On the Feasible Scenarios at the Output of a FIFO Server

Vicent Cholvi, Juan Echagüe, and Jean-Yves LeBoudec

***Abstract*— In this paper we consider the case of a FIFO multiplexer fed by flows that are individually constrained by piecewise linear concave arrival curves. We show that, contrary to what happens at the input, at the output not all valid scenarios in accordance with the worst case arrival curves can occur. This implies that taking an iterative approach to characterize the arrival curves at the output when flows pass throughout several FIFO nodes is suboptimal (in the sense that, although valid, they do not necessarily have to be the best arrival curves that can be found).**

***Index Terms*— FIFO, aggregate scheduling, differentiate services, network calculus.**

I. Introduction

Aggregate packet scheduling has attracted a lot of attention in the networking community. For instance, in the Differentiated Services framework [1]–[3], a required per–hop behavior is provided on an aggregate basis. Additionally, front ends to optical switches require aggregated multiplexing if they are to be performed [4]. Thus, there is a need for a deeper understanding of the effects caused by traffic aggregation.

Recent work [1], [5], [6] has shown that the delay bounds with a FIFO network depend on the level of utilization and the number of hops. Moreover, the effect of multiplexing several flows into a FIFO scheduler has been tightly quantified and shown to result in a increased burstiness at the output of the FIFO server [7], [8].

In this paper, we show that, contrary to what happens at the input, not all valid scenarios (which will be formally defined below) can occur at the output of a FIFO server. As a consequence of this, we demostrate that iteratively applying the "optimal" output burstiness bounds when flows pass through several FIFO nodes does not guarantee that the overall burstiness bound will be tight. Such a result has some potential applications, such as its use in the Expedited Forwarding Service (EF) [9]. The goal of the EF (a service which has been developed in the Differentiated Services Working Group of IETF [10]) is to provide an aggregate of flows with some hard delay guarantees by ensuring that, at each hop, the aggregate requiring EF treatment receives a service rate exceeding the total bandwidth requirements of all flows from the aggregate at each hop.

Manuscript received October 10, 2004. The associate editor coordinating the review of this letter and approving it for publication was Carla-Fabiana Chiasserini.

J. Echagüe and V. Cholvi are with the Universitat Jaume I, Spain (e-mail: vcholvi@uji.es).

Jean–Yves Le Boudec is with the Laboratory for Computer Communications and Applications, École Polytechnique Fédérale de Lausanne, Switzerland.

Digital Object Identifier 10.1109/LCOMM.2005.05006.

The rest of the paper is organized as follows. In Section II we give our assumptions and notation. In Section III we show the suboptimality of taking an iterative approach to characterize the worst case arrival curves. Finally, in Section IV we present our conclusions.

II. Preliminaries

Integrated services in the Internet architecture: an overview In this section we describe our model and assumption. We take a fluid approach and 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 [11]

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

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

$$v(t) = \sup\{s \text{ such that } s \leq t \text{ and } A(s) \leq B(t)\} \quad (1)$$

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:

$$B_i(t) = A_i(v(t)) \quad (2)$$

We assume that input flow i is constrained by an arrival curve α_i; in other words [12]

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

Without loss of generality, we can focus on flow $i = 1$ and consider the set of all flows $j \neq i$ 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 the combination of leaky buckets and are common in practice.

Given the above mentioned assumptions, the following result appears in [8].

Theorem 2.1: Consider a FIFO system serving two flows with the above mentioned assumptions. Define

$$\alpha_1^*(x) = \min\{Rx, \alpha_1(x + a_1(x))\} \textit{ for all } x \geq 0 \quad (4)$$

where $a_1(x)$ is the maximum value for a from the set of couples $(a \geq 0, b \geq 0)$ that solve (5):

$$\alpha_1(b+a+x) - \alpha_1(a+x) + \alpha_2(b) - R(a+b) = 0 \quad (5)$$

Then, α_1^* is the best arrival curve for the output flow $B_1(t)$ that can be found under these assumptions.

We call *scenario* to any arbitrary collection of functions, $(A_i(t))_{1 \leq i \leq I}$, that are wide-sense increasing and non-negative, and that satisfy (3). For convenience, when necessary, we use a super-index to identify a scenario.

Whereas all scenarios at the input of the FIFO server that are in accordance with the arrival function given by (3) can, by definition, occur (hereafter called *feasible*), it remains to be seen whether the same happens for the scenarios at the output of the FIFO server (i.e., if all scenarios that are in accordance with the arrival function given by (4) are feasible). The main contribution of this communication is the formal proof that this is not the case.

III. Feasible Scenarios

Before we proceed with our main result, we introduce a preliminary lemma regarding the form of a_1.

Lemma 3.1: For all x, x' such that $0 \leq x' \leq x$ then $a_1(x) \leq a_1(x')$.

Proof: Consider some arbitrary but fixed time interval $[s, t]$. Given a scenario β, we use the notation $s^\beta = v^\beta(s)$ and denote as s'^β the start of the busy period[1] which last, at least, until s^β.

Denote $x = t - s$ and consider the following scenario β:

1) $A_1^\beta(t) - A_1^\beta(s^\beta) = \alpha_1(t - s^\beta)$ and for any other scenario γ such that $A_1^\gamma(t) - A_1^\gamma(s^\gamma) = \alpha_1(t - s^\gamma)$ then $s^\beta \leq s^\gamma$.
2) Flow 1 injects bits in a greedy fashion[2] in time interval $[s^\beta, t]$ and injects $\alpha_1(t - s'^\beta) - \alpha_1(t - s^\beta)$ bits in time interval $[s'^\beta, s^\beta)$.
3) Flow 2 injects $\alpha_1(s^\beta - s'^\beta)$ bits in time interval $[s'^\beta, s^\beta)$ and stops injecting bits after time instant s^β.

It has been also shown in [8] that $a_1(x) = s - s^\beta$. Let us now take a value of $x' = t' - s$ such that $0 \leq x' \leq x$. Define the scenario γ such that:

1) Flow 1 injects bits in a greedy fashion in time interval $[s^\beta, t']$ and injects $\alpha_1(t' - s'^\beta) - \alpha_1(t' - s^\beta)$ bits in time interval $[s'^\beta, s^\beta)$.
2) Flow 2 behaves as in scenario β.

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

Since α_1 is concave then the number of bits injected in the interval $[s'^\beta, m]$ (for all $m : s'^\beta \leq m \leq t'$) will be greater (or equal) in scenario γ than in scenario β. Therefore, s^γ will be lower than (or equal to) s^β and consequently $a_1(x) \leq a_1(x')$. ∎

Fig. 1 provides a numerical application that shows the value of $a_1(x)$.

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

[2]We say than in scenario β flow 1 injects bits in a greedy fashion in time interval $[s, t]$ if $\forall m : s \leq m \leq t$ $(A_1^\beta(m) - A_1^\beta(s) = \alpha_1(m - s))$ (resp. for flow 2). We also extend this definition to the output functions.

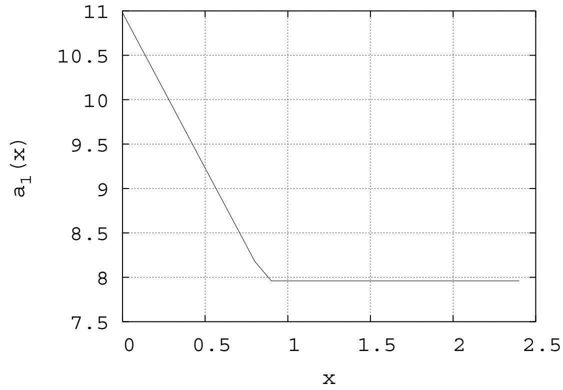

Fig. 1. Input flow 1 has arrival curve $\alpha_1(x) = \min\{10x, 15 + 3x\}$ and input flow 2 has arrival curve $\alpha_2(x) = \min\{8x, 10 + 3x\}$. The server rate is 7. As it can be seen, the value of $a_1(x)$ never grows with x and, for sufficiently large x, it remains constant.

The following theorem shows us that when a FIFO server is considered, not all valid scenarios (at the output) in accordance with the worst case arrival curve can occur (contrary to what happens at the input).

Theorem 3.1: Consider a FIFO server serving two flows (with the assumptions in Section II). Then, at the output, not all valid scenarios in accordance with the worst case arrival curve are feasible.

Proof: By counter–example. Let us focus on a system where $R = 10$, $\alpha_1(x) = \min\{4x, 11 + 3.5x\}$ and $\alpha_2(x) = \min\{7x, 1 + 6x\}$ (which corresponds to the variable bit rate case, or T-SPEC, used by the IETF [13], [14]). Denote as x_1 the point where α_1 changes the value of its linearity (numerically, $x_1 = 22$). This implies that the length of the maximum time interval during which flow 1 can continuously inject bits at the highest rate is x_1.

Take a scenario, denoted β_1^*, in which flow 1 is greedy in time interval $[s, s + x_1]$ at the output. Clearly, by definition of the arrival curve (see (3)), β_1^* is a valid scenario. Assume, by way of contradiction, that it is also feasible. Let us consider what happens at two given time instants at the input of the FIFO server:

- *Time instant* $s + x_1$*:* By solving (5) we have that $a_1(x_1) = 0.5$. Therefore, $Rx_1 > \alpha_1(x_1 + a_1(x_1))$.
  Consequently, by observing the form of α_1^* (see (4)), flow 1 must inject $\alpha_1(x_1 + a_1(x_1))$ bits during time interval $[s - a_1(x_1), s + x_1]$. Since the length of such a time interval is lower than x_1 then, during that period flow 1 must use all its burst. Fig. 2a illustrates this situation.
- *Time instant* $s + x_0$ *with* $x_0 = 21.091$*:* By solving (5) we have that $a_1(x_0) = 0.909$ (note that $x_0 < x_1$ and $a_1(x_1) < a_1(x_0)$, which satisfies Lemma 3.1). Therefore, $Rx_0 > \alpha_1(x_0 + a_1(x_0))$.
  Consequently, by observing the form of α_1^* (see (4)), flow 1 must inject $\alpha_1(x_0 + a_1(x_0))$ bits during time interval $[s - a_1(x_0), s + x_0]$. Since the length of such a time interval is equal to x_1 (note that $a_1(x_0) + x_0 = x_1$) then, during that period flow 1 must inject at its highest rate. Fig. 2B illustrates this situation.

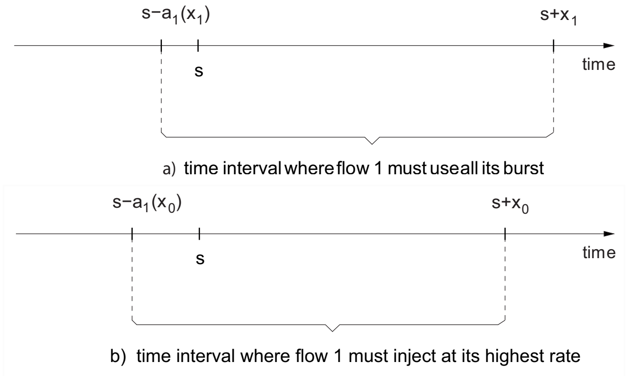

Fig. 2. Example that illustrates the proof of Theorem 3.1: a) situation for time $s + x_1$; b) situation for time $s + x_0$.

As a consequence of the above mentioned requirements, flow 1 cannot use all its burst during time interval $[s - a_1(x_1), s+x_1]$ and inject at its highest rate during time interval $[s - a_1(x_0), s + x_0]$ at the same time.

Therefore, β_1^* cannot be greedy during time interval $[s, s + x_1]$, thus contradicting the hypothesis that it is feasible. ■

This apparent paradox can be explained if we take into account that, in order the output function to reach the maximum value for a particular x, we must have a specific scenario at the input. And that scenario can be incompatible with the scenario needed to reach the maximum value for another different value of x.

As a matter of fact, if we take into account (by Theorem 2.1) that

$$\sup_{\beta,t,x} (B_1^{\beta}(t + x) - B_1^{\beta}(t)) = \alpha_1^*(x)$$

Theorem 3.1 seems to suggest that the β at which the $\sup$ is obtained depends on t and x. This conjecture is strengthened by the fact that it is possible for a flow i to be α-smooth (being sub–additive) with $A_i(t_0) = \alpha(t_0)$ for some time t_0, but $A_i(t) < \alpha(t)$ for all $t < t_0$. For example, by time–reversing a greedy flow (see Chapter 3 in [14]), we can obtain an α-smooth flow which is not greedy.

On the other hand, Theorem 3.1 shows that B_1 is constrained, not only by α_1, but also by some other constraints. Clearly, any feasible scenario will also fulfill that the requirements that the aggregate of the two flows at the output must be bounded by the aggregate of the two flows at the input. One may then ask if this new constraint is enough to characterize the whole set of feasible scenarios. Unfortunately, as the next Corollary shows, this is not the case.

Corollary 1: Consider a FIFO system serving two flows (with the assumptions in Section II). Define $\alpha_0^*(x) = \min\{Rx, \alpha_1(x) + \alpha_2(x)\}$ for all $x \geq 0$. Then, at the output, not all valid scenarios in which flow 1 is in accordance with $\alpha_1^*(x)$, flow 2 is in accordance with $\alpha_2^*(x)$ and the aggregate of the two flows is in accordance with $\alpha_0^*(x)$ are necessarily feasible.

The proof is immediate since the scenario β used in the proof of Theorem 3.1 fulfills the above mentioned constraints and, as has been shown, is not feasible.

IV. Conclusions

In the previous section it has been shown that the scenarios at the output of a FIFO server are, in general, more restrictive that those at the input. Whereas this result may seem straightforward (although *a posteriori*), we provided a formal evidence, which explains the well-known inefficiency involved in finding performance bounds by iteratively applying output burstiness bounds [14], [15].

References

[1] A. Charny and J.-Y. Le Boudec, "Delay bounds in a network with aggregate scheduling," *First International Workshop on Quality of future Internet Services*, pp. 1-13, Sept. 2000.

[2] J. C. R. Bennett, K. Benson, A. Charny, W. F. Courtney, and J.-Y. Le Boudec, "Delay jitter bounds and packet scale rate guarantee for expedited forwarding," *ACM/IEEE Trans. Networking*, vol. 10, pp. 529–540, Aug. 2002.

[3] V. Firoiu, J.-Y. Le Boudec, D. Towsley, and Z.-L. Zhang, "Theories and models for internet quality of service," in *Proc. IEEE*, vol. 90, pp. 1565–1591, Sept. 2002.

[4] L. Thylén, G. Karlsson, and O. Nilsson, "Switching technologies for future guided wave optical networks: potentials and limitations of photonics and electronics," *IEEE Commun. Mag.*, vol. 34, pp. 106-113, Feb. 1996.

[5] Z.-L. Zhang and Z. Duan, "Fundamental trade-offs in aggregate packet scheduling," *SPIE*, vol. 4526, Aug. 2001.

[6] J. Echagüe, V. Cholvi, and A. Fernández, "Universal stability results for low rate adversaries in packet switched networks," *IEEE Commun. Lett.*, vol. 7, pp. 578-580, Dec. 2003.

[7] M. Fidler, V. Sander, and W. Klimala, "Traffic shaping in aggregate-based networks: implementation and analysis," to be published in *Elsevier Computer Communications Journal*, special issue on Quality of Future Internet Services 2004.

[8] V. Cholvi, J. Echagüe, and J.-Y. Le Boudec, "Worst case burstiness increase due to FIFO multiplexing," *Performance Evaluation*, vol. 49, pp. 491-506, Sept. 2002.

[9] V. Jacobson, K. Nichols, and K. Poduri, "An expedited forwarding PHB," June 1999, RFC 2598, IETF.

[10] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An architecture for differentiated services," Dec. 1998, RFC 2475, IETF.

[11] C. S. Chang, "On deterministic traffic regulation and service guarantee: a systematic approach by filtering," *IEEE Trans. Inform. Theory*, vol. 44, pp. 1096-1107, Aug. 1998.

[12] R. L. Cruz, "Quality of service guarantees in virtual circuit switched networks," *IEEE J. Sel. Areas Commun.*, vol. 13, pp. 1048-1056, Aug. 1995.

[13] B. Braden, D. Clark, and S. Shenker, "Integrated services in the Internet architecture: an overview," June 1994, RFC 1633, IETF.

[14] J.-L. Le Boudec and P. Thiran, *Network Calculus*, Springer Verlag Lecture Notes in Computer Science, vol. 2050 (available online at http://icawww.epfl.ch), July 2001.

[15] M. Fidler, "Extending the network calculus pay bursts only once principle to aggregate scheduling," in *QoS-IP*, vol. LNCS 2601, Feb. 2003, pp. 19-34.