Worst case burstiness increase due to FIFO multiplexing

Vicent Cholvi [a,*], Juan Echagüe [b], Jean-Yves Le Boudec [c]

[a] *Universitat Jaume I (DLSI), 12071 Castelló, Spain*
[b] *Universitat Jaume I (DICC), 12071 Castelló, Spain*
[c] *EPFL (ICA), CH-1015 Lausanne, Switzerland*

Abstract

We consider a FIFO multiplexer fed by flows that are individually constrained by arrival curves, and look for the best possible arrival curve for every output flow. This problem arises in scenarios where aggregate multiplexing is performed, such as differentiated services or front ends to optical switches. We obtain an exact result for a fluid model and for piecewise linear concave arrival curves, which are common in practice and correspond to combinations of leaky buckets.
© 2002 Elsevier Science B.V. All rights reserved.

Keywords: FIFO; Network calculus; Differentiated services; Aggregate scheduling

1. Introduction

We consider a FIFO multiplexer fed by flows that are individually constrained by arrival curves. This scenario arises in scenarios where aggregate multiplexing is performed such as: Internet differentiated services [2,4,8,13], or front ends to optical switches [18]. Multiplexing several flows into a FIFO scheduler causes an increase in the burstiness of every flow. Capturing this effect is important in order to properly dimension buffers in complex scenarios where multiplexers are interconnected. However, it is not easy to capture the burstiness increase due to FIFO multiplexing, and this does not appear to be done in a general setting. Partial results indicate that, on one hand, in some cases, FIFO multiplexing may lead to instability, even when the maximum utilization is less than 1 [1,12]. On the other hand, under some strict conditions on source rate or on multiplexing architecture, one can find explicit delay and burstiness bounds for a FIFO ATM network [9,10,16].

Our problem is to quantify the worst case burstiness increase due to FIFO multiplexing. More precisely, given the set of arrival curve constraints for the input flows, we would like to find arrival curve constraints that apply to the output flows, and that are as tight as possible. In this paper, we present a first step in this direction. We consider piecewise linear concave arrival curves, which are common in practice and

---

correspond to combinations of leaky buckets. We take a fluid approach, and leave packetization effects for further study; these effects are likely to impact our results by one maximum packet size [7,15]. We find a worst case bound when the FIFO node is a constant rate server. We illustrate our bound numerically and by simulation.

We compare our bound to previous ones. A method based on a service curve approach, was proposed in [10] and further developed in [17], Chapter 6. If the arrival curve constraints are defined by a single leaky bucket, these bounds coincide with ours. Furthermore, in the general case where we have for example both peak rate and sustainable rate limitations, we show that the method of service curves could also be used to provide tight bounds for FIFO multiplexing, in general.

The paper is organized as follows. Section 2 gives our assumption and notation. Section 3 gives our main result. In Section 4 we give some simulation results that confirm our results. Section 5 compares our result with the state of the art. The proof of our main result relies on a number of technical lemmas, which are given in the Appendix A.

2. FIFO aggregate scheduling: model and notation

In this section we describe our model and assumption. We consider I flows, served as one aggregate in a constant rate server, with rate R. Aggregation of all flows is done in a FIFO manner. Call $A_i(t)$ the input function, which is defined as the number of bits observed on flow i at the *input* between 0 and t. Similarly, let $B_i(t)$ be the output function. We assume that $A_i(t)$ is left-continuous, which does not appear to be a loss of generality. In this framework, the input–output characterization of our system is as follows. Let $A(t) = \sum_{i=1}^{I} A_i(t)$ be the aggregate input function; the aggregate output function $B(t) = \sum_{i=1}^{I} B_i(t)$ is given by [6]

$$B(t) = \inf_{0 \le s \le t} A(s) + R(t - s).$$

For any time t, define $v(t)$ by

$$v(t) = \sup\{s \text{ such that } s \le t \text{ and } A(s) \le B(t)\}. \tag{1}$$

The time $v(t)$ is interpreted as the minimum of t and the arrival time of the first bit leaving after t. Then the input–output characterization for all i is:

$$B_i(t) = A_i(v(t)). \tag{2}$$

We assume that input flow i is constrained by an arrival curve α_i, in other words [11]

$$\text{for all } t, s \text{ such that } s \le t : \quad A_i(t) - A_i(s) \le \alpha_i(t - s), \tag{3}$$

and our problem is, for a given set of arrival curves $\alpha_i(t)$, to find the best possible arrival curves for the output functions $B_i(t)$, under the constraints that Eq. (3) is satisfied.

Without loss of generality, we can focus on flow $i = 1$ and consider the set of all flows $j \ne i$ as one aggregate flow. Thus we can limit ourselves to the case $I = 2$ and find an arrival curve for the output of flow 1.

In this paper, we focus on the case where the arrival curves α_i are concave piecewise linear, which correspond to constraints imposed by combination of leaky buckets. In Proposition 2, we focus on the

case $\alpha_1(x) = \min\{p_1 x, b_1 + r_1 x\}$. This corresponds to the variable bit rate case, or T-SPEC, used by the IETF [5,17] (we neglect the MTU, consistent with our fluid model assumption); p_1 is the peak rate, r_1 the sustainable rate (we assume that $p_1 \geq r_1$) and b_1 is the burst tolerance, or burstiness, of flow 1. We also assume that $\alpha_2(x)$ is concave and piecewise linear, which is consistent with the fact that, flow 2 represents the aggregate of all flows other than 1.

We fix the collection of parameters p_1, r_1, b_1 and the function $\alpha_2(x)$ and call *scenario* any arbitrary collection of functions $(A_i(t))_{1 \leq i \leq I}$ that are wide-sense increasing and non-negative, and that satisfy Eq. (3). The corresponding output functions $B_i(t)$ are given by Eq. (2). For convenience, when necessary, we use a super-index to identify a scenario. For example, for scenario γ, $B_i^\gamma(t)$ is the output function of flow i and $v^\gamma(t)$ is the minimum of t and the arrival time of the first bit leaving after t.

Let Γ be the set of all scenarios. Our problem is now to find the best possible arrival curve $\alpha_1^*(x)$ for the output flow $B_1(t)$, in other words, we should have, for any scenario $\gamma \in \Gamma$:

for all t, s such that $s \leq t$: $\quad B_1^\gamma(t) - B_1^\gamma(s) \leq \alpha_1^*(t - s)$, $\hfill (4)$

and α_1^* should be as small as possible.

Call $B_{\mathrm{req}} := \sup_{x \geq 0} [\alpha_1(x) + \alpha_2(x) - Rx]$ the worst case buffer required for a loss-free operation. We assume the finiteness condition

$$B_{\mathrm{req}} < \infty. \hfill (5)$$

Otherwise, it can easily be seen that our problem has no finite solution.

3. Arrival curve for the output flow

The following theorem gives the solution to our problem.

Theorem 1. *Consider a FIFO system serving two flows, with the assumptions in Section 2. Define*

$$\alpha_1^*(x) = \min\{Rx, \alpha_1(x + a_1(x))\},$$

where $a_1(x)$ is the maximum value for a from the set of couples $(a \geq 0, b \geq 0)$ that solve Eq. (6):

$$\alpha_1(b + a + x) - \alpha_1(a + x) + \alpha_2(b) - R(a + b) = 0. \hfill (6)$$

Then

(1) α_1^* *is an arrival curve for the output flow $B_1(t)$;*
(2) *it is the best arrival curve that can be found under these assumptions.*

The theorem requires the solution of a one-dimensional maximization problem, in order to find $a_1(x)$. We give later in this section an algorithm (see Proposition 2) to perform this when α_1 has the form $\alpha_1(x) = \min\{p_1 x, b_1 + r_1 x\}$. It can easily be generalized to the case where α_1 is concave, piecewise linear.

Proof. First, note that the finiteness condition (5) implies that for any non-negative (a, b) satisfying (6) we have $a \leq B_{\mathrm{req}}/R$, which in turn implies that $a_1(x)$ is well defined and is unique.

Second, consider some arbitrary but fixed time interval $[s, t]$. To simplify the writing, we use the notation $s^\beta = v^\beta(s)$. Define

- $\Lambda_1(s, t)$ as the set of scenarios $\beta \in \Gamma$ such that $A_1^\beta(t) - A_1^\beta(s^\beta) = \alpha_1(t - s^\beta)$,
- s_1 the minimum value of s^β among all scenarios in $\Lambda_1(s, t)$,
- $\Psi_1(s, t)$ as the set of scenarios in $\beta \in \Lambda_1(s, t)$ such that $\forall \beta \in \Psi_1(s, t)$ $(s^\beta = s_1)$,
- $\Phi_1(s, t)$ as the set of scenarios $\beta \in \Lambda_1(s, t)$ such that in time interval $[s'^\beta, s^\beta)$ (where s'^β denotes the start of the busy period[1] which last, at least, until s^β) flow 1 injects $\alpha_1(t - s'^\beta) - \alpha_1(t - s^\beta)$ bits and flow 2 injects $\alpha_2(s^\beta - s'^\beta)$ bits.

Part (1). It follows from Lemma A.3 that the number of bits output by flow 1 in $[s, t]$ is upper bounded by $\alpha_1(t - s_1)$. We now show that $a := s - s_1$ is equal to $a_1(x)$ defined in the theorem. This will establish statement (1).

From Lemma A.7, we know that s_1 is the minimum value of s^β among all scenarios in $\Phi_1(s, t)$. Now for any scenario $\beta \in \Phi_1(s, t)$, the buffer occupancy at time s^β (denoted q) is:

$$\alpha_1(t - s'^\beta) - \alpha_1(t - s^\beta) + \alpha_2(s^\beta - s'^\beta) - R(s^\beta - s'^\beta) = q.$$

Furthermore, since $[s^\beta, s]$ is a busy period then $q = (s - s^\beta)R$. Thus, we have that:

$$\alpha_1(b + a + x) - \alpha_1(a + x) + \alpha_2(b) - Rb = aR \quad \text{with } b = s^\beta - s'^\beta \text{ and } a = s - s^\beta,$$

which shows that a satisfies Eq. (6) for some $b \geq 0$. Conversely, Lemma A.8 shows that for any non-negative (a, b) satisfying Eq. (6), there is some scenario $\beta \in \Phi_1(s, t)$ such that $b = s^\beta - s'^\beta$ and $a = s - s^\beta$. Thus the minimum s_1 of all s^β is $s - a_1(t - s)$.

Part (2). Follows immediately from Lemma A.4.   □

Proposition 2. *In the case $\alpha_1(x) = \min\{p_1 x, b_1 + r_1 x\}$, $a_1(x)$ defined in Theorem 1 can be computed with the following algorithm*:

Step 1: *Define extract (V) as the function that, for a given set of intervals V returns the lower and upper values of each interval (including $+\infty$) and the angular points of α_2 (i.e., the points where α_2 changes the value of its linearity) that lie within V.*

Step 2: *Solve $\alpha_2(b) + p_1 b = (x_1 - x)R$ where $x_1 = b_1/(p_1 - r_1)$ and $b \geq 0$ is the unknown.*

- *Case where there is no solution: solve $\alpha_2(b) = (x_1 - x)R + (R - r_1)b$ where $b > 0$ is the unknown.*
  - *Case where there is no solution: $V_1 = extract([0, \infty])$, $V_2 = \emptyset$ and $V_3 = \emptyset$.*
  - *Case where there is one solution, denoted v: $V_1 = extract([0, v])$, $V_2 = \emptyset$ and $V_3 = extract([v, \infty])$.*
  - *Case where there is more than one solution: this cannot happen.*
- *Case where there is one solution, denoted v: solve $\alpha_2(b) = (x_1 - x)R + (R - r_1)b$ where $b \geq 0$ is the unknown.*
  - *Case where there is no solution: $V_1 = \emptyset$, $V_2 = extract([0, v])$ and $V_3 = extract([v, \infty])$.*
  - *Case where there is one solution, denoted v': $V_1 = extract([v', \infty])$, $V_2 = extract([0, v])$ and $V_3 = extract([v, v'])$.*
  - *Case where there are two solutions, denoted v'' the minimum value and v''' the maximum one: $V_1 = extract([v'', v'''])$, $V_2 = extract([0, v])$ and $V_3 = extract([v, v''], [v''', \infty])$.*
  - *Case where there are more than two solutions: this cannot happen.*

---

[1] A *busy period* is a period where the server buffer is non-empty.

- *Case where there is more than one solution: this cannot happen.*

Step 3: Define the following functions:

- $f_1(b) = (\alpha_2(b) + (r_1 - R)b)/R$;
- $f_2(b) = (\alpha_2(b) + (p_1 - R)b)/R$;
- $f_3(b) = (\alpha_2(b) + (r_1 - R)b + (p_1 - r_1)(x_1 - x))/(R + p_1 - r_1)$;

where $a_1(x)$ is the maximum value of $f_i(b)$ for all $b \in V_i$ (with $i = 1, 2, 3$).

Proof. With some algebra, it can be shown that $\alpha_1(x + a + b) - \alpha_1(x + a) = r_1(b - x'') + p_1 x''$, where $x'' = \max\{0, \min\{b, x_1 - (x + a)\}\}$ and $x_1 = b_1/(p_1 - r_1)$.

Substituting this result in Eq. (6) we have:

$$\max\{\alpha_2(b) + (r_1 - R)b, \min\{\alpha_2(b) + (p_1 - R)b, \alpha_2(b) + (r_1 - R)b + (p_1 - r_1)(x_1 - (x + a))\}\} = aR. \tag{7}$$

Now, by doing some algebra on Eq. (7), we have:

$$a = \begin{cases} \text{if } \dfrac{\alpha_2(b) + (r_1 - R)b}{R} + x \geq x_1, & \text{then } \dfrac{\alpha_2(b) + (r_1 - R)b}{R}, \\[2ex] \text{if } \dfrac{\alpha_2(b) + p_1 b}{R} + x \leq x_1, & \text{then } \dfrac{\alpha_2(b) + (p_1 - R)b}{R}, \\[2ex] & \text{else } \dfrac{\alpha_2(b) + (r_1 - R)b + (p_1 - r_1)(x_1 - x)}{R + p_1 - r_1}. \end{cases} \tag{8}$$

Note that, since x_1 is fixed, only one of the situations Eq. (8) will occur (and consequently only one solution is possible).

In order to obtain which values of b make a maximum we will take into account that the three functions in Eq. (8) are piecewise linear. Therefore, they take their maximum value at their angular points (which are in fact the same as in $\alpha_2(b)$) or at the points that bound the domain on b of each function. Consequently, to find $a_1(x)$ we only need to check the value of a at those points and take the maximum one.

The domain on b where each one of the three abovementioned functions must be used can be obtained by considering the points where the straight line $(x_1 - x)R$ intersects with $\alpha_2(b) + p_1 b$ (which is a concave increasing function) and where the straight line $(x_1 - x)R + (R - r_1)b$ intersects with $\alpha_2(b)$ (which is also a concave increasing function). Fig. 1(a) and (b) provide two graphical representations of the solutions in Step 2. □

Fig. 2(a) and (b) provide two numerical applications that show the worst case arrival α_1^* for the output of flow 1 predicted by Theorem 1.

4. Simulation results

In this section we perform a simulation study of the model described in Section 2. The simulations have been performed by means of a discrete event program that simulates the system at bit level.
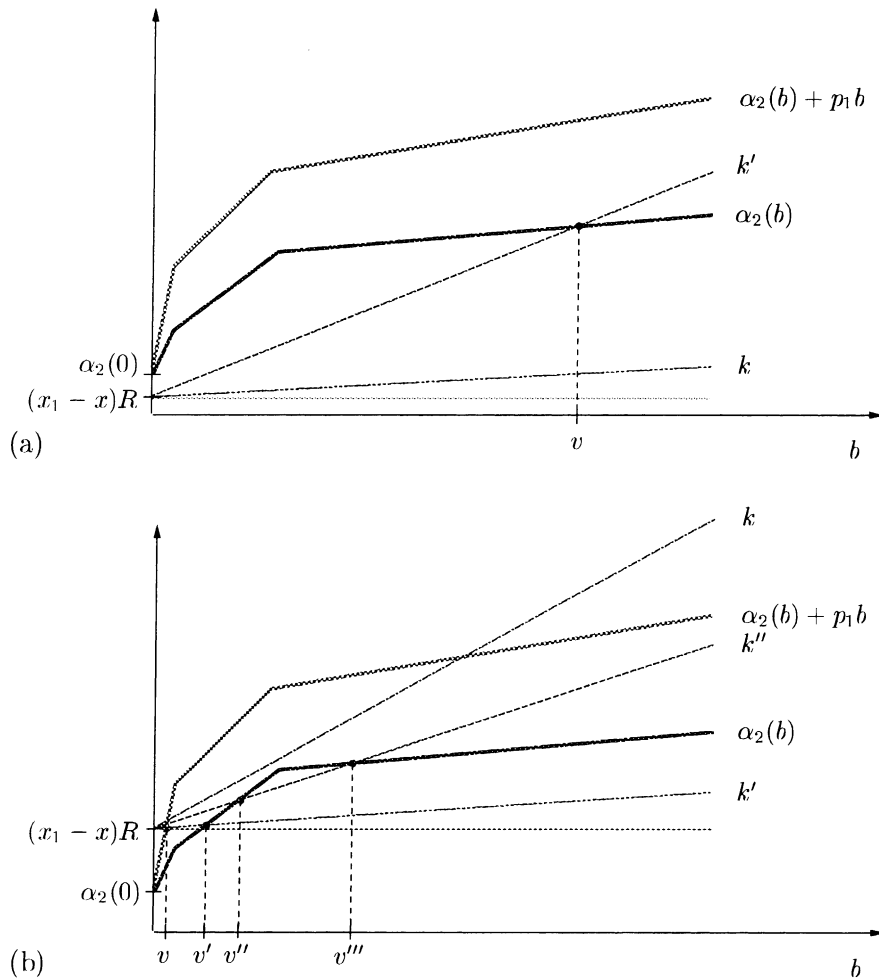
Fig. 1. Representation of the solutions in Step 2 in Proposition 2. (a) Case where there is no solution for $\alpha_2(b) + p_1 b = (x_1 - x)R$ for $b \geq 0$. k and k' represent two values of the line $(x_1 - x)R + (R - r_1)b$ which correspond to the case where $\alpha_2(b) \neq (x_1 - x)R + (R - r_1)b$ for all b and $\alpha_2(b) = (x_1 - x)R + (R - r_1)b$ for one value of b. (b) Case where there is one solution for $\alpha_2(b) + p_1 b = (x_1 - x)R$ for $b \geq 0$. k, k' and k'' represent three values of the line $(x_1 - x)R + (R - r_1)b$ which correspond to the case where $\alpha_2(b) \neq (x_1 - x)R + (R - r_1)b$ for all b, $\alpha_2(b) = (x_1 - x)R + (R - r_1)b$ for one value of b and $\alpha_2(b) = (x_1 - x)R + (R - r_1)b$ for two values of b.

In the first simulation, each input flow is conformed by a two-leaky bucket shaper that constraints them to match the arrival curve defined by equation $\alpha_i(x) = \min\{p_i x, b_i + r_i x\}$, $\forall i \in I$. Note that, instead of considering only two flows where flow 2 is an aggregate flow, we model them independently. Injected bits are served in FIFO order at a rate R.

An extensive simulation has been performed by considering many different scenarios and measuring, for each time interval, the number of bits belonging to each flow observed at the output. The arrival of packets to the traffic shaper has been implemented by means of a normal distribution with mean r_i and standard deviation p_i.
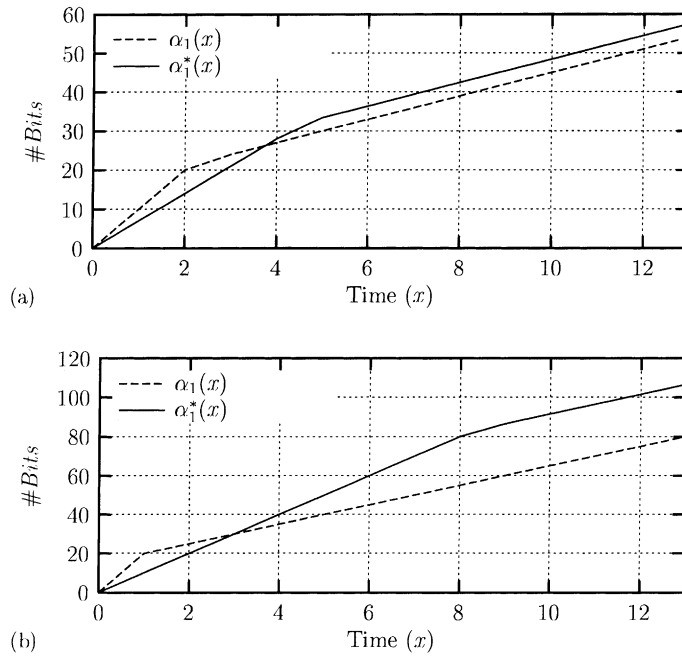
Fig. 2. Examples that show the worst case arrival α_1^* for the output of flow 1 predicted by Theorem 1. (a) Input flow 1 has arrival curve $\alpha_1(x) = \min\{10x, 15 + 3x\}$ and input flow 2 has arrival curve $\alpha_2(x) = \min\{8x, 10 + 3x\}$. The server rate is 7. (b) Input flow 1 has arrival curve $\alpha_1(x) = \min\{11x, 10 + x\}$ and input flow 2 (which is the aggregate of three flows) has arrival curve $\alpha_2(x) = \min\{11x, 10 + x\} + \min\{11x, 20 + x\} + \min\{11x, 30 + x\}$. The server rate is 10.

Fig. 3 shows the same numerical application than in Fig. 2(a). It can be readily seen that all results are bounded by our theoretical result. Furthermore, we found that such a theoretical result is also reached, which is consistent with the fact that our bound is a worst case bound.

In a second experiment (Fig. 4), we analyze how the sustainable rate of flow 2 affects $\alpha_1^*(x)$. As it was expected due to our theoretical result, the increment in the sustainable ratio of flow 2 increases the value of $\alpha_1^*(x)$.
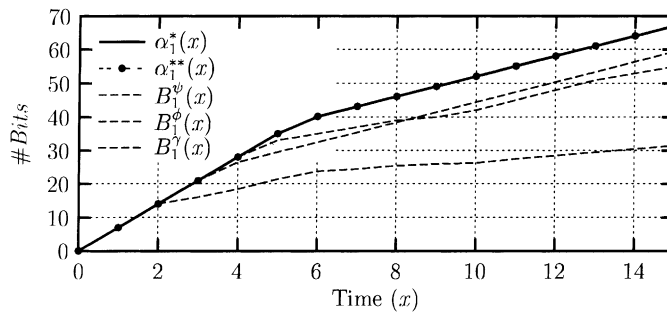

Fig. 3. Worst case arrival curve α_1^* for the output of flow 1, predicted by Theorem 1, and arrival curve α_1^{**} for the output of flow 1 obtained by simulation. The server rate is 7, $\alpha_1(x) = \min\{10x, 15 + 3x\}$ and $\alpha_2(x) = \min\{8x, 10 + 3x\}$. The curves $B_1^i(x)$ show the number of bits observed at the output for three different scenarios (ψ, ϕ, γ).
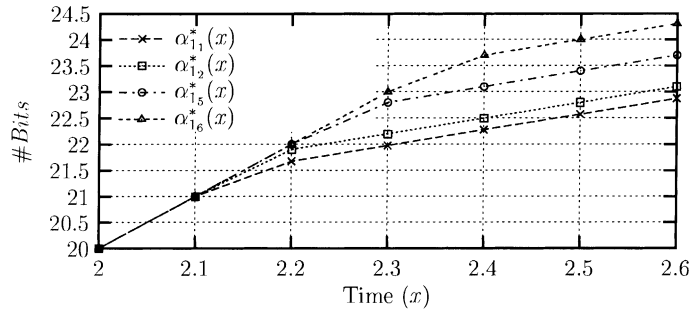
Fig. 4. Worst case arrival curves for the output of flow 1, predicted by Theorem 1 when varying the sustainable rate of flow 2. The server rate is 10, $\alpha_1(x) = \min\{10x, 15 + 3x\}$ and $\alpha_2(x) = \min\{8x, 10 + jx\}$ for $j \in \{1, 2, 5, 6\}$. $\alpha_{1_j}^*(x)$ represents the arrival curves for the output flow 1 when $\alpha_2(x) = \min\{8x, 10 + jx\}$.

5. Previous work

The state of art for aggregate multiplexing in general is surprisingly poor and so is the work done to obtain output arrival curves for FIFO multiplexing. The only sources that we are aware of is a result by Cruz in [10], which is reported and further elaborated in the book by Le Boudec and Thiran [17].

The main result that relates to our work can be summarized as follows. Take the same setting as in this paper, but assume the peak rate of flow 1 is infinite, in other words, $\alpha_1(t) = r_1 t + b_1$. Otherwise, there is no special assumption in α_2. It is shown in [17], Chapter 6, that an arrival curve for the output of flow 1 is given by

$$\alpha_1^{**}(x) = \min\left\{ Rx, b_1 + r_1 \max_{\forall u \geq 0}\left(\frac{\alpha_2(u) + (r_1 - R)u}{R} \right) + r_1 x \right\}. \tag{9}$$

It can easily be shown, after some easy but tedious algebra, that Eq. (9) coincides with the same bound that we find in this paper. It is shown in [17] that (9) is the best bound that can be found under the assumption that the peak rate of flow 1 is infinite, which is consistent with our result. Note that in the simple case where all flows (not only flow 1) are constrained by a single leaky bucket, Eq. (9) gives the formula (where $\alpha_2(t) = r_2 t + b_2$):

$$\alpha_1^{**}(x) = \min\left\{ Rx, b_1 + r_1\frac{b_2}{R} + r_1 x \right\}, \tag{10}$$

which is interesting by its simplicity.

The method, introduced by Cruz in [10], first finds a family of service curves β_θ, indexed by a real valued parameter θ, and applies traditional network calculus results to derive an arrival curve for each value of θ. Eq. (9) is then obtained as by minimizing over θ.

If we remove the assumption that the peak rate of flow 1 is infinite, by using Theorem 6.4.1 in [17] we can derive an arrival curve that can be expressed as:

$$\alpha_1^{**}(x) = \min\{Rx, \min\{b_1^* + p_1 x, b_1^{**} + r_1 x\}\}, \tag{11}$$

with $b_1^* = p_1 \max_{\forall u \geq 0}((\alpha_2(u) + (p_1 - R)u)/R)$ and $b_1^{**} = b_1 + r_1 \max_{\forall u \geq 0}((\alpha_2(u) + (r_1 - R)u)/R)$.
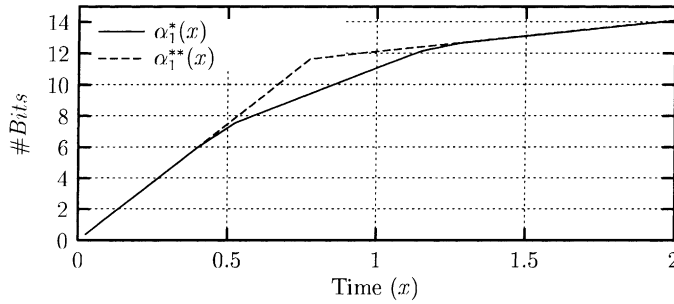
Fig. 5. Worst case arrival curve α_1^* for the output of flow 1, predicted by Theorem 1 and arrival curve α_1^{**} for the output of flow 1 predicted by Eq. (11). The server rate is 15, $\alpha_1(x) = \min\{10x, 10 + 2x\}$ and $\alpha_2(x) = \min\{50x, 1 + 10x\}$.

This bound, contrary to the previous cases, is not tight. Fig. 5 provides a numerical example. However, there is the following relationship.

Proposition 3. *Let $\alpha_1^*(x)$ be the output of flow* 1, *predicted by Theorem* 1 *and $\alpha_1^{**}(x)$ the output of flow* 1 *predicted by Eq. (11). If x is large enough then $\alpha_1^*(x) = \alpha_1^{**}(x)$.*

Proof. First of all, note that both b_1^* and b_1^{**} in Eq. (11) are constants. Since $p_1 x$ grows quickly than $r_1 x$ then, for x large enough, $b_1^* + p_1 x$ will be bigger than $b_1^{**} + r_1 x$. Therefore, in this case we have that

$$\alpha_1^{**}(x) = \min\{Rx, b_1^{**} + r_1 x\}.$$

On the other hand, if x is large enough then we have that $((\alpha_2(b) + (r_1 - R)b)/R) + x \geq x_1$ for all $b \geq 0$. By using the Eq. (8) in Proposition 2, we know that $a_1(x) = \max_{\forall b \geq 0}((\alpha_2(b) + (r_1 - R)b)/R)$.

Substituting this value in Theorem 1, we have that:

$$\alpha_1^*(x) = \min\left\{Rx, \alpha_1\left(x + \max_{\forall b \geq 0}\left(\frac{\alpha_2(b) + (r_1 - R)b}{R}\right)\right)\right\}$$

$$= \min\left\{Rx, b_1 + r_1 \max_{\forall b \geq 0}\left(\frac{\alpha_2(b) + (r_1 - R)b}{R}\right) + r_1 x\right\} = \min\{Rx, b_1^{**} + r_1 x\} = \alpha_1^{**}(x). \qquad \square$$

This last result does not imply that network calculus cannot give the worst case bound. Indeed, in accordance with Proposition 6.4.2 in [17] we can derive an arrival curve that can be expressed as:

$$\alpha_1^{**}(x) = \min\left\{Rx, \min_{\theta \geq 0}\left(\max_{u \geq 0}(\alpha_1(x + u) - \beta_\theta^1(u))\right)\right\}$$

$$= \min\left\{Rx, \min_{\theta \geq 0}\left(\max_{u \geq 0}(\alpha_1(x + u) + \alpha_2(u - \theta) - Ru)\right)\right\}.$$

Let us conjecture that the values of θ and u that solve the previous equation are $\theta = s - s_1$ and $u = s - s_1'$. Then

$$\alpha_1^{**}(x) = \min\{Rx, \alpha_1(x + s - s_1') + \alpha_2(s_1 - s_1') - R(s - s_1')\}. \tag{12}$$

After some algebra, it can be readily seen that Eq. (12) coincides with the same bound we find in this paper. This shows that the network calculus could be used to give a worst case bound.

6. Conclusions

In this paper we have analyzed the impact of FIFO multiplexing in the case where the arrival curve constraints for the input flows are concave piecewise linear functions (which are common in practice and correspond to combinations of leaky buckets). We found a worst case bound when the FIFO node is a constant rate server. Furthermore, we also provide a numerical algorithm to compute such a worst case bound in the case where α_1 corresponds to the variable bit rate case (or T-SPEC) used by the IETF.

Comparing our bound to previous ones (based on a service curve approach [10,17]), we found that if the arrival curve constraints are defined by a single leaky bucket, these bounds coincide with ours. Furthermore, in the general case, we show that network calculus could also be used to give the worst case bound.

Our results have some potential applications. In particular, they may be relevant for the Expedited Forwarding Service (EF) [14], a service which has been developed in the *Differentiated Services Working Group of IETF* [3]. The goal of the EF is to provide to an aggregate of flows some hard delay guarantees by means of ensuring that, at each hop, the aggregate requiring EF treatment receives service rate exceeding the total bandwidth requirements of all flows in the aggregate at this hop.

Some issues require further study. In this work, we used a fluid approach and, even though we know that packetization effects are likely to impact our results by one maximum packet size [7,15], understanding these effects appears to be an important issue.

Appendix A. Proof of lemmas

The following lemma shows that, for any non-greedy scenario β for flow 1 in time interval $[s^\beta, t]$, there is another greedy scenario γ for flow 1 in time interval $[s^\gamma, t]$ (with $s^\gamma > s^\beta$) that injects more flow 1 bits than scenario β in time interval $[s^\beta, t]$.

Lemma A.1. *Let β be a scenario in Γ such that, in time interval $[s^\beta, t]$, flow 1 injects $K < \alpha_1(t - s^\beta)$ bits. Then, there is a scenario γ (with $s^\gamma > s^\beta$) such that, in time interval $[s^\gamma, t]$, flow 1 injects $K = \alpha_1(t - s^\gamma)$ bits.*

Proof.

(1) Case $K \leq \alpha_1(t - s)$: take a scenario γ such that:
   (a) Flow 2 does not inject any bit.
   (b) Flow 1 injects $\alpha_1(t - s)$ bits in time interval $[s, t]$ and no bit in the rest of intervals.
   In such a scenario $s^\gamma = s$ and the number of flow 1 injected bits in time interval $[s^\gamma, t]$ is $\alpha_1(t - s^\gamma)$.
(2) Case $K > \alpha_1(t - s)$: take a scenario γ (see Fig. 6) such that:
   (a) Flow 2 behaves as in scenario β.
   (b) Flow 1 behaves as in scenario β until time instant immediately before s^β.
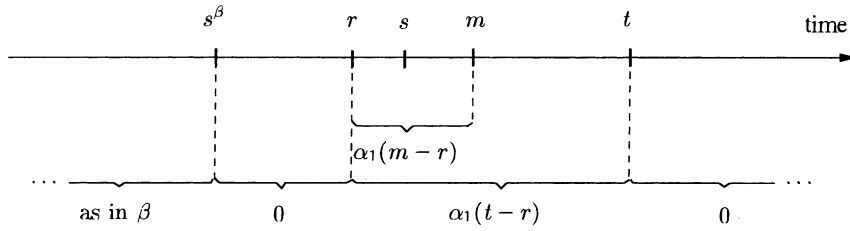   (c) Flow 1 injects no bits after time instant t.

Fig. 6. Scenario γ for Lemma A.1 showing the amount of flow 1 bits injected in different time intervals.

(d) $\exists r : s^\beta < r < s \ (A_1^\gamma(t) - A_1^\gamma(r) = \alpha_1(t - r) = K)$.
(e) $\forall m : r \le m \le t \ (A_1^\gamma(m) - A_1^\gamma(r) = \alpha_1(m - r))$.

In such a scenario, it can be readily seen that:

- s^γ is located at the first flow 1 bit injected after s^β. That is, it is located at time instant r.
- The number of flow 1 bits injected in time interval $[s^\gamma, t]$ is $\alpha_1(t - r)$. As $s^\gamma = r$ then $\alpha_1(t - s^\gamma)$.

Therefore, we have that, in time interval $[s^\gamma, t]$, flow 1 injects $K = \alpha_1(t - s^\gamma)$ bits, which proves the lemma.

However, it is also necessary to prove that γ is a valid scenario in accordance with the constraint curve for the arrival function. That is, it must be proved the following holds

$$\forall a, b : a \le b \quad (A_1^\gamma(b) - A_1^\gamma(a) \le \alpha_1(b - a)).$$

(a) Case $b < s^\beta$: immediate, since $A_1^\gamma(b) - A_1^\gamma(a) = A_1^\beta(b) - A_1^\beta(a) \le \alpha_1(b - a)$.
(b) Case $s^\beta \le b \le t$:
   (i) Case $r \le a$: immediate.
   (ii) Case $r > a$.
        By contradiction. Assume that $A_1^\gamma(b) - A_1^\gamma(a) > \alpha_1(b - a)$. That is, $\alpha_1(b - r) + d > \alpha_1(b - a)$, where d is the number of flow 1 bits injected in time interval $[a, r)$.
        As α_1 is concave,[2] we have that $\forall l \ge 0$: $\alpha_1(b - r + l) + d > \alpha_1(b - a + l)$. If we take $l = t - b$ then $\alpha_1(t - r) + d > \alpha_1(t - a)$.
        By construction of γ, we have that $A_1^\gamma(t) - A_1^\gamma(a) = \alpha_1(t - r) + d > \alpha_1(t - a)$.
        However, by construction of γ we also have that $A_1^\gamma(t) - A_1^\gamma(a) = A_1^\beta(t) - A_1^\beta(a) \le \alpha_1(t - a)$. We reach a contradiction.
(c) Case $t < b$:
   (i) Case $a > t$: immediate since $A_1^\gamma(b) - A_1^\gamma(a) = 0 \le \alpha_1(b - a)$.
   (ii) Case $a \le t$: immediate since $A_1^\gamma(b) - A_1^\gamma(a) = A_1^\gamma(t) - A_1^\gamma(a) \le \alpha_1(t - a) \le \alpha_1(b - a)$. □

Now, we state the next lemma which shows that the number of flow 1 bits injected for any scenario β in time interval $[s^\beta, t]$ must be, at most, $\alpha_1(t - s_1)$.

---

[2] Note that, if f is concave, then the increment $f(y + l) - f(x + l)$ is wide-sense decreasing with l (where $x < y$ and $l > 0$), thus $f(x) + d > f(y)$ implies $\forall l \ge 0$: $f(x + l) + d > f(y + l)$.

Lemma A.2. *For any time interval* $[s, t]$, *we have that* $\forall \beta \in \Gamma$ $(A_1^\beta(t) - A_1^\beta(s^\beta) \leq \alpha_1(t - s_1))$.

Proof. By contradiction. Assume there is a scenario γ, such that, for some time interval $[s, t]$, $A_1^\gamma(t) - A_1^\gamma(s^\gamma) > \alpha_1(t - s_1)$.

- Case $s^\gamma \geq s_1$: since $A_1^\gamma(t) - A_1^\gamma(s^\gamma) \leq \alpha_1(t - s^\gamma)$ and $\alpha_1(t - s^\gamma) \leq \alpha_1(t - s_1)$ then $A_1^\gamma(t) - A_1^\gamma(s^\gamma) \leq \alpha_1(t - s_1)$. We reach a contradiction.
- Case $s < s^\gamma < s_1$:
  - Case $A_1^\gamma(t) - A_1^\gamma(s^\gamma) = \alpha_1(t - s^\gamma)$: this implies that $A_1^\gamma(t) - A_1^\gamma(s^\gamma) \leq \alpha_1(t - s_1)$. We reach a contradiction.
  - Case $A_1^\gamma(t) - A_1^\gamma(s^\gamma) < \alpha_1(t - s^\gamma)$: from Lemma A.1, there is a scenario ϕ such that $A_1^\gamma(t) - A_1^\gamma(s^\gamma) \leq A_1^\phi(t) - A_1^\phi(s^\phi) = \alpha_1(t - s^\phi)$, being $s^\phi > s^\gamma$.
    However, this implies that $A_1^\phi(t) - A_1^\phi(s^\phi) \leq \alpha_1(t - s_1)$ and consequently $A_1^\gamma(t) - A_1^\gamma(s^\gamma) \leq \alpha_1(t - s_1)$. We reach a contradiction.
- Case $s^\gamma = s$: since $A_1^\gamma(t) - A_1^\gamma(s^\gamma) \leq \alpha_1(t - s)$ and $\alpha_1(t - s) \leq \alpha_1(t - s_1)$ then $A_1^\gamma(t) - A_1^\gamma(s^\gamma) \leq \alpha_1(t - s_1)$. We reach a contradiction. $\square$

From the previous lemma, we can derive the following lemma. Roughly speaking it states that $\alpha_1^*(t - s) = \min\{R(t - s), \alpha_1(t - s_1)\}$ is a valid arrival curve for the function B_1.

Lemma A.3. *For any time interval* $[s, t]$, *we have that* $\forall \beta \in \Gamma$ $(B_1^\beta(t) - B_1^\beta(s) \leq \min\{R(t - s), A_1^\beta(t) - A_1^\beta(s^\beta)\} \leq \min\{R(t - s), \alpha_1(t - s_1)\})$.

Proof. Immediate. On the first hand, given a time interval $[s, t]$, the FIFO buffer cannot transmit at a rate higher than R (i.e., $R(t - s)$). Furthermore, the FIFO server cannot transmit, in time interval $[s, t]$, more that what is injected in time interval $[s^\beta, t]$, which, as shown in Lemma A.2, is at most $\alpha_1(t - s_1)$. $\square$

Now, we prove that the abovementioned arrival curve is optimal in the sense that, for each flow and time interval, there is a scenario for which the formula in Definition 4 is exactly an equality.

Lemma A.4. *For any time interval* $[s, t]$, *we have that* $\exists \beta \in \Psi_1(s, t)$ $(B_1^\beta(t) - B_1^\beta(s) = \min\{R(t - s), \alpha_1(t - s_1)\})$.

Proof. Take some scenario $\gamma \in \Psi_1(s, t)$. Now, take another scenario β such that:

(1) Flow 2 behaves as in scenario γ until time instant immediately before s_1 and then stops injecting.
(2) Flow 1 behaves as in scenario γ until time instant immediately before s_1.
(3) Flow 1 injects no bits after time instant t.
(4) $\forall m: s_1 \leq m \leq t$ $(A_1^\beta(m) - A_1^\beta(s_1) = \alpha_1(m - s_1))$.

First, we prove that β is a valid scenario in accordance with the constraint curve for the arrival function. That is, we will prove that the following holds:

$$\forall a, b : a \leq b \quad (A_1^\beta(b) - A_1^\beta(a) \leq \alpha_1(b - a)).$$

(1) Case $b < s_1$: immediate, since $A_1^\beta(b) - A_1^\beta(a) = A_1^\gamma(b) - A_1^\gamma(a) \le \alpha_1(b - a)$.

(2) Case $s_1 \le b \le t$:

    (a) Case $s_1 \le a$: immediate (by condition 4).

    (b) Case $s_1 > a$.

        By contradiction. Assume that $A_1^\beta(b) - A_1^\beta(a) > \alpha_1(b-a)$. That is, $\alpha_1(b-s_1)+d > \alpha_1(b-a)$, where d is the number of flow 1 bits injected in time interval $[a, s_1]$.

        As α_1 is concave we have that $\forall l \ge 0$: $\alpha_1(b - s_1 + l) + d > \alpha_1(b - a + l)$. If we take $l = t - b$ then $\alpha_1(t - s_1) + d > \alpha_1(t - a)$.

        By construction of β, we have that $A_1^\beta(t) - A_1^\beta(a) = \alpha_1(t - s_1) + d > \alpha_1(t - a)$.

        However, by construction of γ we also have that $A_1^\beta(t) - A_1^\beta(a) = A_1^\gamma(t) - A_1^\gamma(a) \le \alpha_1(t - a)$. We reach a contradiction.

(3) Case $t < b$:

    (a) Case $a > t$: immediate since $A_1^\beta(b) - A_1^\beta(a) = 0 \le \alpha_1(b - a)$.

    (b) Case $a \le t$: immediate since $A_1^\beta(b) - A_1^\beta(a) = A_1^\beta(t) - A_1^\beta(a) \le \alpha_1(t - a) \le \alpha_1(b - a)$.

Therefore, we have that $\beta \in \Psi_1(s, t)$ and that it behaves in a greedy fashion in time interval $[s_1, t]$. Consequently, the buffer content at time instant $r \in (s, t]$ will be $q(r) = \max\{0, \alpha_1(r - s_1) - R(r - s)\}$. Thus, we have that:

- Case $R(t - s) < \alpha_1(t - s_1)$: at time instant t, we have that $q(t) > 0$. This means that not all injected packets have been transmitted and since FIFO is work conserving then $R(t - s)$ flow 1 bits have be transmitted.
- Case $R(t - s) \ge \alpha_1(t - s_1)$: since $R(t - s) \ge \alpha_1(t - s_1)$ then, at time instant t, we have that $q(t) = 0$. This means that all bits injected in time interval $[s_1, t]$ have been transmitted at time t. Namely, $\alpha_1(t - s_1)$. $\qquad\square$

Lemma A.5. $\Psi_1(s, t) \subseteq \Phi_1(s, t)$.

Proof. By contradiction. Consider a scenario $\gamma \in \Psi_1(s, t)$ such that in time interval $[s'^\gamma, s_1)$ either flow 1 injects $K < \alpha_1(t - s'^\gamma) - \alpha_1(t - s_1)$ bits or flow 2 injects $K' < \alpha_2(s_1 - s'^\gamma)$ bits (or both things). Clearly $\gamma \notin \Phi_1(s, t)$.

Now, take a scenario β such that:

(1) Flows 1 and 2 start injecting bits after at time s'^γ.

(2) Flow 1 injects no bits after time instant t.

(3) Flow 2 injects no bits after time instant s_1.

(4) $\forall m$: $s'^\gamma \le m \le t$ $(A_1^\beta(t) - A_1^\beta(m) = \alpha_1(t - m))$.

(5) $\forall m$: $s'^\gamma \le m \le s_1$ $(A_2^\beta(m) - A_2^\beta(s'^\gamma) = \alpha_2(m - s'^\gamma))$.

Clearly, β is a valid scenario in accordance with the constraint curve for the arrival function.

Now, we can see that the number of bits injected in time interval $[s'^\gamma, s_1)$ is bigger in β than in γ. Namely, in β is $\alpha_1(t - s'^\gamma) - \alpha_1(t - s_1) + \alpha_2(s_1 - s'^\gamma)$ and in γ is $K + K'$. Consequently $s^\beta < s_1$.

As (by Condition 4) $\forall m$: $s'^\gamma \le m \le t$ $(A_1^\beta(t) - A_1^\beta(m) = \alpha_1(t - m))$, we reach a contradiction just taking $m = s^\beta$. $\qquad\square$

Lemma A.6. *For each scenario $\beta \in \Phi_1(s, t)$, there is another scenario $\gamma \in \Phi_1(s, t)$ with $s'^\gamma = s'^\beta$ and $s^\gamma = s^\beta$ such that in time interval $[s^\gamma, t]$ flow 1 injects $\alpha_1(t - s^\gamma)$ bits.*

Proof. Take a scenario γ such that:

(1) Flow 2 behaves as in scenario β.
(2) Flow 1 starts injecting bits at time instant s'^β.
(3) Flow 1 injects no bits after time instant t.
(4) Flow 1 injects $\alpha_1(t - s'^\beta) - \alpha_1(t - s^\beta)$ bits in time interval $[s'^\beta, s^\beta)$ in a greedy fashion.
(5) $\forall m: s^\beta \leq m \leq t \ (A_1^\gamma(t) - A_1^\gamma(m) = \alpha_1(t - m))$.

It can be readily seen that γ is a valid scenario in accordance with the constraint curve for the arrival function. Indeed, regarding flow 2 scenario γ behaves as in scenario β. Regarding flow 1, we will prove that the following holds:

$$\forall a, b : a \leq b \quad (A_1^\gamma(b) - A_1^\gamma(a) \leq \alpha_1(b - a)).$$

(1) Case $b < s^\beta$: immediate.
(2) Case $s^\beta \leq b \leq t$: let c be the first time instant in time interval $[s'^\beta, s^\beta)$ where flow 1 stops injecting bits. Remember that flow 1 in such a interval injects a given number of bits (namely $\alpha_1(t-s'^\beta)-\alpha_1(t-s^\beta)$) in a greedy fashion.
   (a) Case $c \leq a$: immediate.
   (b) Case $c > a$: by contradiction. Assume that $A_1^\gamma(b) - A_1^\gamma(a) > \alpha_1(b - a)$. By construction of γ we have that $A_1^\gamma(t) - A_1^\gamma(s'^\beta) = \alpha_1(a - s'^\beta) + \alpha_1(t - b) + A_1^\gamma(b) - A_1^\gamma(a) > \alpha_1(a - s'^\beta) + \alpha_1(t - b) + \alpha_1(b - a) > \alpha_1(t - s'^\beta)$. However, by construction of γ we also know that $A_1^\gamma(t) - A_1^\gamma(s'^\beta) = \alpha_1(t - s'^\beta)$. We reach a contradiction.
(3) Case $t < b$:
   (a) Case $a > t$: immediate since $A_1^\gamma(b) - A_1^\gamma(a) = 0 \leq \alpha_1(b - a)$.
   (b) Case $a \leq t$: immediate since $A_1^\gamma(b) - A_1^\gamma(a) = A_1^\gamma(t) - A_1^\gamma(a) \leq \alpha_1(t - a) \leq \alpha_1(b - a)$.

Because of the γ definition (see Condition 4), the number of flow 1 bits injected in time interval $[s'^\beta, m]$ (for all $s'^\beta \leq m < s^\beta)$) in γ is greater or equal than in β. Therefore, $[s'^\beta, s^\beta)$ is also a busy period in scenario γ. Now, since the number of flow 1 injected bits in $[s'^\beta, s^\beta)$ is the same in both scenarios (namely, it is $\alpha_1(t - s'^\beta) - \alpha_1(t - s^\beta)$), we have that $s^\gamma = s^\beta$ and $s'^\gamma = s'^\beta$.

Furthermore, as $A_1^\gamma(s^\beta) - A_1^\gamma(s'^\beta) = \alpha_1(t - s'^\beta) - \alpha_1(t - s^\beta)$ and $A_1^\gamma(t) - A_1^\gamma(s'^\beta) = \alpha_1(t - s'^\beta)$ then $A_1^\gamma(t) - A_1^\gamma(s^\beta) = \alpha_1(t - s^\beta)$. This ends the proof. □

Lemma A.7. *s_1 is the minimum value of s^β among all scenarios in $\Phi_1(s, t)$.*

Proof. By contradiction. Assume that the minimum value of s^β among all scenarios in $\Phi_1(s, t)$, denoted s_2, is not s_1.

- Case $s_2 < s_1$: by Lemma A.6, we can obtain a scenario $\gamma \in \Phi_1(s, t)$ such that, in time interval $[s_2, t]$ flow 1 injects $\alpha_1(t - s_2)$. Thus $\gamma \in \Lambda_1(s, t)$. Consequently, s_1 is not the minimum value of s^β among all scenarios in $\Lambda_1(s, t)$ and (by definition of s_1) we reach a contradiction.

- Case $s_2 > s_1$: by Lemma A.5, $\Psi_1(s, t) \subseteq \Phi_1(s, t)$. We reach a contradiction. $\qquad\square$

Lemma A.8. *For any non-negative (a, b) satisfying Eq. (6), there is some scenario $\beta \in \Phi_1(s, t)$ such that $b = s^\beta - s'^\beta$ and $a = s - s^\beta$.*

Proof. For a given time interval $[s, t]$, take a scenario β such that:

(1) Flows 1 and 2 start injecting bits at time instant $s - (a + b)$.
(2) Flows 1 and 2 stop injecting bits at time instant $s - a$.
(3) Flow 1 injects $\alpha_1(b + a + x) - \alpha_1(a + x)$ bits in time interval $[s - (a + b), s - a)$ in a greedy fashion.
(4) Flow 2 injects $\alpha_2(b)$ bits in time interval $[s - (a + b), s - a)$ in a greedy fashion.
(5) (a, b) satisfy Eq. (6).

Clearly β is a valid scenario in accordance with the constraint curve for the arrival function. We must prove that $a = s - s^\beta$ and $b = s^\beta - s'^\beta$.

Since in time interval $[s - (a + b), s - a)$ both flow 1 and 2 are greedy and since $\alpha_1(b + a + x) - \alpha_1(a + x) + \alpha_2(b) - Rb = Ra > 0$ then time interval $[s - (a + b), s - a)$ is a busy period. Furthermore, the buffer occupancy at time $s - a$ will be Ra. Consequently, by time s all those bits will be transmitted.

This shows that $s^\beta = s - a$ and $s'^\beta = s - (a + b) = s^\beta - b$, which proves the lemma. $\qquad\square$

References

[1] M. Andrews, Instability of fifo in session-oriented networks, in: Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2000), January 2000.
[2] J.C.R. Bennett, K. Benson, A. Charny, W.F. Courtney, J.-Y. Le Boudec, Delay jitter bounds and packet scale rate guarantee for expedited forwarding, in: Proceedings of the Infocom, April 2001.
[3] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, An architecture for differentiated services, RFC 2475, IETF, December 1998.
[4] R. Boorstyn, A. Burchard, J. Liebeherr, C. Oottamakorn, Statistical service assurances for traffic scheduling algorithms, IEEE J. Selected Areas Commun. 18 (12) (2000) 2651–2664 (special issue on Internet QoS).
[5] B. Braden, D. Clark, S. Shenker, Integrated services in the Internet architecture: an overview, RFC 1633, IETF, June 1994.
[6] C.S. Chang, On deterministic traffic regulation and service guarantee: a systematic approach by filtering, IEEE Trans. Inform. Theory 44 (1998) 1096–1107.
[7] C.S. Chang, Performance Guarantees in Communication Networks, Springer, New York, 2000.
[8] A. Charny, J.-Y. Le Boudec, Delay bounds in a network with aggregate scheduling, in: Proceedings of the First International Workshop on Quality of Future Internet Services, Berlin, Germany, September 2000.
[9] I. Chlamtac, A. Faragó, H. Zhang, A. Fumagalli, A deterministic approach to the end-to-end analysis of packet flows in connection oriented networks, IEEE/ACM Trans. Netw. 4 (6) (1998) 422–431.
[10] R.L. Cruz, Sced+: efficient management of quality of service guarantees, in: Proceedings of the IEEE Infocom'98, San Francisco, March 1998.
[11] R.L. Cruz, Quality of service guarantees in virtual circuit switched networks, IEEE J. Selected Areas Commun. 13 (6) (1995) 1048–1056.
[12] B. Hajek, Large bursts do not cause instability, IEEE Trans. Auto. Control 45 (2000) 116–118.
[13] S.D. Patek, J. Liebeherr, A. Burchard, A calculus for end-to-end statistical service guarantees, Technical Report CS-2001-19, Department of Computer Science, University of Virginia, August 2001.
[14] V. Jacobson, K. Nichols, K. Poduri, An expedited forwarding phb, RFC 2598, IETF, June 1999.
[15] J.-Y. Le Boudec, Some properties of variable length packet shapers, in: Proceedings of the ACM Sigmetrics/Performance'01, 2001.

[16] J.-Y. Le Boudec, G. Hebuterne, Comment on a deterministic approach to the end-to-end analysis of packet flows in connection oriented network, IEEE/ACM Trans. Netw. 8 (1) (2000).

[17] J.-Y. Le Boudec, P. Thiran, Network Calculus, Lecture Notes in Computer Science, Vol. 2050, Springer, Berlin, July 2001 (available online at http://icawww.epfl.ch).

[18] L. Thylén, G. Karlsson, O. Nilsson, Switching technologies for future guided wave optical networks: potentials and limitations of photonics and electronics, IEEE Commun. Mag. 34 (2) (1996) 106–113.

Vicent Cholvi graduated from the University of Valencia (Spain) and received his doctorate in 1994 from the Polytechnic University of Valencia (Spain). In 1995, he joined the Jaume I University in Castelló where he is currently an associate professor. He is in charge of the Distributed Systems Group and his interests are in distributed and communication systems.

Juan Echagüe received the B.S. and M.S. degrees, both in Computer Sciences, from the Polytechnic University of Valencia. He is currently a Ph.D. student at the Distributed Systems Group at the Jaume I University (Spain). His current interests are related to provide with quality of service to broadband networks.
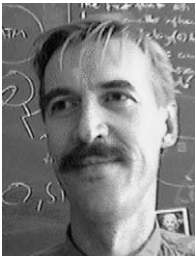
Jean-Yves Le Boudec graduated from Ecole Normale Superieure de Saint-Cloud, Paris, received his doctorate in 1984 from the University of Rennes, France. In 1987 he joined Bell Northern Research, Ottawa, Canada, as a member of scientific staff in the Network and Product Traffic Design Department. In 1988, he joined the IBM Zurich Research Laboratory where he was manager of the Customer Premises Network Department. In 1994 he became professor at EPFL, where he is now full professor. His interests are in the architecture and performance of communication systems.