t-s) - B)\mathbf{1}_{Q(t)=B} \\ & \leq (A(t) - A(s) - \beta(t-s) - B) \vee 0 \\ & \leq (\max_{0 \leq s \leq t} \{A(t) - A(s) - \beta(t-s)\} - B) \vee 0 \\ & = (\tilde{Q}(t) - B) \vee 0. \end{split}$$

Thus, $\mathbb{P}(L(t) - L(t-1) > u) \leq \mathbb{P}(\tilde{Q}(t) \lor B > u + B) = \mathbb{P}(\tilde{Q}(t) > u + B)$, for $u \geq 0$. Then

$$\mathbb{E}[L(t) - L(t-1)] \le \int_0^{v(\alpha,\beta)-B} \mathbb{P}(\tilde{Q}(t) > u + B) du,$$

where the upper boundary in the integral is due to $0 \leq \tilde{Q}(t) \leq v(\alpha, \beta)$, for all t.

The bound on the long-run loss ratio (9) is immediately obtained by observing $\overline{l} = \mathbb{E}[L(t) - L(t-1)]/\mathbb{E}[A(t) - A(t-1)]$, and by definition $\overline{\rho} = \mathbb{E}[A(t) - A(t-1)]$. This completes the proof of the theorem.