Elements of Probabilistic Network Calculus for Packet Scale Rate Guarantee Nodes

Milan Vojnović and Jean-Yves Le Boudec
EPFL
CH-1015 Lausanne
Switzerland

Abstract

Packet Scale Rate Guarantee (PSRG) is a node model used by IETF for Expedited Forwarding, a priority service defined in the context of Differentiated Services [10]. We review probabilistic performance bounds that apply to PSRG nodes, and improve a previous bound for loss probability.

1 Introduction

Differentiated Services. Expedited Forwarding (EF) is a Per-Hop Behavior (PHB) of Differentiated Services (DiffServ), designed to offer a low-loss, low-latency, low-jitter service [10]. In an EF network, packet flows are shaped individually at the network ingress points, but are served as aggregates by the nodes in the network. The definition of EF PHB relies on Packet Scale Rate Guarantee (PSRG) [10] (defined later). Some sample-path properties of PSRG nodes are established, e.g., in [3, 6, 5]. In our prior work [23, 24] we obtained some probabilistic bounds on the performance that apply to PSRG nodes. It is of network engineering interest to have tight bounds on the performance of PSRG nodes, which would enable one to dimension the network such that some notion of quality-of-service is guaranteed.

Our approach is based on combining some deterministic results of network calculus (see, e.g., [5] and the references therein) with stochastic majorization such that we obtain bounds on performance that hold in probability. While accounting for statistical multiplexing gain, the probabilistic bounds are typically much tighter estimates of the actual performance than the worst-case deterministic bounds. We discuss some related probabilistic approaches later.

Abstract Node Models: PSRG, GR, Service Curves. Consider a node that serves packets and let $(T_n)_n$, $n \in \mathbb{Z}^+$, (resp. $(T'_n)_n$), be the sequence of packet arrival (resp. departure) times. Denote with L_n the size in bits of the packet that arrives at T_n.

A node is said to offer PSRG with rate r and latency e, if for all $n \in \mathbb{Z}^+$, $T'_n \leq V_n + e$; where $V_0 = 0$, and else

$$V_n = \max\left\{T_n, \min[V_{n-1}, T'_{n-1}]\right\} + \frac{L_n}{r},\ n > 0.$$

An equivalent definition [6] is that for all $j \leq n$

$$T'_n \leq \left\{e + T'_j + \frac{L_{j+1} + ... + L_n}{r}\right\} \vee \bigvee_{k=j+1}^{n} \left\{e + T_k + \frac{L_k + ... + L_n}{r}\right\}. \tag{1.1}$$

Another node model is Guaranteed Rate (GR) [12]. We say that the node is GR with rate r and latency e [12], if, similarly, $T'_n \leq V_n + e$ where $V_0 = 0$ but the recursion for V_n is replaced by

$$V_n = \max\left[T_n, V_{n-1}\right] + \frac{L_n}{r},\ n > 0. \tag{1.2}$$

An equivalent definition of GR is that for all n

$$T'_n \leq \bigvee_{k=1}^{n} \left\{ e + T_k + \frac{L_k + ... + L_n}{r} \right\}. \tag{1.3}$$

Both PSRG and GR capture how much a real node differs from a hypothetical minimum rate server: the GR model puts a bound on how much *later* the real node is, whereas PSRG puts a bound on how much *earlier or later* it is. The difference is important if a node multiplexes several flows into one single aggregate, thus PSRG is used in the context of aggregate scheduling, while GR is in the context of per-flow scheduling. It can easily be seen that PSRG is a stricter definition than GR, i.e., if a node is PSRG with rate r and latency e, then it is also GR with rate r and latency e (the converse is not true); see [3] for a detailed comparison of PSRG and GR.

We next recall the concept of service curves [21], and then discuss the relationship between PSRG and GR, on one hand, and service curves, on the other hand. We need some further notation. Let $R(t)$ be number of bits observed in $[0, t]$ at the node input. Likewise, $R^*(t)$ be the number of bits observed in $[0, t]$ at the node output.

A node is said to offer the service curve β (positive-valued, wide-sense increasing), if

$$R^* \geq R \otimes \beta. \tag{1.4}$$

where $\otimes$ is the (min,+) convolution: $R \otimes \beta(t) = \inf_{u \in [0,t]}[R(u) + \beta(t-u)]$.

A special case of service curve ("strict" service curve) is when the node guarantees that for an interval of length v that is contained in a busy period, the service offered by the node is at least $\beta(v)$. Note however that is not generally true; in contrast, it is a merit of this definition (like other node models mentioned in this paper) to apply to complex nodes, possibly non work conserving, where the concept of busy period is not relevant (consider for example a node made of a delay element followed by a server of a constant rate r, fed with a flow of constant rate $\epsilon \leq r$; the amount of service received during a busy period of duration t is ϵt which can be arbitrarily small).

Service curve and GR are, roughly speaking, equivalent, when the service curve is a rate-latency function (i.e., has the form $\beta(t) = r(t-e)^+$, with $(\cdot)^+ := \max[\cdot, 0]$). A more precise statement ([5], Chapter 2) is as follows: a FIFO node that offers a service curve $\beta(t) = r(t-e)^+$ is GR with rate r and latency e; conversely, a GR node with rate r and latency e offers the service curve $\beta(t) = r(t - e - \frac{L_{max}}{r})^+$, (even if it is non-FIFO), where L_{max} is the maximum packet size of the input. Also note that the definition of GR in (1.3) is the max-plus equivalent of the min-plus definition (1.4) of service curve.

Another node abstraction is adaptive guarantee [18, 1] (also called adaptive service curve [5], Chapter 7). A node is said to offer the adaptive service curve β if for all $t \geq 0$, and all $s \leq t$,

$$R^*(t) \geq [R^*(s) + \beta(t-s)] \wedge \inf_{u \in [s,t]} [R(u) + \beta(t-u)]. \tag{1.5}$$

The adaptive service curve is the min-plus equivalent of the max-plus definition of PSRG in (1.1), and there is the same type of relationship between PSRG and adaptive service curve ([5], Chapter 7): a FIFO node that offers an adaptive service curve $\beta(t) = r(t-e)^+$ is PSRG with rate r and latency e; conversely, a PSRG node with rate r and latency e offers the service curve $\beta(t) = r(t - e - \frac{L_{max}}{r})^+$. Similarly also, adaptive service curve is a stronger property than service curve.

Probabilistic Bounds. In this paper we focus on probabilistic performance bounds for an isolated PSRG or adaptive service curve node. We review our previous results on service curve nodes [23, 24]; from the discussion above, it follows that all the results obtained for service curve nodes also apply to PSRG nodes. Further, we give an improved loss bound that exploits the adaptive service curve or PSRG definition, and does not seem to be readily available for a service curve element. The study of probabilistic guarantees for a network of nodes with aggregate scheduling turns out to be a non-trivial problem, and we leave it aside here; see for example [24] for some indications.

In Section 2, we review the main lines of a bounding method, which has two phases. First, we exploit events that are expressed in terms of arrival flows only. Then we find bounds on the probabilities of these events using one of several Hoeffding's inequalities [13]. We explain how this method extends, and simplifies, previously known bounds. In Section 3, we explain in this framework our previous bounds for backlog, delay and loss, that are all based on the service curve property. We also give a new bound, which is an improved bound on loss that exploits the adaptive service curve property. We show some numerical results for the loss bound in Section 4.

Notation. Let $Q(t)$ denote the number of bits present in a node at time t (we call this backlog); also let $L(t)$ be the number of bits lost in $[0,t]$ at the node input, and $R'(t) = R(t) - L(t)$. For convenience, suppose $R(0) = 0$ and $Q(0) = 0$. We say R is α-smooth (or α is an arrival curve for R, or R conforms to α), if α is a positive-valued wide-sense increasing function such that $R \leq R \otimes \alpha$, i.e. for all $s \leq t$, $R(t) - R(s) \leq \alpha(t-s)$. If the input R is a superposition of I number of the flows, then with $R_i(t)$ we denote number of the bits observed in $[0,t]$ that belong to the flow i, $i = 1, 2, \ldots, I$. In the case of regulated flows, we call α_i an arrival curve for flow i. Next, for two functions α and β, let $v(\alpha, \beta)$ be their maximum vertical deviation. A classical result is that if a node offers the service curve β and the input is α-smooth, then $v(\alpha, \beta)$ is an upper bound on the backlog at any time.

2 A Bounding Method

Phase 1: Reduction to Arrival Events. All bounds we present in this paper use the following reduction, which was already found for example in [16, 26, 22, 7]. Events of the interest such as "steady-state backlog larger than a given value", "waiting time larger than a given value", "loss for a lossy node larger than a given value", are shown to be subsets of inclusions of the following events; for $s \le t$ that belong to some countable set:

$$\left\{ \sum_{i=1}^{I} R_i(t) - R_i(s) \ge g(t-s) \right\}, \tag{2.6}$$

where g is some appropriately defined positive-valued non-decreasing function. The details are given in Section 3.

Phase 2: Bounding the Probability of Arrival Events (2.6). A classical approach that was used in the past consisted in posing some a-priori bounds for the probability of (2.6): Kurose [16] assumes that, for each fixed $s \le t$, and $1 \le i \le I$, $R_i(t) - R_i(s)$ is stochastically majorized by some known random variable; Yaron and Sidi [26] suppose that the probability of (2.6) is upper bounded by a time-invariant exponential and then build calculus under this assumption; this is later extended by Starobinski and Sidi [22]; Chang assumes in [7] that, for each fixed $s \le t$, the moment generating function of $R(t) - R(s)$ is exponentially bounded and then establishes another calculus under this assumption.

Hoeffding's Bounds for Phase 2. Instead of a-priori imposing a bound on the probability of the event in (2.6), we can establish a bound from the following assumptions: (A1) R_i, $1 \le i \le I$, are mutually independent; (A2) R_i is α_i-smooth, all $1 \le i \le I$; (A3) we know ρ such that $\mathsf{E}[R(t) - R(s)] \le \rho(t-s)$, for all $s \le t$.

We argue that the above assumptions make sense for EF networks. Firstly, we expect at the network ingress points that the independence assumption (A1) would be verified. Secondly, with EF, it is supposed the flows are individually shaped at the network access, thus (A2) would hold. Thirdly, without loss of generality, we can suppose α_i is sub-additive; thus, for all $1 \le i \le I$, we can take

$$\rho_i = \lim_{t\to\infty} \frac{\alpha_i(t)}{t} = \inf_{t>0} \frac{\alpha_i(t)}{t},$$

thus, in (A3), one can always use $\rho = \sum_{i=1}^{I} \rho_i$.

Now the event (2.6) is the deviation of the sum $\sum_{i=1}^{I} R_i(t) - R_i(s)$ from its mean $\mathsf{E}[R(t) - R(s)]$ by $g(t-s) - \mathsf{E}[R(t) - R(s)]$. By assumption, the sum is of independent random variables (A1), with bounded support (A2), and known upper bound on $\mathsf{E}[R(t) - R(s)]$ (A3). This is exactly the framework of the bounds known as Hoeffding's inequalities [13].

For instance, one Hoeffding's inequality is for the case of a of random variables with the same support; in our setting, this corresponds to the homogeneous case $\alpha_i = \alpha_1$, for all $1 \le i \le I$. This gives, for $\rho t \le g(t) \le \alpha(t)$,

$$\mathsf{P}(\sum_{i=1}^{I} R_i(t) - R_i(0) \geq g(t)) \leq \left(\frac{\rho t}{g(t)}\right)^{-I\frac{\alpha(t)}{g(t)}} \left(\frac{\alpha(t)-g(t)}{\alpha(t)-\rho t}\right)^{-I\left(1-\frac{g(t)}{\alpha(t)}\right)}. \tag{2.7}$$

The bound is tight; it is attained for $R_i(t) - R_i(0) = \alpha_i(t)$, with probability $p_i = \frac{\rho t}{\alpha(t)}$, and $R_i(t) - R_i(0) = 0$, with probability $1 - p_i$. As pointed out by Hoeffding [13] this bound is the best one can obtain from Chernoff's bound under the given assumptions. Hoeffding [13] also gives bounds for the sum of random variables with non-identical supports, but those are looser; see [24] for the application of these bounds to our case.

A related work that combines the regulation constraints with stochastic majorization is that of Elwalid et al [11], Lo Presti et al [19], and Massoulié and Busson [17]. In [4, 8, 14] the authors impose similar assumptions as those in (A1)–(A3). We demonstrate in [23, 24] that most of the results derived under these assumptions can be obtained as an application of Hoeffding's inequalities.

The Bounds in this Paper. In the following section we give bounds in terms of the probability of the events in (2.6). Combined with a direct application of (2.7), this gives final bounds for the homogeneous case. The alert reader will enjoy making up the corresponding final formulas. For the heterogeneous case, replace (2.7) by the corresponding Hoeffding's inequality (see [23, 24]).

It is worth mentioning that it is also possible to combine the bounds in the following section with a-priori bounds on the probability of (2.6), as in [16, 26, 22], thus replacing assumptions (A1) to (A3) by direct assumptions on the probability of (2.6). This effectively provides a generalization of many results in [16, 26, 22] to the case of service curve or PSRG nodes.

3 Bounds on Backlog, Delay, and Loss

3.1 A Bound on Backlog for Service Curve Nodes

Define

$$\tau = \sup[v | v \geq 0,\ \alpha(v) \geq \beta(v)].$$

An intuitive, but over-simplifying, interpretation of τ is a (deterministic) bound on the duration of the busy period in a node with service curve β, fed by an α-smooth flow.

Proposition 3.1. *Consider a node that offers the service curve β and has a buffer of size $b \geq v(\alpha, \beta)$. For any fixed $t > 0$, any $n \in \mathbb{N}$ and any $0 = s_0 \leq s_1 \leq \cdots \leq s_{n-1} \leq s_n = t - (t-\tau)^+$,*

$$\mathsf{P}(Q(t) \geq q) \leq \sum_{k=0}^{n-1} \mathsf{P}(R(t) - R(t - s_{k+1}) \geq \beta(s_k) + q). \tag{3.8}$$

The condition on buffer means that it is large enough for loss-free operation.

Proof. (sketch) Fix any $t > 0$. By definition of the service curve β,

$$Q(t) \leq \sup_{s\in[0,t]} [R(t) - R(t-s) - \beta(s)].$$

Notice for α-smooth input and τ as defined above, we can restrict the supremum in the above inequality to $[(t-\tau)^+, t]$. Next, for any fixed $n \in \mathbb{N}$ and any $0 = s_0 \leq s_1 \leq \cdots \leq s_{n-1} \leq s_n = t - (t-\tau)^+$,

$$\{Q(t) \geq q\} \subset \bigcup_{k\in[0,n-1]} \{R(t) - R(t - s_{k+1}) \geq \beta(s_k) + q\}.$$

(3.8) can then be obtained by the union bound $\mathsf{P}(\bigcup_n A_n) \leq \sum_n \mathsf{P}(A_n)$. □

3.2 A Bound on Delay for Service Curve Nodes

In order to accommodate the stationary framework, locally to this section, we extend the definition of the GR property to hold on $\mathbb{R}$; the whole real line. We say that a node offers the GR with rate r and latency e, if there exists a sequence, defined by the recursion in (1.2), such that for all $n \in \mathbb{Z}$, there exists some $m \leq n$ with $V_m \leq T_m$. We suppose $(T_n)_n$ and $(T'_n)_n$, $n \in \mathbb{Z}$, are defined on $\mathbb{R}$.

Consider a GR node with rate r and latency e. Let D_n be delay incurred at the node by a packet labeled with n. We show in [25] that

$$D_n \leq \frac{\tilde{Q}(T_n)}{r} + e, \tag{3.9}$$

where $\tilde{Q}$ is the backlog of the hypothetical work-conserving server with constant rate r, which is fed with the same input as our original system. It is known that for PSRG node with rate r and latency e, the delay at any time t is bounded with $\frac{Q(t)}{r} + e$; this is known as the delay-from-backlog property, and it holds for both FIFO and non-FIFO PSRG nodes [6]. In contrast, the delay cannot be bounded from backlog for GR nodes. Notice that by (3.9) we establish the delay-from-backlog property for GR nodes, but with respect to the backlog of the hypothetical system. This observation has a merit in deriving probabilistic bounds on delay as we will see later.

Suppose R and Q are right-continuous with left-hand limits. From (1.3), and $D_n = T'_n - T_n \leq V_n - T_n + e$, we obtain:

$$D_n \leq \max_{k\leq n}\left[\frac{R(T_n) - R(T_k-)}{r} - (T_n - T_k)\right] + e.$$

Define $\tilde{Q}(t) = \sup_{s\leq t}[R(t) - R(s) - r(t-s)]$, and then combining with the last inequality we directly obtain (3.9). A reflection to Reich's formula [20] justifies the interpretation of $\tilde{Q}$ as the backlog of a work-conserving server with constant service rate r.

Now to consider the waiting-time distribution we need to introduce Palm probability P^0_T [2], which is in simple terms probability of an event at time 0 with $T_0 = 0$, i.e. there exists an arrival at 0. Suppose the increments of R are stationary, and $\mathsf{E}[R(t)] = \bar{\rho} t$. Then, from distributional Little's law [15] and (3.9), we obtain, for any $u \geq 0$ (see [25], Eq. (10)),

$$\mathsf{P}^0_T(D_0 \geq u) \leq \frac{r}{\bar{\rho}} \mathsf{P}(\tilde{Q}(0) \geq r(u - e - \frac{L_{max}}{r})^+). \tag{3.10}$$

Notice the bound is in terms of the stationary distribution of the backlog $\tilde{Q}$, which can be further bounded as shown in Section 3.1.

Lastly, we briefly comment an alternative approach [24]. We could use the delay-from-backlog bound of PSRG nodes and obtain (3.10) with $\tilde{Q}$ replaced by Q, the backlog of the original system. Then, again we bound the right-hand side in (3.10) as in Section 3.1. However, note that this would amount to accounting for the latency e twice, which will give a loser bound than if we use (3.10) instead.

3.3 A Bound on Loss for Service Curve Nodes

Assume time is discrete ($t \in \mathbb{Z}^+$). Fix any $m \geq 0$. Redefine

$$\tau(m) = \max\{v \in \mathbb{Z}^+ | \alpha(v) \geq \beta(v) + b + m\}. \tag{3.11}$$

For example, with $\alpha(t) = \rho t + \sigma$, such that $\rho < r$, we have $v(\alpha, \beta) = \sigma + \rho e$, and

$$\tau(m) = \frac{v(\alpha, \beta) - b - m}{r - \rho} + e.$$

Proposition 3.2. *Consider a node that offers the service curve β and has a buffer of size $b < v(\alpha, \beta)$. Then, for any $t > 0$ and $m > 0$,*

$$\mathsf{P}(L(t) - L(t-1) \geq m) \leq \sum_{s=(t-\tau(m))^+}^{t-1} \mathsf{P}(R(t) - R(s) \geq \beta(t-s) + b + m), \tag{3.12}$$

and $\mathsf{P}(L(t) - L(t-1) \geq 0) = 1$.

Remark 3.1. *The following may help in positioning the result. The bound comes from majorizing $L(t) - L(t-1)$ with $L(t) - L(s)$, where $s < t$ is the beginning of the busy period in which t is contained. This is rigorously true if β would be a strict service curve.*

Another observation from [24], which can be also observed from the sketch of the proof below, is that for all $t > 0$,

$$L(t) - L(t-1) \leq [\tilde{Q}(t) - b] 1_{Q(t)=b} \leq (\tilde{Q}(t) - b)^+,$$

where $\tilde{Q}$ is the backlog of the hypothetical system with input R and output $R^* = R \otimes \beta$ (min-plus linear system [5]), with buffer of size $\tilde{b} \geq v(\alpha, \beta)$. Notice, $L(t) - L(t-1) \leq v(\alpha, \beta) - b$, for any $t > 0$.

We have [24]:

$$\mathsf{E}[L(t)-L(t-1)] \le \mathsf{E}[(\tilde{Q}(t)-b)^+] \le [v(\alpha,\beta)-b]\mathsf{P}(\tilde{Q}(t)\ge b). \tag{3.13}$$

Suppose R is with stationary increments and $\mathsf{E}[R(u)] = \bar{\rho}u$. Then, for any fixed $t \ge \tau$, a bound on the long-run loss ratio $\bar{l}$ is [24]:

$$\bar{l} = \frac{\mathsf{E}[L(t)-L(t-1)]}{\bar{\rho}} \le \frac{v(\alpha,\beta)-b}{\bar{\rho}}\mathsf{P}(\tilde{Q}(t)\ge b).$$

Proof. (sketch) We have $L(t)-L(t-1) = [L(t)-L(t-1)]1_{Q(t)=b} \le [L(t)-L(s)]1_{Q(t)=b}$, for any $s \in [0, t-1]$. It follows from the definition of the service curve β that there exists some $s \in [0,t]$ such that

$$L(t)-L(s) \le R(t)-R(s)-\beta(t-s)-Q(t).$$

If $s=t$, then $Q(t) \le -\beta(0) \le 0$, and thus $L(t)-L(t-1)=0$. As a result, we can write: for some $s \in [0, t-1]$,

$$\begin{aligned} L(t)-L(t-1) &\le [R(t)-R(s)-\beta(t-s)-b]1_{Q(t)=b} \\ &\le [R(t)-R(s)-\beta(t-s)-b] \vee 0. \end{aligned}$$

Thus, for any $m>0$,

$$\{L(t)-L(t-1)\ge m\} \subset \bigcup_{s\in[(t-\tau(m))^+,t-1]} \{R(t)-R(s)\ge \beta(t-s)+b+m\}. \tag{3.14}$$

Above we reduce the union of events from $[0, t-1]$ to $[(t-\tau(m))^+, t-1]$. This follows directly from the definition of $\tau(m)$ (3.11). Suppose $t > \tau(m)$. Then, for all $s \in [0, t-\tau(m)-1]$, $\{R(t)-R(s) \ge \beta(t-s)+b+m\}$ is not true. (3.12) is obtained from (3.14) by the union bound. □

3.4 A Bound on Loss for Adaptive Service Curve Nodes

We refine Proposition 3.2 under the stricter assumption that the node offers an adaptive guarantee. Thus it applies in particular to a PSRG node, whereas the bound in Proposition 3.2 applies more generally to a GR node. Time is discrete ($t \in \mathbb{Z}^+$). By convention, $[x,y]=\emptyset$, for any $x>y$.

Theorem 3.1. *Consider a node that offers the adaptive service curve β and has a buffer of size $b < v(\alpha,\beta)$. Then, for any $t>0$ and $m>0$,*

$$\mathsf{P}(L(t)-L(t-1)\ge m) \le \sum_{z=(t-\tau(m))^+}^{t-1} \mathsf{P}(R(t)-R(z)\ge \beta(t-z)+b+m) \wedge B(z), \tag{3.15}$$

and $\mathsf{P}(L(t)-L(t-1)\ge 0)=1$, where

$$B(z) = \min_{u\in[z+1,t-1]} \mathsf{P}(R(t)-R(u)\ge \beta(t-u)+m).$$

Remark 3.2. *Compare with Proposition 3.2 (which is under a weaker assumption that the node offers the service curve β). The bound in (3.15) may be tighter, and this is confirmed numerically in Section 4. An alert reader may also find useful to relate the result with that of Cruz and Liu [9]. As an aside, note that the result can be readily extended to a more general definition of the adaptive guarantee (def. 7.2.1 in [5], Chapter 7), but we defer that for brevity.*

From the proof below we can deduce the following counterpart to the last bound in (3.13):

$$\mathsf{E}[L(t) - L(t-1)] \le ([v(\alpha, \beta) - b] \wedge [\alpha(1) - \beta(1)])^+ \, \mathsf{P}(\tilde{Q}(t) \ge b),$$

where $\tilde{Q}$ is as defined in (3.13). This bound is clearly tighter than the last bound in (3.13).

Proof. The definition of the adaptive service curve β is equivalent to: for all $t \ge 0$, and all $s \in [0, t]$, either

$$R^*(t) - R^*(s) \ge \beta(t - s), \tag{3.16}$$

or

$$\exists \ u \in [s, t] : \ R^*(t) \ge R'(u) + \beta(t - u). \tag{3.17}$$

Note that the inequality in (3.16) can be written as:

$$R(t) - R(s) \ge \beta(t - s) + Q(t) - Q(s) + L(t) - L(s), \tag{3.18}$$

and, the inequality in (3.17) as:

$$R(t) - R(u) \ge \beta(t - u) + Q(t) + L(t) - L(u). \tag{3.19}$$

Fix any $t > 0$. Let z be the largest integer in $[0, t]$ such that $R^*(t) \ge R'(z) + \beta(t - z)$. There always exists such an integer; note that both (3.16) and (3.17) are verified for $u = s = 0$. Notice that by definition of β and z, (3.16) is verified for all $s \in [z, t]$.

Further, clearly, for any $u \in [0, t - 1]$,

$$L(t) - L(t-1) = [L(t) - L(t-1)]1_{Q(t)=b} \le [L(t) - L(u)]1_{Q(t)=b}.$$

By definition of β and z, from (3.19), we have

$$L(t) - L(z) \le R(t) - R(z) - \beta(t - z) - Q(t),$$

and, from (3.18), for all $u \in [z, t]$,

$$L(t) - L(u) \le R(t) - R(u) - \beta(t - u) - Q(t) + Q(u).$$

From the last three inequalities, we conclude:

$$L(t) - L(t-1) \le \left[(R(t) - R(z) - \beta(t - z) - b) \wedge \bigwedge_{u \in [z+1, t-1]} (R(t) - R(u) - \beta(t - u)) \right] 1_{Q(t)=b}.$$

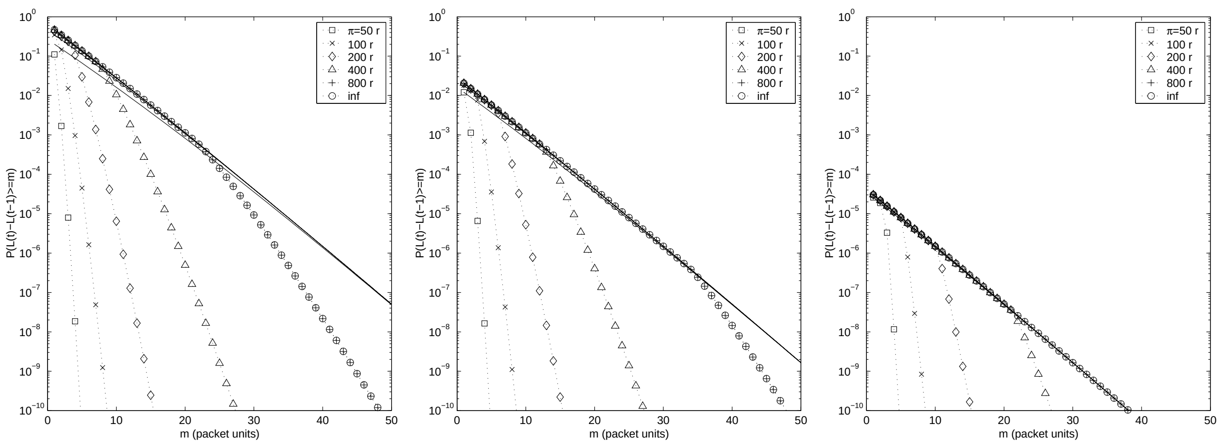

$b = 10L$ $\qquad$ $b = 20L$ $\qquad$ $b = 40L$

Figure 1: Bounds on the loss probability with varying peak-rate constraint π and the buffer size b; L is packet length. The marks denote the bound in (3.15). The solid lines show (3.12); the lines coincide in almost all the cases with an exception for $\pi = 50r$.

If $z = t$, then from (3.19), $Q(t) = 0$. Thus, we can write $L(t) - L(t-1) \leq$

$$\leq \bigvee_{z\in[0,t-1]} \left[(R(t) - R(z) - \beta(t-z) - b) \wedge \bigwedge_{u\in[z+1,t-1]} (R(t) - R(u) - \beta(t-u)) \right] 1_{Q(t)=b}$$
$$\leq \bigvee_{z\in[0,t-1]} \left[(R(t) - R(z) - \beta(t-z) - b) \wedge \bigwedge_{u\in[z+1,t-1]} (R(t) - R(u) - \beta(t-u)) \right] \vee 0.$$

And, for any $m > 0$, we obtain $\{L(t) - L(t-1) \geq m\} \subset$

$$\subset \bigcup_{z\in[(t-\tau(m))^+,t-1]} \{R(t) - R(z) \geq \beta(t-z) + b + m\} \bigcap_{u\in[z+1,t-1]} \{R(t) - R(u) \geq \beta(t-u) + m\},$$

where the reduction of the union clearly follows from the definition of $\tau(m)$ (3.11).

(3.15) follows by the union bound and the intersection bound $\mathsf{P}(\bigcap_n A_n) \leq \min_n \mathsf{P}(A_n)$. $\square$

4 Numerical Examples

We show a numerical example for the bound on loss probability in (3.15) and compare with (3.12). We consider a node that offers the adaptive service curve $\beta(t) = rt$, with $r = 150$ Mb/sec. The buffer size of the node is b and is varied. The input R is a superposition of $I = 100$ flows; each with constant packet length $L = 1500$ bytes. The flow $i \in \{1, 2, \ldots, I\}$ is regulated by the dual leaky-bucket $\alpha_i(t) = \min[\pi t, \rho_1 t + \sigma_1]$, where $\rho_1 = I^{-1}\alpha r$, $\sigma_1 = 8L$, and π is varied; α is set to 0.2 (it is an upper bound on the load). Notice π is the peak-rate constraint; ρ_1 is an upper bound on the sustained rate; σ_1 is the burstiness parameter. Time is discrete with each discrete time slot being equal to $\frac{L}{r}$ of real time (one slot equal to packet transmission time). We use Hoeffding's bound (2.7) on the probabilities in (3.12) and (3.15).

Numerical results are shown in Fig. 1. Observe that the bound (3.12) remains fairly insensitive to the value of π. In contrast, we observe, smaller π is, smaller the bound (3.15) is. This is to be expected. Recall that the term $B(z)$ in the bound (3.15) accounts for the arrival overflow over small timescales. It is exactly for these small timescales where the peak-rate constraint may act a role, and may decrease the term $B(z)$. It is to be expected from (3.15), larger b is, lessen the impact of $B(z)$ is; this can be observed from the numerical results.

References

[1] R. Agrawal, R. L. Cruz, C. Okino, and R. Rajan, *Performance bounds for flow control protocols*, IEEE/ACM Trans. on Networking **7** (1999), no. 3, 310–323.

[2] Francois Baccelli and Pierre Brémaud, *Elements of queueing theory*, vol. 26, Applications of Mathematics, Springer-Verlag, 1991.

[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*, Proc. of IEEE INFOCOM'2001, March 2001.

[4] Robert Boorstyn, Almut Burchard, Jorg Liebeherr, and Chaiwat Ootamakorn, *Statistical service assurances for traffic scheduling algorithms*, Journal of Selected Areas in Communications **18** (2000), no. 12, 2651–2664.

[5] J.-Y. Le Boudec and P. Thiran, *Network calculus*, Springer-Verlag, 2001, (also available on-line at http://ica1www.epfl.ch/http://ica1www.epfl.ch/PS_files/NetCal.htm).

[6] Jean-Yves Le Boudec and Anna Charny, *Packet scale rate guarantee for non-fifo nodes*, Proc. of IEEE Infocom 2002 (New York, USA), June 2002.

[7] Cheng-Shang Chang, *Performance guarantees in communication networks*, Springer-Verlag, 2000.

[8] Cheng-Shang Chang, Wheyming Song, and Yuh ming Chiu, *On the performance of multiplexing independent regulated inputs*, Proc. of Sigmetrics 2001 (Massachusetts, USA), May 2001.

[9] R. L. Cruz and Hai-Ning Liu, *Single server queues with loss: A formulation*, Proc. 1993 CISS, Johns Hopkins University, March 1993.

[10] B. Davie (Editor), A. Charny, F. Baker, J. Bennett, K. Benson, J.-Y. Le Boudec, A. Chiu, W. Courtney, S. Davari, V. Firoiu, C. Kalmanek, K. K. Ramakrishnan, and D. Stiliadis, *An expedited forwarding PHB*, April 2001.

[11] A. Elwalid, D. Mitra, and R. H. Wentworth, *A new approach for allocating buffers and bandwidth to heterogeneous, regulated traffic in an ATM node*, IEEE Journal on Selected Areas in Communications **13** (1995), no. 6, 1115–1127.

[12] Pawan Goyal and Harrick M. Vin, *Generalized guaranteed rate scheduling algorithms: A framework*, ACM/IEEE Trans. on Networking **5** (1997), no. 4, 561–571.

[13] Wassily Hoeffding, *Probability inequalities for sums of bounded random variables*, American Statistical Association Journal (1963), 13–30.

[14] George Kesidis and Takis Konstantopoulos, *Extremal traffic and worst-case performance for queues with shaped arrivals*, Analysis of Communication Networks: Call Centers, Traffic and Performance (edt. D. R. McDonald and S. R. E. Turner), Fields Institute Communications/AMS, ISBN 0-8218-1991-7 (2000).

[15] T. Konstantopoulos and G. Last, *On the dynamics and performance of stochastic fluid systems*, Journal of Applied Probability **37** (2000), 652–667.

[16] J. Kurose, *On computing per-session performance bounds in high-speed multi-hop computer networks*, Proc. of ACM Sigmetrics and Performance '92 (Newport, Rhode Island), June 1992.

[17] Laurent Massoulie and Anthony Busson, *Stochastic majorization of aggregates of leaky bucket-constrained traffic streams*, preprint, http://www.research.microsoft.com/-/users/lmassoul/, 2000.

[18] Clayton M. Okino, *A framework for performance guarantees in communication networks*, Ph.D. Dissertation, UCSD, 1998.

[19] F. Lo Presti, Z. Zhang, D. Towsley, and J. Kurose, *Source time scale optimal buffer/bandwidth trade-off for regulated traffic in an ATM node*, IEEE Trans. on Networking **7** (1999), no. 4, 490–501.

[20] E. Reich, *On the integrodifferential equation of Takacs, i*, Ann. Math. Statist. **29** (1958), 563–570.

[21] H. Sariowan, *A service curve approach to performance guarantees in integrated service networks*, Ph.D. thesis, UCSD, 1996.

[22] D. Starobinski and M. Sidi, *Stochastically bounded burstiness for communication networks*, IEEE Trans. on Information Theory **46** (2000), no. 1, 206–212.

[23] Milan Vojnović and Jean-Yves Le Boudec, *Bounds for independent regulated inputs multiplexed in a service curve network element*, to appear in IEEE Trans. on Communications, a preliminary version in Proc. of IEEE Globecom 2001 (Internet Performance Symposium) (San Antonio, Texas, USA), November 2001.

[24] ———, *Stochastic analysis of some expedited forwarding networks*, Proc. of IEEE Infocom 2002 (New York, USA), June 2002.

[25] ———, *Stochastic bound on delay for guaranteed rate nodes*, Tech. Report IC/2002/22, School of Computer and Communication Sciences, EPFL, http://icwww.epfl.ch/publications/list.php, May 2002.

[26] Opher Yaron and Moshe Sidi, *Performance and stability of communication networks via robust exponential bounds*, IEEE/ACM Trans. on Networking **1** (1993), no. 3, 372–384.