

On the Impact of Finite Buffers on Per-Flow Delays in FIFO Queues

Yashar Ghiassi-Farrokhfal

Department of Electrical and Computer Engineering
University of Toronto

Florin Ciucu

T-Labs / TU Berlin

Abstract—The literature on queueing systems with finite buffers addresses mostly asymptotic performance metrics on an aggregate flow, and/or generally relies on a convenient, but provably inaccurate, approximation of the loss probability by the overflow probability in an infinite size buffer. This paper addresses *non-asymptotic per-flow* metrics in a multi-flow queueing system with finite buffer and FIFO scheduling. The analysis dispenses with the above approximation, and lends itself to several interesting insights on the impact of finite buffers on per-flow metrics. Counterintuitively, the per-flow delay distribution is not monotonous in the buffer size, and such an effect is especially visible in high burstiness regimes. Another observation is that buffer dimensioning becomes insensitive to the type of SLA constraint, e.g., fixed violation probability on either loss or delay, in high multiplexing regimes. In the particular case of aggregate scheduling, the results on the aggregate input flow significantly improve upon existing results by capturing the manifestation of bufferless multiplexing in regimes with many flows.

I. INTRODUCTION

Buffer memory is an important component in core routers. Its main role is to absorb traffic fluctuations, especially during periods of high bursts, thus preventing packet loss and domino effects like packet retransmissions, high delays, etc. Buffer overprovisioning, however, has recently become a topic of technical debate [20] and of concern regarding its detrimental impact on congestion control mechanisms [1]. In effect, there is a need for a deep understanding of sizing buffer memories, e.g., to reduce unnecessarily high costs from over-provisioning buffer memories.

The problem of sizing buffer memory can be represented with a queueing model with finite buffer. Unlike queues with infinite buffers which have been extensively analyzed using large deviation techniques (see Chapter 4 in [16]), queues with finite buffers are much less understood. In fact, one of the key metrics of interest in queues with buffers of size K , i.e., the *loss probability* (denoted by P_K^{loss} , a.k.a. the cell loss ratio) is often mapped to the *virtual overflow probability* (denoted by P_K^∞) of exceeding a threshold K in an infinite buffer queue. The direct, and very convenient approximation

$$P_K^{loss} \approx P_K^\infty$$

is generally very inaccurate. Several correction terms, e.g., $\frac{1}{u}P_K^\infty$ [22], $\frac{1-u}{u}P_K^\infty$ [4], $\frac{(1-u)P_K^\infty}{1-uP_K^\infty}$ [10], or $\frac{P_0^{loss}}{P_0^\infty}P_K^\infty$ [12], where u is the utilization factor, were shown to be accurate in some cases. Other improved approximation for

P_K^{loss} have been proposed for multiplexed On-Off [2], Markov-modulated [19], and generally stationary and ergodic sources [13], [14] in many sources asymptotic regimes by simultaneously scaling the number of sources, the service capacity, and the buffer sizes.

The loss calculation problem in data networks is closely related to the problem of dimensioning telephone networks in which “buffers”, with the physical interpretation of circuits, are to be dimensioned according to some target blocking probability for incoming calls. The classical Erlang blocking probability exact formula for Poisson arrivals and general call holding times, at a single link, has been extensively generalized to account for multi-rate(class) systems, loss networks, dynamic routing, etc. (see [24] for a recent review). The blocking probability and the loss probability P_K^{loss} have been related through the notion of effective bandwidth in [15].

While P_K^{loss} is well characterized for Poisson traffic arrivals, the case of more general arrivals is mostly restricted to aggregate results in asymptotic regimes; per-flow results are available for static priority (SP) [9] and GPS in [3]. Aggregate delay can be obtained from buffer overflow probabilities in the FIFO case [15]. The per-flow delay distributions in finite buffer queues are obtained asymptotically for GPS [17] and SP [18].

To fill the apparent lack of non-asymptotic per-flow delay results in queues with finite buffers, and for broad classes of arrivals, this paper makes the following contributions:

- analyzes a multi-flow queue with finite buffer, FIFO scheduling, and Markov modulated arrivals.
- provides bounds on the per-flow, and in particular also on the aggregate, distributions of loss probability and delay. This aspect is particularly important for service differentiation or, conversely, for buffer dimensioning subject to per-flow delay/loss constraints.
- derives results in non-asymptotic regimes, i.e., both the buffer size and the number of sources can be arbitrarily set. Moreover, the obtained results are expressed in terms of upper bounds, and do not rely on convenient technical assumptions or approximations.

The obtained results are in closed-form and explicit up to the computation of an infimum operator. In this way, the results are applicable to the practical problem of buffer dimensioning for delay-sensitive applications. Through a numerical study of the obtained results, this paper provides several interesting insights. For instance, increasing the buffer size does not




Fig. 1: Two flows at a constant-rate link with finite buffer

necessarily lead to a delay decrease, and that is due to the interplay between delay due to losses (unbounded by convention) and queuing (more exactly, fewer losses compensates for increasing queuing delay). Moreover, numerical results indicate that buffer provisioning is insensitive to the type of constraint (on loss and delay) in high multiplexing regimes. An interesting consequence is that the problem of buffer provisioning subject to per-flow delay constraints can be reduced to the much simpler problem of buffer provisioning under loss constraints.

The rest of the paper is organized as follows. The next section describes the system model. Sec. III derives upper bounds on the aggregate and per-flow loss probabilities. Sec. IV derives per-flow delay bounds by using the results on loss. These results are numerically illustrated in Sec. V. Finally, Sec. VI concludes the paper.

II. SYSTEM MODEL

The time model is discrete with events (e.g., traffic arrivals) occurring at time instants $t = 0, 1, 2, \dots$. For some arrival process, the instantaneous arrivals at time t is denoted by $a(t)$, and the total arrivals during $(s, t]$ is denoted by $A(s, t)$. Note that $a(t) = A(t - 1, t)$ and denote $A(t) = A(0, t)$.

We consider the queuing scenario depicted in Fig. 1. A through (arrival) flow A_0 and a cross flow A_c share a FIFO link with capacity C and a buffer with size K . To define the arrival processes, in great generality, we use the bounding approach of the Stochastic Network Calculus [6], [11]. For some generic arrival process $A(t)$ (standing for both $A_0(t)$ and $A_c(t)$), we consider a model with a *statistical sample path envelope* \mathcal{G} and bounding function $\varepsilon(\sigma)$, satisfying for all $t, \sigma \geq 0$

$$\Pr\left\{\max_{s \leq t} (A(s, t) - \mathcal{G}(t - s; \sigma)) \geq 0\right\} \leq \varepsilon(\sigma). \quad (1)$$

Such an envelope can be constructed for the broad class of Exponentially Bounded Burstiness (EBB) traffic arrivals [23]. An arrival process A is said to be EBB with parameters (M, ρ, α) , and is represented by $A \sim (M, \rho, \alpha)$, if for all $0 \leq s \leq t$ and $\sigma \geq 0$

$$\Pr\{A(s, t) \geq \rho(t - s) + \sigma\} \leq Me^{-\alpha\sigma}. \quad (2)$$

Indeed, if $A \sim (M, \rho, \alpha)$, then for any $\gamma > 0$ (see [11], The-

orem 3.13)

$$\mathcal{G}(t; \sigma) = (\rho + \gamma)t + \sigma; \quad \varepsilon(\sigma) = \frac{Me^{-\alpha\sigma}}{1 - e^{-\alpha\gamma}} \quad (3)$$

is a statistical sample path envelope satisfying Eq. (1).

The performance metric of interest is the *per-flow* delay, in particular of A_0 ; auxiliary and helpful metrics are the total backlog and loss processes (i.e., of the aggregate $A_0 + A_c$). To define these processes, suppose that A is some generic arrival process at a node (link + buffer) and D is the corresponding departure process. The backlog at some time is the total number of stored units in the buffer due to insufficient link capacity. In turn, the delay of a traffic unit from A is defined as the time difference between the arrival and departure of that particular unit. If the unit is lost, due to buffer overflow, then the delay is infinite by convention.

Clearly, no losses occur in a node with unlimited buffer capacity. This assumption considerably simplifies the backlog and delay analysis, since the backlog $B(t)$ and virtual delay $W(t)$ processes, at some time t , can be simply expressed as

$$B(t) = A(t) - D(t) \quad (4)$$

and

$$W(t) = \inf\{s \geq 0 \mid A(t) \leq D(t + s)\}. \quad (5)$$

Moreover, the network calculus provides backlog and delay bounds for various scheduling algorithms. The following delay bound is helpful in this paper:

Theorem 1 (Per-flow delay for FIFO and infinite buffer [5]): Consider the queuing scenario from Fig. 1 with the two flows being served in FIFO order, and assume that $K = \infty$. Denote $A_{tot} = A_0 + A_c$ and assume that it has a statistical sample path envelope \mathcal{G}_{tot} with bounding function ε_{tot} , as in Eq. (1). Then the delay process $W_0(t)$ of the through flow satisfies for all $t \geq 0$

$$\Pr\{W_0(t) > d(\sigma)\} \leq \varepsilon_{tot}(\sigma),$$

where $d(\sigma) = \inf\{s \mid \forall \tau \geq 0 : C(\tau + s) > \mathcal{G}_{tot}(\tau; \sigma)\}$.

If the buffer size is finite, however, then the expressions for $B(t)$ and $W(t)$ from Eqs. (4)-(5) get more compounded since an inherent loss process must be accounted for. Our approach to analyze the delay process, in which we are particularly interested in, will be to separate the loss process from the rest. The delay of the loss process is unbounded, whereas the delay of the rest of the arrival process can be bounded with Theorem 1; Section IV will elaborate on this approach.

III. LOSS PROBABILITY

In this section we provide upper bounds on the loss probability in a finite buffer system; such bounds will be used for the main per-flow delay analysis. We first treat the aggregate loss and then the per-flow loss.

A. Aggregate loss probability

Consider the queueing scenario from Fig. 1 at the aggregate level, i.e., we are interested in the loss probability for the aggregate $A = A_0 + A_c$. Assume that A is EBB with $A \sim (M, \rho, \alpha)$. To analyze the corresponding loss process $L(t)$, we need the backlog process $B(t)$ defined recursively as

$$B(t) = \min([B(t-1) + a(t) - C]_+, K) \quad \forall t \geq 1,$$

with the initial condition $B(0) = 0$. Cruz and Liu [8] converted this recursion into the non-recursive expression

$$B(t) = \min_{0 \leq u \leq t} \left(\max_{u \leq s \leq t} \left(A(s, t) - C(t-s), \right. \right. \\ \left. \left. A(u, t) - C(t-u) + K \right) \right). \quad (6)$$

Since $L(t) = [B(t-1) + a(t) - C - K]_+$, a non-recursive expression for $L(t)$ immediately follows from Eq. (6) [8]

$$L(t) = \min_{0 \leq u < t} \left(\max_{u \leq s < t} \left([A(s, t) - C(t-s) - K]_+, \right. \right. \\ \left. \left. A(u, t) - C(t-u) \right) \right). \quad (7)$$

Next we use these deterministic expressions for $B(t)$ and $L(t)$ to derive loss and backlog bounds probabilities. For some fixed $0 \leq u \leq t$, define $j = t - u$, $x = Me^{-\alpha(C-\rho)}$, and $y = e^{-\alpha K}$. Then, from Eq. (7), for $\ell \geq 0$

$$\Pr\{L(t) \geq \ell\} \\ \leq \min_{1 \leq j < t} \Pr \left\{ \max_{t-j \leq s < t} \left(A(s, t) - C(t-s) - K, \right. \right. \\ \left. \left. A(t-j, t) - Cj \right) \geq \ell \right\} \quad (8)$$

$$\leq \min_{1 \leq j < t} \left(\Pr\{A(t-j, t) \geq Cj + \ell\} \right. \\ \left. + \sum_{s=t-j+1}^{t-1} \Pr\{A(s, t) \geq C(t-s) + K + \ell\} \right) \quad (9)$$

$$\leq e^{-\alpha \ell} \min \left(x, \min_{1 < j < t} (x^j + x^{j-1}y + \dots + xy) \right). \quad (10)$$

In Eq. (8) we used $P(X \cap Y) \leq \min(P(X), P(Y))$ for any events X and Y , and in Eq. (9) we used Boole's inequality. In the last line we used the EBB definition of A .

If $x + y > 1$, then the optimal choice of j in Eq. (9) is $j = 1$; otherwise, $j = t - 1$. Then, from Eq. (10), we have

$$\Pr\{L(t) \geq \ell\} \leq Me^{-\alpha \ell} e^{-\alpha(C-\rho)} X(\alpha, K), \quad (11)$$

where

$$X(\alpha, K) = \min \left(1, \frac{e^{-\alpha K}}{1 - e^{-\alpha(C-\rho)}} \right). \quad (12)$$

We point out that Eq. (11) can be recovered using the analysis from [21], and only partially recovered using the analysis from [8] with $X(\alpha, K) = \frac{e^{-\alpha K}}{1 - e^{-\alpha(C-\rho)}}$ (i.e., the second term in the minimum from Eq. (12)). Our bound thus improves upon the corresponding bound obtained from [8] by implicitly capturing the loss probability in a bufferless regime, i.e., for the values of K for which the minimum in Eq. (12) evaluates

to 1. The improvement can be quite sharp in multiplexing scenarios, and is of practical interest as it manifests itself for small buffer sizes.

The loss probability from Eq. (10) lends itself to an upper bound on the average loss per time unit, i.e.,

$$E[L(t)] = \sum_{\ell=1}^{\infty} \Pr\{L(t) \geq \ell\} \\ \leq \frac{e^{-\alpha}}{1 - e^{-\alpha}} Me^{-\alpha(C-\rho)} X(\alpha, K). \quad (13)$$

This bound further lends itself to a bound on the loss probability P_K^{loss} , i.e.,

$$P_K^{loss} \leq \frac{E[L(t)]}{\rho^{av}}, \quad (14)$$

where ρ^{av} denotes the average rate of arrivals.

A useful result for later is a bound on the backlog process $B(t)$. Following the same argument to obtain Eq. (11) from Eq. (7), but working directly on Eq. (6), we get for all $b \geq 0$

$$\Pr\{B(t) \geq b\} \leq Me^{-\alpha(b-K)} e^{-\alpha(C-\rho)} X(\alpha, K). \quad (15)$$

B. Per-flow loss probability

Now we consider the queueing scenario from Fig. 1 at the per-flow level, and derive an upper bound on the *per-flow* (i.e., A_0) loss probability $P_{K,0}^{loss}$. A trivial bound can be obtained by first computing the average loss $E[L_{tot}(t)]$ for the aggregate $A_0 + A_c$ as in Eq. (13). Then, considering the losses of the aggregate flow as an upper bound for the losses of the through flow A_0 , we get

$$P_{K,0}^{loss} \leq \frac{E[L_{tot}(t)]}{\rho_0^{av}}, \quad (16)$$

where ρ_0^{av} is A_0 's average rate. This bound, however, is conceivably loose because of the worst-case assumption on the through flow's losses.

In the following we derive tighter per-flow loss probability bounds than the one from Eq. (16), by carefully accounting for per-flow losses and properties of FIFO scheduling. We distinguish two cases: a general case which dispenses with any statistical independence assumptions on the flows, and a second case which considers such additional assumptions.

1) *General case (no statistical independence assumptions):* Consider the scenario depicted in Fig. 1. Both the through and cross flows are EBB with parameters $A_0 \sim (M_0, \rho_0, \alpha_0)$ and $A_c \sim (M_c, \rho_c, \alpha_c)$. The aggregate $A_{tot} = A_c + A_0$ is also EBB with parameters $A_{tot} \sim (M_{tot}, \rho_0 + \rho_c, \alpha_{tot})$ [23], where

$$M_{tot} = M_0 + M_c, \quad \alpha_{tot} = \frac{\alpha_0 \alpha_c}{\alpha_0 + \alpha_c}.$$

We now make the technical observation that a necessary condition to have at least ℓ losses of the through flow at time slot t is that the total loss and the through flow arrivals in that time slot are both larger than ℓ . This joint condition can be formally expressed as

$$\{L_0(t) \geq \ell\} \subset \left(\bigcap \left\{ \begin{array}{l} \{a_0(t) \geq \ell\} \\ \{L_{tot}(t) \geq \ell\} \end{array} \right. \right). \quad (17)$$

Applying the inequality $P(X \cap Y) \leq \min(P(X), P(Y))$ for some events X and Y , Eq. (17) yields:

$$\begin{aligned} & \Pr\{L_0(t) \geq \ell\} \\ & \leq \min\left(\Pr\{a_0(t) \geq \ell\}, \Pr\{L_{tot}(t) \geq \ell\}\right) \\ & \leq \begin{cases} M_{tot}e^{-\alpha_{tot}(\ell+C-\rho_c-\rho_0)}X(\alpha_{tot}, K) & \text{if } \ell < \ell_{cr} \\ M_0e^{-\alpha_0(\ell-\rho_0)} & \text{if } \ell \geq \ell_{cr} \end{cases}. \end{aligned} \quad (18)$$

In the last line we used the EBB definition of the through flow along with Eq. (11), and we also computed the minimum from Eq. (18). ℓ_{cr} is given as the value of ℓ for which the two terms in Eq. (19) are equal, i.e.,

$$\ell_{cr} = \frac{\log(We^{-\alpha_0\rho_0}) - \log(M_0)}{\alpha_{tot} - \alpha_0},$$

where

$$W = M_{tot}e^{-\alpha_{tot}(C-\rho_c-\rho_0)}X(\alpha_{tot}, K).$$

Using the obtained bound on the distribution of $L_0(t)$, an upper bound on the through flow's average instantaneous loss (per time unit) is

$$\begin{aligned} E[L_0(t)] &= \sum_{\ell=1}^{\infty} \Pr\{L_0(t) \geq \ell\} \\ &\leq \sum_{1 \leq \ell < \ell_{cr}} We^{-\alpha_{tot}\ell} + \sum_{\ell \geq \ell_{cr}} M_0e^{-\alpha_0(\ell-\rho_0)} \\ &= \frac{e^{-\alpha_{tot}} - e^{-\alpha_{tot}\ell_{cr}}}{1 - e^{-\alpha_{tot}}}W + \frac{M_0e^{-\alpha_0(\ell_{cr}-\rho_0)}}{1 - e^{-\alpha_0}}. \end{aligned} \quad (20)$$

Finally, an upper bound on the per-flow loss probability is

$$P_{K,0}^{loss} \leq \frac{E[L_0(t)]}{\rho_0^{av}}. \quad (21)$$

2) *Additional statistical independence assumptions:* Concretely, here we assume that $A_0(t)$ and $A_c(t)$ are statistically independent, and that they also have independent increments. These additional assumptions conceivably lend themselves to an improvement on the loss probability bound from Eq. (21).

Indeed, using the independence of $A_0(t)$ and $A_c(t)$, a tighter EBB envelope can be first obtained for the aggregate $A_{tot}(t)$ [23]. Denote for convenience by the subscript *min* (*max*) as the subscript of the flow with the smaller (larger) α value. For instance, if $\alpha_0 < \alpha_c$, then $\alpha_{min} = \alpha_0$ and $M_{min} = M_0$. Then, $A_{tot} = A_0 + A_c$ is EBB with parameters $A_{tot} \sim (M_{tot}, \rho_0 + \rho_c, \alpha_{tot})$, where M_{tot} and α_{tot} are chosen such that or all $\sigma \geq 0$

$$M_{max} + 2M_{min} + \sigma M_{min}M_{max}\alpha_{min} \leq M_{tot}e^{(\alpha_{min}-\alpha_{tot})\sigma}.$$

Next, let us observe that the total traffic that can be served without loss at slot t is $[K - B(t-1)]_+ + C$. Suppose that the cross flow arrivals $a_c(t)$ leave only k units available for the through flow to be served without loss. Then, in order to have at least ℓ through flow losses at time slot t , the through flow arrivals $a(t)$ must exceed $\ell + k$. Varying k over all possible

values provides a necessary condition for the through flow loss event:

$$\{L_0(t) \geq \ell\} \subset \left(\exists 0 \leq k \leq K + C : \bigcap \left\{ \begin{aligned} & a_0(t) \geq \ell + k \\ & B_{tot}(t-1) + a_c(t) \geq Z \end{aligned} \right\} \right) \quad (22)$$

where $Z(k) = K + C - k$. Since $a_0(t)$ and $a_c(t) + B_{tot}(t-1)$ are independent events, according to the enforced statistical independence assumptions, Eq. (22) implies that

$$\Pr\{L_0(t) \geq \ell\} \leq \max_{0 \leq k \leq K+C} \left(\Pr\{a_0(t) \geq \ell + k\} \times \Pr\{B_{tot}(t-1) + a_c(t) \geq Z\} \right). \quad (23)$$

Moreover, since $a_c(t)$ and $B_{tot}(t-1)$ are also independent events, we get the bound

$$\Pr\{B_{tot}(t-1) + a_c(t) \geq Z(k)\} \leq U(k), \quad (24)$$

by applying Lemma 6.1 in [11] and Eq. (15), where

$$\begin{aligned} U(k) &= 1 - Y(K) \left(1 + \frac{\alpha_{tot}e^{-\alpha_c(Z(k)-\rho_c)}}{\alpha_{tot} - \alpha_c} \right. \\ &\quad \left. - \left(1 + \frac{\alpha_{tot}M_c}{\alpha_c - \alpha_{tot}} \right) e^{-\alpha_{tot}(Z(k)-\rho_c)} \right). \end{aligned}$$

and $Y(K) = M_{tot}e^{-\alpha_{tot}(C-\rho_c-\rho_0-K)}X(\alpha_{tot}, K)$. Inserting Eq. (24) into Eq. (23), and using the EBB definition of the through flow, we further get

$$\begin{aligned} \Pr\{L_0(t) \geq \ell\} &\leq \max_{0 \leq k \leq K+C} \left(\Pr\{a_0(t) \geq \ell + k\}U(k) \right) \\ &\leq e^{-\alpha_0\ell} \max_{0 \leq k \leq K+C} \left(e^{-\alpha_0(k-\rho_0)}U(k) \right). \end{aligned}$$

Integrating the tail bounds yields

$$\begin{aligned} E[L_0(t)] &= \sum_{\ell \geq 1} \Pr\{L_0(t) \geq \ell\} \\ &\leq \frac{e^{-\alpha_0}}{1 - e^{-\alpha_0}} \max_{0 \leq k \leq K+C} \left(e^{-\alpha_0(k-\rho_0)}U(k) \right). \end{aligned} \quad (25)$$

Inserting Eq. (25) into Eq. (21) finally gives an upper bound on the through flow loss probability; note that all these bounds are invariant to the time t , and thus hold in steady-state.

IV. PER-FLOW DELAY ANALYSIS

In this section we reconsider the queueing scenario from Fig. 1 and derive the main result from this paper, i.e., per-flow (i.e., for A_0) bounds on the delay distribution.

The delay at a finite buffer link must account for the delay of data losses (infinite by convention) and the queueing delay of the rest of the data. Denoting the through flow's delay process by $W_0(t)$, the total probability law yields for all $d_0 \geq 0$

$$\begin{aligned} \Pr\{W_0(t) > d_0\} &\leq \Pr\{W_0(t) = \infty\} + \Pr\{W_0(t) < \infty\} \times \\ &\quad \Pr\{W_0(t) > d_0 \mid W_0(t) < \infty\}, \end{aligned} \quad (26)$$

where we used $\Pr\{W_0(t) > d_0 \mid W_0(t) = \infty\} = 1$ in the first line. Note also that $\Pr\{W_0(t) = \infty\} = P_{K,0}^{loss}$, for which an




Fig. 2: Modelling the link with finite buffer from Fig. 1 by concatenating a clipper and a link with unlimited buffer

upper bound is available from the previous section. The second probability can be bounded with $\Pr\{W_0(t) < \infty\} \leq 1$.

In turn, to compute an upper bound on the third probability in Eq. (26), we use *clippers*. These are bufferless network elements which take as input an arrival process, drop the data units which violate a predefined (deterministic) envelope function, and output a process conforming to the envelope. In other words, the output $D(t)$ of a clipper associated with an envelope function $G(t)$ satisfies for all $t \geq 0$

$$\max_{0 \leq s \leq t} (D(s, t) - G(t - s)) \leq 0.$$

In our case, the FIFO link with capacity C and finite buffer K can be modelled by concatenating a clipper with envelope function $G(t) = Ct + K$ and a link with capacity C and infinite buffer, as depicted in Fig. 2 (see [7]). We denote the clipper's output, i.e., the total (aggregate) traffic which is not lost, by $A'_{tot}(t)$. Moreover, according to the definition of a clipper,

$$A'_{tot}(s, t) \leq C(t - s) + K \quad \forall 0 \leq s \leq t. \quad (27)$$

Recall that if $A_0 \sim (M_0, \rho_0, \alpha_0)$ and $A_c \sim (M_c, \rho_c, \alpha_c)$ are EBB, then their aggregate $A_{tot} = A_0 + A_c$ is EBB as well with $A_{tot} \sim (M_{tot}, \rho_0 + \rho_c, \alpha_{tot})$. Following the construction from Eq. (3), and accounting for Eq. (27), we obtain the following sample path envelope for A'_{tot} for any $\gamma > 0$

$$G'_{tot}(t; \sigma) = \min((\rho_0 + \rho_c + \gamma)t + \sigma, Ct + K)$$

with bounding function $\varepsilon(\sigma) = \frac{M_{tot}e^{-\alpha_{tot}\sigma}}{1 - e^{-\alpha_{tot}\sigma}}$. We can now apply Theorem 1 and get the following probabilistic delay bound for $A'_{tot}(t)$

$$d_0(\sigma) = \frac{\min(\sigma, K)}{C}, \quad (28)$$

in the sense that

$$\Pr\{W_0(t) > d_0(\sigma) \mid W_0(t) < \infty\} \leq \frac{M_{tot}e^{-\alpha_{tot}\sigma}}{1 - e^{-\alpha_{tot}\sigma}}$$

for any $0 < \gamma \leq C - \rho_0 - \rho_c$. Finally, by inserting the above result into Eq. (26), we get that $d(\sigma)$ from Eq. (28) is also an upper bound on the per-flow delay in a finite buffer queueing system in the sense that

$$\Pr\{W_0(t) > d_0(\sigma)\} \leq P_{K,0}^{loss} + \frac{M_{tot}e^{-\alpha_{tot}\sigma}}{1 - e^{-\alpha_{tot}\sigma}}. \quad (29)$$

We point out that $P_{K,0}^{loss}$ can be bounded according to underlying statistical independence assumptions.

V. NUMERICAL EXAMPLES

In this section we illustrate numerically several derived results so far. We consider discrete-time Markov modulated On-Off (MMOO) sources, which can be represented with a two state discrete and homogenous Markov chain. At each time slot, an MMOO source is either in state '0' during which the source is idle, or in state '1' during which the source generates traffic at constant rate P . The transition probabilities for $0 \rightarrow 1$ and $1 \rightarrow 0$ are $1 - p_{00}$ and $1 - p_{11}$, respectively.

If A is an aggregate of n statistically independent MMOO sources, then A is EBB with parameters $A \sim (1, nEb(\alpha), \alpha)$ for any $\alpha > 0$ [6], where

$$Eb(\alpha) \leq \frac{1}{\alpha} \log(p_{00} + p_{11}e^{\alpha P}) + \sqrt{(p_{00} + p_{11}e^{\alpha P})^2 - 4(p_{00} + p_{11} - 1)e^{\alpha P}}. \quad (30)$$

To account for different levels of burstiness, we consider two types of MMOO sources with parameters displayed in Table I. The average rates of a source is $\rho^{av} = \frac{P(1-p_{00})}{2-p_{00}-p_{11}}$, and is set to 0.15 Mbps for both types. An indicator of the burstiness of an MMOO source is the average cycle time T to return to the same state in the underlying Markov chain; this is given by $T = \frac{2-p_{00}-p_{11}}{(1-p_{00})(1-p_{11})}$. We set $T = 100$ ms for Type 1 and $T = 500$ ms for Type 2, which means that a Type 2 source is more bursty than a Type 1 source.

	P (Mbps)	ρ^{av} (Mbps)	T (ms)	p_{00}	p_{11}
Type 1	1.5	0.15	100	0.989	0.9
Type 2	1.5	0.15	500	0.9978	0.9802

TABLE I: Traffic parameters.

To fit the two flows in the queueing scenario from Fig. 1, we assume that there are N MMOO sources at the link, out of which N_0 form the through flow A_0 , and the rest $N_c = N - N_0$ form the cross flow A_c . We do not make any statistical assumption between the through and cross flows, and thus we use Eqs. (20) and (21) to compute loss probabilities, and Eq. (29) to compute the probabilistic delay bounds in all examples. Note, however, that the MMOO sources forming N_0 and N_c are locally independent for the EBB characterization of A to hold.

The discrete time unit is set to 1 ms. The link capacity is $C = \frac{N\rho^{av}}{u_{tot}}$ Mbps, where u_{tot} is the total link utilization; this is fixed to 90% regardless of N and the traffic types. In all examples we numerically optimize over the parameter α , for both through and cross flows, which stems from Eq. (30) and appears in all the bounds.

A. The improvement on the aggregate loss probability

The first example concerns with the improvement of our aggregate loss probability bound from Eq. (14) relative to 1) the corresponding bound obtained by applying the analysis from [8], and 2) the overflow probability approximation $\frac{P_{K,0}^{loss}}{u_{tot}}$

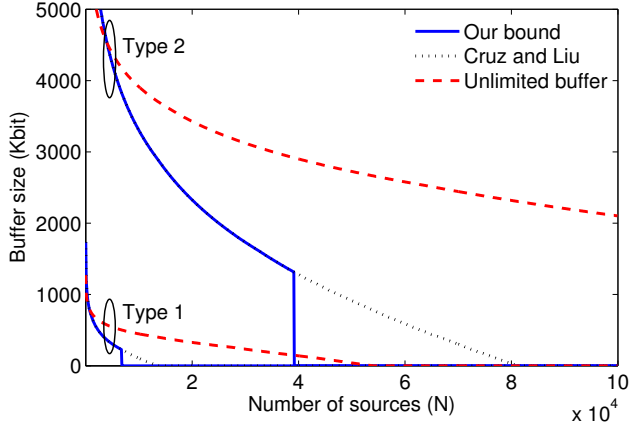


Fig. 3: Comparison of three aggregate loss probability bounds. We fix the loss probability $P^{loss} = 0.1$, the utilization = 90%, and plot the required buffer size as a function of the total number N of sources, for the two types from Table I.

from [22]. Fig. 3 illustrates the required buffer size, as a function of the number of sources N , to satisfy a loss probability $P_K^{loss} = 0.1$. There are two sets of curves corresponding to the two types of sources from Table I. En passant, we observe that the accuracy of $\frac{P_K^\infty}{u_{tot}}$, compared to the two other bounds, degrades drastically as N increases. This was also observed and justified analytically in [13], which showed in particular that P_K^∞ scales differently (in N) than the loss probability.

More interestingly, we remark the sharp improvement of our bound relative to the one from [8], especially for very bursty traffic (i.e., Type 2). The difference stems from the fact that, in addition to the finite buffer overflow in [8], our bound also captures the loss probability of bufferless multiplexing. As N increases, bufferless multiplexing eventually satisfies the loss probability constraint. The transition point can be visually perceived in Fig. 3 where our bound sharply drops to zero; formally, the transition occurs for the smallest buffer size K for which $X(\alpha, K) = 1$ in Eq. (14).

B. The impact of finite buffer on per-flow delays

Figs. 4 and 5 illustrate the difference between the per-flow delay bounds computed with Eq. (29) (which accounts for finite buffers) and with Theorem 1 (which assumes an unlimited buffer capacity).

Each curve in the plots starts from a different buffer size; note that, for values smaller than the starting point, $P_{K,0}^{loss}$ in Eq. (29) is larger than ε and thus the corresponding delay bound associated with ε is unbounded. Once the buffer size is large enough to guarantee a loss probability smaller than ε , then the system can guarantee a probabilistic delay bound. After that point, it is interesting to observe that the delay is *not* monotonous in the buffer size. At the beginning, the delay increases since the queuing delay has a smaller weight than the (infinite) delay due to losses. In that case, increasing the buffer size leads to accommodating more bursty traffic which in turn increases the delay. As the buffer size increases, after

a threshold value, however, the delay due to losses starts to play a lesser role (i.e., less and less traffic is dropped) and the delay starts to decrease.

A further interesting observation is that the delay in a finite buffer system can be smaller than the delay in an unlimited buffer system. This is because large bursts may be dropped at small buffer sizes; accordingly, high delays may be cut especially for bursty arrivals, e.g., Type 2 sources. Moreover, Figs. 4 and 5 suggest that depending on the buffer size, estimating the delay using the infinite buffer approximation can be misleading. For instance, in Fig. 5b with $N_c = 10$, the inaccuracies are 27%, 22%, and 19%, respectively, for $\varepsilon = 10^{-1}$, $\varepsilon = 10^{-2}$, and $\varepsilon = 10^{-3}$; the corresponding curve from Theorem 1 is not shown in the figure for the sake of clarity.

Let us next more closely compare the delay bounds from Figs. 4 and 5 in a finite versus infinite buffer size queue, by focusing on the impact of the number of flows N and the traffic mix.

• The impact of the number of flows N :

Increasing N is equivalent to increasing the capacity while fixing the utilization. As shown in Fig. 4b, the difference between the delay bound in an infinite buffer queue and that of a finite buffer queue can be considerable for smaller values of N , i.e., $N = 300$. However, as N increases, statistical multiplexing kicks in and the difference decays. In other words, the loss probability decays and the total delay in a finite buffer system converges to the queuing delay in an unlimited buffer system.

• The impact of traffic mix:

Traffic mix is the ratio between the number of through and cross flows in the aggregate flow. We have included three different traffic mixes $N_c = 1$, $N_c = 10$, and $N_0 = 1$ in Figs. 4 and 5. The cases $N_0 = 1$ and $N_c = 1$ share the property that they have the same statistical multiplexing gain. Thus, under the unlimited buffer assumption, Theorem 1 yields identical delay bounds for both cases. The case of $N_c = 10$ has the largest required buffer size among the three traffic mixes since it features the smallest multiplexing gain. We have not plotted the corresponding delay bound for $N_c = 10$ for the sake of clarity of the plots. Comparing all curves with identical ε in Figs. 5a and 5b shows that the effect of traffic mixes varies with the burstiness of the input traffic. When the through flow is tiny ($N_0 = 1$), and for Type 2 bursty traffic, smaller buffer sizes lead to smaller delays (compared to those of Type 1 traffic). In particular, study this property for $\varepsilon = 10^{-1}$. This is happening because smaller buffer sizes cut the massive bursts of the cross flow which increase the delay of the through flow. Another important observation is that in both Figs. 5a and 5b, decreasing ε works in favor of $N_c = 1$ with respect to $N_0 = 1$. This is justified by noting that the corresponding curves to larger ε have finite delay bounds for larger values of buffer sizes; moreover, larger buffer size is equivalent to accommodating larger burstiness, which is more beneficial to the through flow when it has a larger share in the traffic mix.

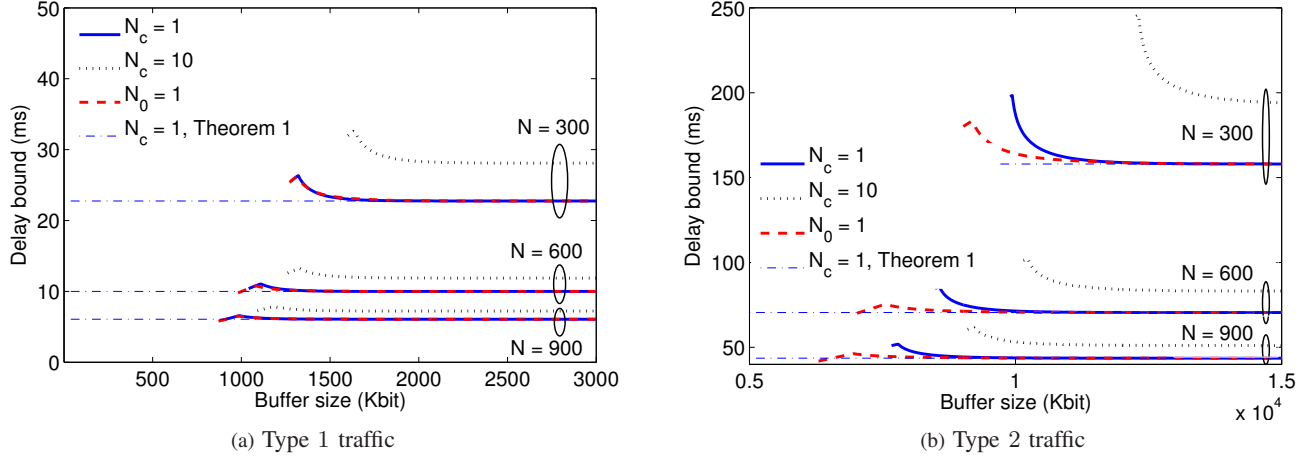


Fig. 4: The impact of considering finite buffers on per-flow delay computations; $\varepsilon = 0.1$, $N = 300, 600, 900$ with $N_0 + N_c = N$, A_0 is either very tiny (when $N_0 = 1$) or very large (when $N_c = 1, 10$), relative to the aggregate $A_0 + A_c$, utilization = 90%, and the source types from Table I (Type 1 (less bursty) in (a) and Type 2 (more bursty) in (b)); for $N_c = 1$, the infinite buffer approximation result (with Theorem 1) is also shown.

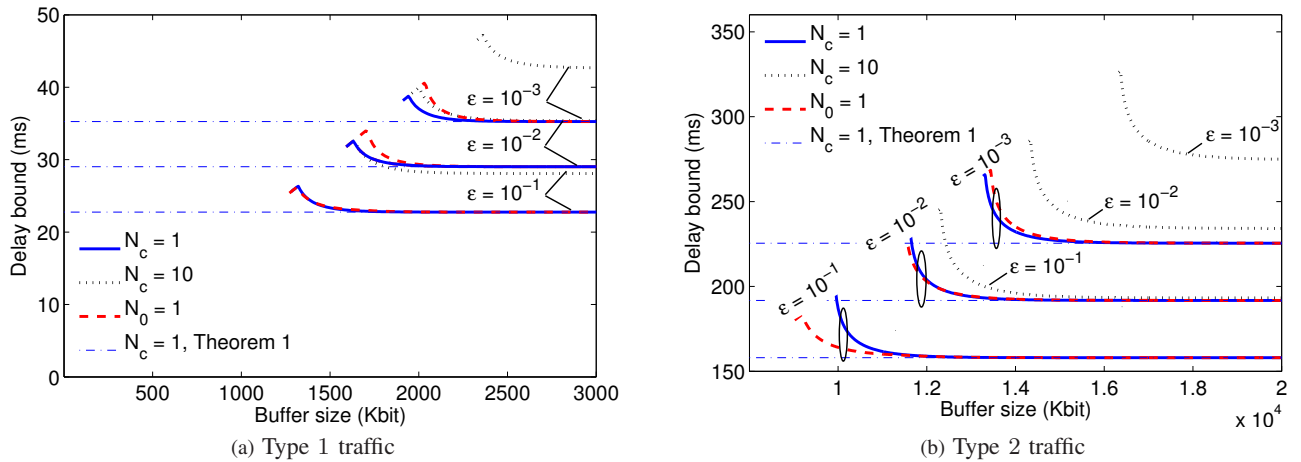


Fig. 5: The impact of considering finite buffers on per-flow delay computations; $\varepsilon = 10^{-1}, 10^{-2}, 10^{-3}$, $N = 300$, $N_0 + N_c = N$, A_0 is either very tiny (when $N_0 = 1$) or very large (when $N_c = 1, 10$), relative to the aggregate $A_0 + A_c$, utilization = 90%, and the source types from Table I (Type 1 (less bursty) in (a) and Type 2 (more bursty) in (b)); for $N_c = 1$, the infinite buffer approximation result (with Theorem 1) is also shown.

C. Buffer provisioning

Most of existing buffer provisioning results are based on a loss constraint. In this paper, a buffer provisioning method based on a delay constraint can be extracted from our delay analysis. Fig. 6a compares two buffer provisioning approaches with constraints on both delay and loss. We fix the violation probability of both methods to ε and compute the minimum buffer that satisfies the loss and delay constraints, separately. The delay constraint in this example is set to be $1.01d_0$, where d_0 is the delay in an unlimited buffer queue (we based this choice on the observations from Figs. 4 and 5 regarding delays in finite vs. infinite buffers). Fig. 6a shows that at small values of N , the required buffer for the delay constraint is

considerably larger than that of the loss constraint. However, as N increases, both provisioning methods yield the same results, which suggests that buffer provisioning is insensitive to the type of constraint in high multiplexing regimes.

VI. CONCLUSION

In this paper we have derived non-asymptotic per-flow delay bounds at a link with FIFO scheduling and finite buffer, and for a broad class of arrival processes covering Markov arrivals. The non-asymptotic aspect is particularly important since the obtained results can be applied in scenarios with a small to moderate number of flows and buffer sizes. Equally importantly, since the obtained results are derived at the per-flow level, they can be applied for service differentiation at

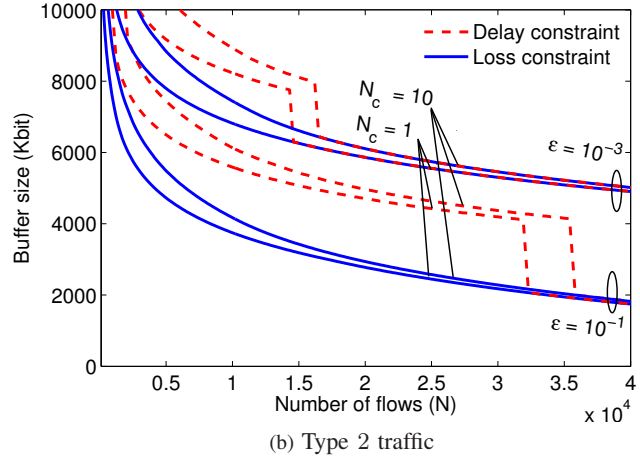
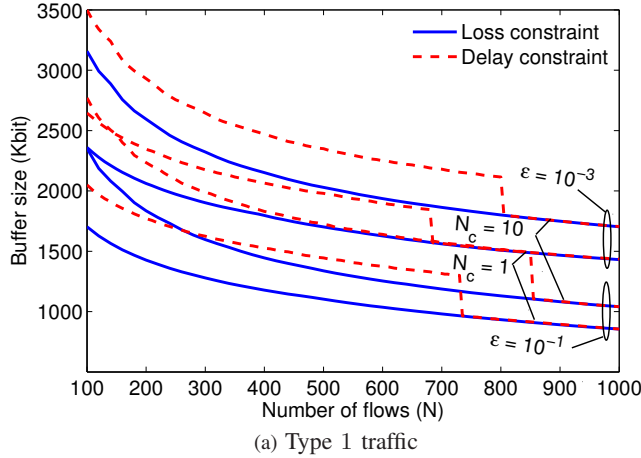


Fig. 6: Comparison of buffer size provisioning methods based on through flow loss constraint (i.e., $P_{K,0}^{loss} = \varepsilon$) and delay constraint as a function of total number of flows $N = N_0 + N_c$, utilization=90%, and the source types from Table I (Type 1 (less bursty) in (a) and Type 2 (more bursty) in (b))

a FIFO link, in particular for buffer provisioning for delay sensitive applications. Numerical illustrations revealed several interesting insights related to the dependency of the per-flow delay on the buffer size. Perhaps the most striking and counter-intuitive observation is that per-flow delay is not monotonous in the buffer size, as an effect of the interplay between delays due to losses and queueing. Another observation is that the type of probabilistic constraint (loss or delay) appears not to matter for the buffer provisioning problem in high multiplexing regimes, whereby per-flow delay constraints can be replaced by conceivably easier to handle loss constraints.

REFERENCES

- [1] BufferBloat: what's wrong with the Internet? A discussion with Vint Cerf, Van Jacobson, Nick Weaver, and Jim Gettys. *Communications of the ACM*, 55(2):40 – 47, February 2012.
- [2] A. Baiocchi, N. B. Melazzi, M. Listanti, A. Roveri, and R. Winkler. Loss performance analysis of an ATM multiplexer loaded with high-speed on-off sources. *IEEE Journal on Selected Areas in Communications*, 9(3):388 – 393, April 1991.
- [3] D. Bertsimas, I. C. Paschalidis, and J. N. Tsitsiklis. Asymptotic buffer overflow probabilities in multiclass multiplexers: an optimal control approach. *IEEE Transactions on Automatic Control*, 43(3):315 – 335, March 1998.
- [4] C. Bisdikian, J. S. Lew, and A. N. Tantawi. On the tail approximation of the blocking probability of single server queues with finite buffer capacity. In *Proc. of the Second International Conference on Queueing Networks with Finite Capacity*, pages 267 – 280, May 1992.
- [5] R. R. Boorstyn, A. Burchard, J. Liebeherr, and C. Ottamakorn. Statistical service assurances for traffic scheduling algorithms. *IEEE Journal on Selected Areas in Communications*, 18(12):2651 – 2664, December 2000.
- [6] C. S. Chang. *Performance guarantees in communication networks*. Springer Verlag, 2000.
- [7] C. S. Chang, J.-Y. Cruz, R. L. Le Boudec, and P. Thiran. A min, + system theory for constrained traffic regulation and dynamic service guarantees. *IEEE Transactions on Networking*, 10(6):805 – 817, December 2002.
- [8] R. L. Cruz and H. N. Liu. Single server queues with loss: A formulation. In *Proc. of CISS*, March 1993.
- [9] S. Delas, R. R. Mazumdar, and C. Rosenberg. Cell loss asymptotics in priority queues accessed by a large number of independent stationary sources. In *Proc. of IEEE INFOCOM*, pages 551 – 558, March 1999.
- [10] F. N. Gouweleeuw and H. C. Tijms. Computing loss probabilities in discrete-time queues. *Operations Research*, 46(1):149 – 154, January 1998.
- [11] Y. Jiang and Y. Liu. *Stochastic Network Calculus*. Springer Verlag, 2008.
- [12] H. S. Kim and N. B. Shroff. Loss probability calculations and asymptotic analysis for finite buffer multiplexers. *IEEE/ACM Transactions on Networking*, 9(6):755 – 768, December 2001.
- [13] N. Likhanov and R. R. Mazumdar. Cell loss asymptotics in buffers fed with a large number of independent stationary sources. In *Proc. of IEEE INFOCOM*, pages 339 – 346, March 1998.
- [14] N. Likhanov and R. R. Mazumdar. Cell loss asymptotics in buffers fed by heterogeneous long-tailed sources. In *Proc. of IEEE INFOCOM*, pages 173 – 180, March 2000.
- [15] N. Likhanov, R. R. Mazumdar, and F. Theberge. Providing QoS in large networks: Statistical multiplexing and admission control. In *E.K. Boukas and R. Malhame (Eds.), Analysis, Control and Optimization of Complex Dynamic Systems*, pages 701 – 726. Kluwer, 2005.
- [16] R. R. Mazumdar. *Performance Modeling, Loss Networks, and Statistical Multiplexing*. Synthesis Lectures on Communication Networks. Morgan & Claypool Publishers, 2009.
- [17] I. C. Paschalidis. Performance analysis and admission control in multimedia communication networks. In *Proc. of the 36th IEEE Conference on Decision and Control*, pages 1786 – 1791, December 1997.
- [18] S. Shakkottai and R. Srikant. Many-sources delay asymptotics with applications to priority queues. *Queueing Systems - Theory and Applications*, 39:183 – 200, October 2001.
- [19] N. B. Shroff and M. Schwartz. Improved loss calculations at an ATM multiplexer. *IEEE/ACM Transactions on Networking*, 6(4):411 – 421, August 1998.
- [20] A. Vishwanath, V. Sivaraman, and M. Thottan. Perspectives on router buffer sizing: recent results and open problems. *SIGCOMM Computer Communication Reviews*, 39(2):34 – 39, March 2009.
- [21] M. Vojnovic and J.-Y. Le Boudec. Elements of probabilistic network calculus for packet scale rate guarantee nodes. In *Int'l Symp. of Mathematical Theory of Networks and Systems (MTNS)*, August 2002.
- [22] A. K. Wong. Queueing analysis for atm switching of continuous-bit-rate traffic—a recursion computation method. In *Proc. of IEEE GLOBECOM*, pages 1438 – 1444, December 1990.
- [23] 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.
- [24] S. Zachary and I. Ziedins. Loss networks. In *R. Bourcherie and N. van Dijk (Eds.), Queueing networks: A fundamental approach*, pages 701 – 726. Springer, 2011.