Marco Listanti (*), Fabio Ricciato (*), Stefano Salsano (**)

An Admission Control with Quantitative Statistical QoS Guarantees for *Expedited Forwarding*

(*) Infocom Dept. – University of Rome "La Sapienza" – Italy
(**) CoRiTeL – Rome - Italy
email: {ricciato,salsano}@coritel.it

Abstract

This paper presents an admission control framework for *Expedited Forwarding* traffic in a Differentiated Service network. The aim is to overcome the limitations, in terms of achievable efficiency, which are proper of a deterministic "worst-case" based on the zero loss assumption. An admission control procedure is defined which provides quantifiable end-to-end QoS guarantees in terms of maximum delay and per-flow loss probability. The flow admission control procedure relies on the analytical bounds for the packet loss probabilities which are evaluated for each node in the flow path. The degradation of the statistical characteristics of the flow due to packet clumping through the network is taken into account. The derivation of the per-flow loss probability bound is based on an extension of an analytical result due to Beneš. In order to exploit it in a range of cases of practical interest, a dropping device is introduced before the FIFO multiplexer. The purpose of the dropper is to discard packets in order to avoid conflicts at burst scale. The loss introduced by the dropper is properly taken into account in the evaluation of the bound for the node loss probability. The interactions between EF traffic and non-EF packets are also considered. Finally, a comparison between analytical bounds and actual performance results obtained by simulations is presented. The results show that the requested QoS targets are largely met and that the achievable efficiency is much higher than that derived from worst case allocation.

Reference Author

Fabio Ricciato
INFOCOM Dept.
Via Eudossiana, 18
00184 Roma

Phone. +39 06 47852302
Fax: +39 06 48906114
E-mail: ricciato@coritel.it

1. INTRODUCTION

At present a growing number of Internet applications require some kind of Quality of Service (QoS) guarantees, such as delay constraints or low packet loss. In the framework of the actual Internet architecture only *best effort* service can be provided, so that additional mechanisms have to be defined in order to bring QoS into the Internet. In this context, two models are under study within the IETF: the Integrated Services and the Differentiated Services models.

The Integrated Services [1] model adopts a per-flow approach, that means each traffic flow is handled separately at each router. This way resources can be allocated individually to each flow and the QoS can be provided on a per-flow basis, i.e. each flow can request a specific QoS. It has been recognized that such a per-flow approach is affected by heavy scalability problems.

The Differentiated Services model is aimed at providing QoS on a per-aggregate basis: a limited set of service classes (called PHB in the Differentiated Services terminology) is supported, each class being associated to a specific QoS level. Any generic traffic flow logically access one PHB and the relevant QoS. The internal routers handle packets according to the PHB identifier, and do not distinguish the individual flows. Then resources are allocated on a per-class basis.

Going from the Integrated Services to the Differentiated Services model the system complexity moves from the packet handling mechanism, which has to be implemented at every network node, towards the resources allocation function, which is logically located at the network edge. Thus the Differentiated Services model is well scalable, but providing quantifiable end-to-end QoS guarantees without per-flow packet handling is generally a very challenging task.

In order to provide end-to-end QoS two preliminary steps must be considered: i) the QoS parameters definition and ii) the QoS performance evaluation given a network load state. The latter point in turns depends on the input traffic characterization. Note that evaluating the QoS performances at the generic k-th stage along a cascade of nodes is generally a more challenging task than evaluating the same parameters at the first stage, as the modifications introduced on the traffic characteristics by the previous k-1 multiplexing stages should be taken into account.

This paper deals with QoS provisioning for the Expedited Forwarding (EF) PHB, which is intended to support real time applications. The aim is to provide an effective flow admission control scheme able to deliver quantifiable end-to-end QoS guarantees to EF. As suggested in [5], we assume that packets belonging to EF-PHB are served at every node with higher non-preemptive priority over all other PHBs packets.

The traditional approach to the EF performances is based on a ‘deterministic worst case’ (DWC) analysis: a zero-loss scenario is assumed (infinite buffer) and the maximum queuing delay is adopted as QoS parameter. The DWC analysis leaves the statistical characteristics of the traffic flows out of consideration. This is a determinant simplification, but at the same time constitutes the main limitation of the DWC approach, as it excludes the possibility to exploit the traffic structure to achieve statistical gain in resources allocation.

In [9] [10] and [11] analytical worst cases are investigated. They aim at evaluating the worst case delay for each single flow, given the traffic matrix (number and rates of flows) and the network topology. All the results show that under the no loss hypothesis the increase of the delay is unacceptable even after few hops. The general procedure is to consider the worst case arrival pattern to the node, in order to compute the buffer space needed at the node and evaluate a worst case pattern on the output link. The “output” worst case patterns can be worst-case combined in the next stages and so on. In [9] strictly CBR flows are considered as input, while [10] and [11] take also into account that the input flows are not necessarily constant bit rate and that a “clumping” of packets on the output link can occur. Due to the clumping effect the short time arrival rates to the next stage nodes can be higher, and the results are drastically worse.

If the worst case analysis is taken as the admission criteria, the utilization by EF traffic must be kept unacceptably low. On the one hand this should be the only solution if deterministic zero loss guarantees are sought. On the other hand the common sense and the simulation results show that the real behavior is far from the worst case, making it not feasible as an admission control rule. In fact, the simulations and the first EF real trials show that quite good performances can be achieved with reasonable efficiency, but a comprehensive analytical model is still missing.

In order to overcome such limitations and increase the utilization efficiency the deterministic zero-loss assumption has to be dropped. Allowing for a little loss probability makes it possible to achieve a higher utilization for EF traffic, by exploiting the statistical characteristics of the traffic, provided that such a loss can be quantified and controlled. In such a framework, we propose an admission scheme for EF able to provide quantitative QoS guarantees in terms of per-flow loss probability and maximum delay. The underlying analytical approach is also presented. It takes into account the sources activities and can manage a heterogeneous scenario with different traffic sources.

The flow admission procedure is based on the evaluation of an analytical bound for the packet loss probability at each node crossed by the flow. These bounds are determined as functions of the node structural parameters (buffer size, link capacity) and of two traffic parameters for each input flow: the average and the peak packet rates. The degradation of the statistical characteristics of a flow along its path is bounded taking into account the structural parameters of the previously crossed nodes, independently of the characteristics of the other flows loading the network. The typical approaches to statistical admission control described in the literature ([12] provides a comprehensive survey) fail to take into account this degradation and limit their analysis to a single stage multiplexer.

The derivation of a bound for the loss probability is based on an analytical achievement found in [15] and [16], which in turn is derived from a general result due to Beneš: substantially it provides a closed formula for the first order statistic of the queue process in a FIFO multiplexer, under the hypothesis that conflicts occur at the packet level only. In order to extend the Beneš approach to consider realistic cases where burst scale congestion can appear, a device (called *dropper*) is introduced before the multiplexer. The purpose of the dropper is to discard packets in order to avoid conflicts at the burst scale at the successive queue, so to be able to bound it by means of the above mentioned formula. The loss probability introduced eventually by the dropper is properly taken into account in the evaluation of the node total loss budget. Finally, the interactions between EF traffic and non-EF one are considered and a simple approach to take into account the relevant impairments on the EF traffic is proposed.

An extensive performance study has been carried out. The comparison between analytical and simulation results shows that the requested QoS targets are largely met. The section 2 presents the general framework of the proposed admission control mechanism, whereas the analytical approach the it is based on is discussed in section 3. The section 4 describes the admission control procedure and comments its implementation aspects. Finally, in section 5 the performance results are discussed.

2. ARCHITECTURE AND ADMISSION CONTROL FRAMEWORK FOR EF SERVICE

The support of EF PHB in an IP DiffServ domain requires the definition of an Admission Control procedure in order to assure the needed QoS to the incoming EF traffic. The procedure can be statically performed at the network provisioning level, or alternatively new flows can dynamically be added and removed. Here we will mainly consider the "dynamic" option; the static one can be seen as a simplified scenario. The network architecture is depicted in Fig. 1. A "Bandwidth Broker" (BB) must be in charge of controlling allocation within the Differentiated Services domain. An EF flow is submitted to an admission control procedure handled by the BB. The DiffServ domain interconnects a number of peripheral domains supporting QoS via heterogeneous mechanisms (e.g. IntServ-RSVP or DiffServ) [13] [14] . Within the Differentiated Service domain, the packets will cross several nodes, so that "end-to-end" (i.e. from the ingress point to the egress point) performance parameters have to be taken into account.

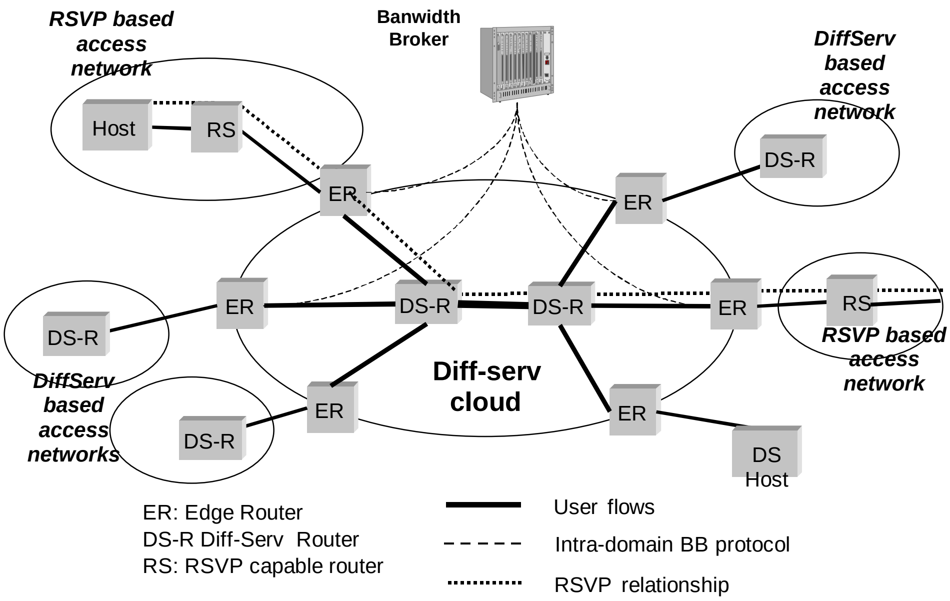

Fig. 1 – Architectural scenario for a centralized admission control in a Differentiated Service domain

We assume that, as suggested in [5], the EF packets are scheduled with non-preemptive priority over non-EF ones at every node. The impact of non-EF traffic can be considered as an impairment to the EF packets, due to the additional delay possibly caused by the non-EF packets that are completing transmission when an EF packets enters the node. The effect of this additional delay is initially neglected, so that the EF traffic is assumed to cross network nodes operating with a simple FIFO discipline. A simple approach to take into account the impact of non-EF traffic is given in section 3.4.

A scenario with fixed packet size (L bits) input traffic is considered. The queue buffers at each node are finite, so that for each packet the maximum end-to-end delay can be trivially derived as the sum of the maximum queuing time along its path.

We consider the generic i-th packet flow through the DiffServ network. Let M_i be the set of the network nodes crossed by the i-th packet flow, and let $M_i=|\mathsf{M}_i|$ be the cardinality of M_i. Let us assume that the elements of M_i are numbered according to the packet crossing order.

We assume that each network node Ω is statically assigned a maximum admitted packet loss probability Π_Ω.

Let $\Pi_i^{(k)}$ (k=1,.., M_i) be the assigned maximum admitted packet loss probability of the k-th node crossed by i-th packet flow along its path, so that any packet of the i-th flow crossing the k-th node will experience a loss probability not greater than $\Pi_i^{(k)}$. Note that a generic network node Ω is associated by a set of couples (i,k) denoted by H_Ω. Therefore, $\Pi_i^{(k)}$ is identical for all of the couples (i,k) belonging to the set H_Ω, i.e. $\Pi_i^{(k)} = \Pi_\Omega \quad \forall (i,k) \in \mathsf{H}_\Omega$. In the following, the notation $X_i^{(k)}$ and X_Ω will be interchangeably used to denote the generic parameter X relevant to the node Ω, being implicit that $(i,k) \in \mathsf{H}_\Omega$.

Denoting with $p_i^{(k)}$ the actual loss probability of the generic i-th flow through the k-th node, it will be guaranteed that:

$$p_i^{(k)} \le \Pi_i^{(k)} \qquad k = 1,..,M_i \qquad \forall i\,. \tag{1}$$

The BB will guarantee that (1) holds at any node inside the DiffServ domain, by opportunely admitting or rejecting packets flows. In order to perform such a decision, the BB must be able to calculate the per-flow loss probabilities $p_i^{(k)}$ at any node, or at least an upper bound for them. Suppose that an upper bound of $p_i^{(k)}$ is computable as a function $f(\cdot)$ of: i) a proper set of node parameters, denoted by G; ii) a proper set of traffic parameters related to all the flows feeding the node; this set will be denoted by F. Formally:

$$p_i^{(k)} \le f\left(\mathsf{G}_i^{(k)}, \mathsf{F}_i^{(k)}\right) \qquad k = 1,..,M_i \qquad \forall i\,. \tag{2}$$

When a new flow requests EF service, the BB determines the path and checks that its admission preserves condition (2) at every node along the path. Such a control can be performed by computing for each of the path nodes the upper bound (2) for the new input flows set, that is the union between the former set and the new flow. The flow can thus be admitted only if the new computed bound remains below the given threshold $\Pi_i^{(k)}$ for every path nodes, formally:

$$f\left(\mathrm{G}_i^{(k)},\mathrm{F}_i^{(k)}\right)\leq \Pi_i^{(k)} \qquad k=1,..,M_i \qquad \forall i\,. \tag{3}$$

The flow will be rejected if, for at least one node of the path, the constraint (3) is not met or the bound (2) is not computable.

In order to follow this approach, the values of the traffic parameters belonging to the set F have to be independent of the current network load, so that they can be evaluated just once in the flow admission phase and do not change during the flow lifetime. As it will be shown hereafter, the presented procedure meets such a requirement, so as to allow each node Ω to be considered as a “black box”, that introduces on the flows crossing it an aleatory queuing delay, bounded by the buffer length, and a loss probability, bounded by Π_Ω.

The end-to-end QoS parameters for any flow in the network are then trivially derived by composing the per node QoS parameters along the path. For the generic *i*-th flow, denote by $C_i^{(k)}$ and $B_i^{(k)}$ (k=1,.. M_i) the output link capacity (bit/s) and the buffer length (bits) of the *k*-th node along its path, respectively. The maximum queuing time experienced by packets of the *i*-th flow at the *k*-th node of its path is

$$d_i^{(k)}=\frac{B_i^{(k)}}{C_i^{(k)}} \tag{4}$$

The end-to-end maximum delay variation for *i*-th flow, denoted by Δ_i^{tot}, is given by the sum of the maximum queuing delay along the path, formally:

$$\Delta_i^{tot}=\sum_{k=1}^{M_i} d_i^{(k)}. \tag{5}$$

An upper bound Π_i^{tot} of the total loss probability experienced by a packet of the *i*-th flow through the network is given, except for some very artificial scenarios, by the following sum:

$$\Pi_i^{tot}=\sum_{k=1}^{M_i}\Pi_i^{(k)} \tag{6}$$

Let's now focus on the determination of function $f(\mathsf{G},\mathsf{F})$. It must be noted that in general the statistical characteristics of a traffic flow through the network are modified by each multiplexing process, so that at any multiplexing stage the values of the flow parameters will change accordingly. The approach proposed here considers just two parameters for each traffic flow, whose modification at each multiplexing stage can be easily taken into account. The considered flow parameters are the average and the peak packet rate.

The average packet rate is defined as the long term averaged flow packet rate. The peak packet rate is defined as the inverse of the minimum arrival interval between two consecutive packets of the same flow. We assume that the average flow packet rate does not change with the multiplexing stage inside the network. In fact it is only diminished by packet loss events, whose effects can be neglected because the involved values of the admitted loss probabilities $\Pi_i^{(k)}$ are hopefully small in practical cases. On the other hand, we can account for modifications in the flow peak packet rate at the *k*-th node (*k*=1,..M$_i$) by simply considering the maximum delay variation introduced by the previous *k*-1 nodes. Denote with $R_i^{(k)}$ and $P_i^{(k)}$ the average and peak packet rate of the *i*-th flow at the *k*-th node (*k*=1,..M$_i$), respectively. Moreover, let $t_i^{(k)} = L/C_i^{(k)}$ be the packet transmission time, and $T_i^{(k)}$ be the minimum arrival interval between two consecutive packets of the *i*-th flow at the *k*-th node. It holds that:

$$
\begin{aligned}
k = 1: &\qquad T_i^{(1)} = L / P_i^{(1)} \\
k \geq 2: &\qquad T_i^{(k+1)} = \max\left\{ T_i^{(1)} - \sum_{m=1}^{k-1} d_i^{(m)}, t_i^{(k)} \right\}.
\end{aligned}
\qquad (7)
$$

Thus, we can write for *k*=1,..M$_i$:

$$
\begin{aligned}
R_i^{(k)} &= R_i^{(1)} \\
P_i^{(k)} &= \frac{1}{T_i^{(k)}}
\end{aligned}
\qquad (8)
$$

It is worth noting that the values of the average and peak packet rate as given by (8) do not depend on other traffic flows parameters but only on physical parameters of the crossed nodes (i.e. buffer size and link capacity).

By applying (8) and knowing the node parameters, the BB is able to evaluate all the arguments of the function $f(\mathsf{G},\mathsf{F})$ for all the nodes in the path of the flow under test. In the next section function $f(\mathsf{G},\mathsf{F})$ will be analytically derived.

3. AN ANALITICAL BOUND FOR THE LOSS PROBABILITY IN A SINGLE NODE

The analytical result derived in [15] and [16], dealing with multiplexing of independent On/Off packet streams onto an ATM multiplexer, is the starting point for our work. This result is in turn based on a theorem on the distribution of work in a G/G/1 queue originally derived by Beneš [17]. Therefore, throughout the paper, we refer to it as "Beneš result". In section 3.1 results from [15] and [16] are summarized, omitting the complete demonstration and highlighting the stringent conditions of applicability. In section 3.2 we investigate how to overcome these limitations in order to exploit the Beneš result in cases of more practical interest so as to achieve a statistical gain. In section 3.3 an upper bound for the packet loss probability in a multiplexer stage is evaluated; finally, the interactions between EF traffic and non-EF one are investigated in section 3.4.

In order to simplify the notations used in this section, as we consider a single multiplexing stage, the indication of the index (k) of the node crossed by the i-th flow will be omitted.

3.1. Queuing analysis of a multiplexing stage: review of the Beneš result

A multiplexer with infinite buffer space, loaded by independent flows of fixed packet size is considered. Let:

- C: the output link capacity (bit/s);
- I: the total number of flows entering the multiplexer;
- R_i: the average rate (bit/s) of the i-th flow ($i=1,2..I$);
- P_i: the peak rate (bit/s) of the i-th flow ($i=1,2..I$);
- $a_i=R_i/P_i$: the activity of the i-th flow ($i=1,2..I$);
- $T_i=L/P_i$: the minimum inter-arrival period between two consecutive packets of the i-th flow ($i=1,2..I$);
- $t=L/C$: the packet service time;
- $A_i(t_1,t_2)$: the number of arrivals from the generic i-th flow in the interval $[t_1,t_2)$.

Let's consider a time interval of duration D smaller than the shortest per-flow inter-arrival period, viz. $D\leq\min_i\{T_i\}$. Let $P_D(N)$ be the probability mass function of the total number of arrivals in a time interval of length D, viz.:

$$P_D(N) = \Pr\left\{\sum_{i=1}^{I} A_i(t-D,t) = N\right\} \quad (9)$$

As the flows feeding the multiplexer are independent, $P_D(N)$ is given by the convolution between the probability mass functions of the number of per-flow arrivals $g_i(k) = \Pr\{A_i(t-D,t) = k\}$:

$$P_D(N) = g_1(k) \otimes g_2(k) \otimes .. g_I(k) = \underset{i=1..I}{conv}\{g_i(k)\} \quad (10)$$

From the definition of *D*, it derives that it is not possible to collect more than one arrival from each flow in a time interval of length *D*, so that any $g_i(k)$ is non zero only for *k*=0,1. As the probability to find one packet belonging to the *i*-th flow in a generic interval of length *D* is equal to the ratio between *D* and the mean per-flow inter-arrival period, we have:

$$g_i(1) = \frac{D}{L/R_i} = \frac{D}{T_i/a_i}, \quad g_i(0) = 1 - g_i(1). \quad (11)$$

Denote by *V* the *virtual waiting time* at a generic instant *t*, i.e. the amount of work in the system at instant *t* expressed in terms of the time required to finish it at rate *C*. For such a system, under the hypothesis

$$I < D/t, \quad (12)$$

the exact expression of the complementary distribution function of *V* is

$$\Pr\{V > x\} = \sum_{N=0}^{I} P_D(N) \cdot Q_D^N(x), \quad (13)$$

wherein $Q_D^N(x)$ represents the probability to have more than *x* backlog at instant *t* given *N* arrivals in interval [*t-D,t*). $Q_D^N(x)$ is computable by:

$$Q_D^N(x) = \sum_{\frac{x}{t} < n \le N} \left[\binom{N}{n} \cdot \left(\frac{n \cdot t - x}{D}\right)^n \cdot \left(1 - \frac{n \cdot t - x}{D}\right)^{N-n} \cdot \left(\frac{D - N \cdot t + x}{D - n \cdot t + x}\right) \right] \quad (14)$$

The proof of (13) , given in [15], is fundamentally based on the two following properties of the overall arrival process at the queue.

Property 1: *in any interval of length D, the total amount of packets entering the multiplexer is less than the number of packets that can be processed in the same interval*;

Property 2: *in any interval of duration D, the packet arrival instants are uniformly distributed in [t-D,t).*

The Property 1 is implied by condition (12) which limits the number of flows in the system. It guarantees that the backlog at the generic instant t does not depend on the arrivals before t-D. The Property 2 derives from the fact that, in a generic interval of length D, any arrival belongs to a different flow ($D \le \min_i\{T_i\}$) and from the reciprocal independence of the flows.

It is worth noting that, condition (12) on the maximum number of flows guarantees that queuing is only due to conflicts at "packet scale", and eliminates the possibility of conflicts at "burst scale", i.e. at any time it is not possible to find in the queue more than one packet for each flow. Thanks to such a condition the first order statistic of the queue process is independent from the correlation structure of the traffic flows and can be expressed by (13), that only requires the average per-flow arrival intervals to be known.

Unfortunately, direct application of (12) as an admission control rule has a poor practical appeal; in fact if we examine its equivalent form:

$$I \cdot \max_i \{P_i\} < C \,, \tag{15}$$

it clearly appears that (15) would represent the deterministic allocation condition in the case that all the I flows were considered to have the same peak rate equal to the highest one. It is evident that (15) would lead to very poor bandwidth utilization efficiency in a heterogeneous scenario with flows with low activities and/or high peak rates.

3.2. Avoiding burst scale congestion via packet dropping: queuing analysis

In order to apply the Beneš result in a multiplexer in which condition (12), or equivalently (14), does not hold, let us consider the scheme of Fig. 2.

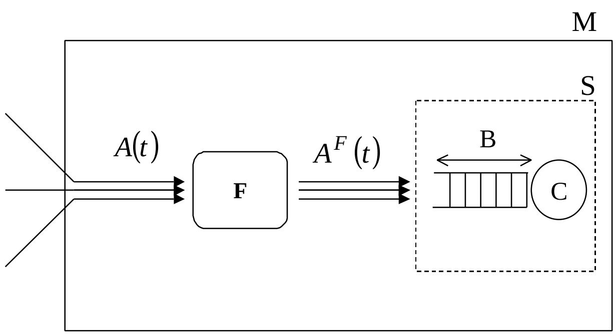

Fig. 2: The proposed scheme of the Multiplexer *M*, showing the dropper F and the queuing subsystem S

The input arrival process $A(t)$ to the multiplexer M, constituted by the superposition of I independent flows, is "filtered" by the dropper F. The resulting process $A^F(t)$ is fed to a queuing subsystem S constituted by a buffer of length B and a link with capacity C.

The dropper F has to guarantee that the arrival process $A^F(t)$ to the subsystem S satisfies Property 1. Packets arriving to F are either admitted (i.e. forwarded to S with no delay) or discarded. Let's define a time window of duration D, called *dropping window*; a packet arriving at a generic instant t is admitted by F if and only if the total number of already admitted packets in the time window $[t-D,t)$ is less than N_F, where N_F is the maximum integer such that:

$$N_F < D/t. \tag{16}$$

Hence, the maximum number of arrivals to the system S in any time interval of duration D is less than the number of packets that can be processed by the system in the same interval, so the Property 1 is satisfied. The duration D of the dropping window is chosen according to the following constraint:

$$D < \min_i \{T_i\} \tag{17}$$

The (16) assures that in a generic interval of length D any arrival of the process $A^F(t)$ belongs to a different flow.

We assume that the packet arrival instants of the process $A^F(t)$ are uniformly distributed in a time window D (Property 2) so neglecting the selective effect of dropper F, i.e. *which* packets are discarded in the dropping window. Therefore, the Beneš result can be applied to the queue system S of Fig. 2.

Comparing this situation with that described in the previous section, we note that the role of the dropper is to avoid the "burst level" congestion, i.e. at a time scale greater than D. Obviously the additional loss introduced by the dropper has to be controlled in order to meet the total loss requirement for the multiplexer. An analytical upper bound of the loss probability at the dropper will be derived in the next section.

Let's assume the buffer size B of the subsystem S is infinite. The results obtained in the following will be utilized to evaluate the loss probability in case of finite buffer.

Let $V^F(t)$ be the *virtual waiting time* in the queue system S at a generic instant t. In order to compute the complementary distribution function of $V^F(t)$, the event $\{V^F(t) > x\}$ can be partitioned with respect to the total number N of arrivals at the queue in [t-D,t); N can not be larger than N_F because of the dropper action. Then it holds that:

$$\Pr\left\{V^F(t) > x\right\} = \sum_{N=0}^{N_F} \Pr\left\{V^F(t) > x \mid A^F(t-D,t) = N\right\} \cdot \Pr\left\{A^F(t-D,t) = N\right\}. \tag{18}$$

Let $P_D^F(N)$ be the probability mass function of $A^F(t-D,t)$, that is the total number of arrivals to the system S in a time window of length D. Similarly to (13), equation (18) can be written:

$$\Pr\left\{V^F(t) > x\right\} = \sum_{N=0}^{N_F} P_D^F(N) \cdot Q_D^N(x). \tag{19}$$

The exact form of $P_D^F(N)$ depends on the correlation structure of the process $A(t)$, which in general is unknown. Therefore the complementary distribution function of $V^F(t)$ cannot be exactly computed. However, it is possible to obtain an upper bound for it. Let $P_D^W(N)$ be the probability mass function of the variable $\min\left(A(t-D,t), N_F\right)$, viz.:

$$P_D^W(N) = \begin{cases} P_D(N) & 0 \le N \le N_F - 1 \\ \sum_{m \ge N_F} P_D(m) & N = N_F \\ 0 & N > N_F \end{cases} \tag{20}$$

As proved in [18], the required bound is:

$$\Pr\left\{V^F(t) > x\right\} \le \sum_{N=0}^{N_F} P_D^W(N) \cdot Q_D^N(x), \tag{21}$$

Thus, for the system of Fig. 2, eq. (21) provides an upper bound for the complementary distribution function of the amount of unfinished work in case of infinite buffer. Eq. (21) can be also intuitively justified by observing, as a matter of example, the Fig. 3 in which a realization of the arrival process at the dropper F (dotted line) and of the corresponding departure process (dashed line) is represented. As the process $\min\left(A(t-D,t), N_F\right)$ (continuos line) is always above the actual departure process, it is intuitive that a conservative approximation is obtained by considering the next FIFO queue fed by the process $\min\left(A(t-D,t), N_F\right)$.

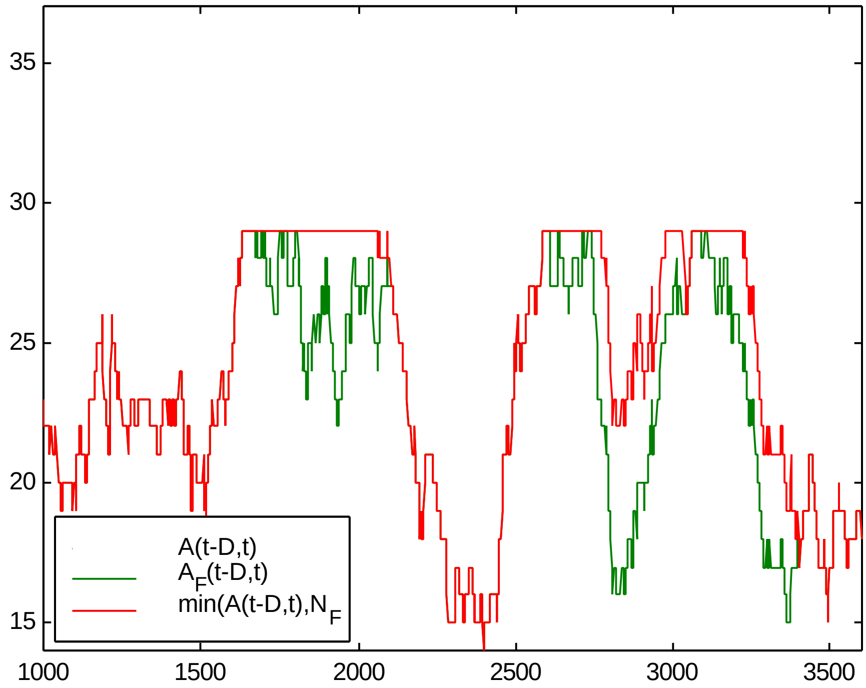

Fig. 3: Realization of arrival process and departure processes at the dropper *F*

The consistency of the bound (21) has been verified by simulations. It has been considered a link of capacity C=10 Mb/s, fed by 804 markovian On/Off flows with peak rate P=32 Kb/s and activity a=0.35. The packet length is L=128 bytes. The dropping window interval is D=31.98 msec and N_F=312. The percentage of packets dropped by F has been $4.1*10^{-4}$. In Fig. 4 the complementary distribution function of $V^F(t)$ obtained by simulation has been compared to the analytical bound (21). It can be noted that the analytical bound tightly approaches the actual system behavior. The distance between the two curves strictly depends on the loss introduced by the dropper.

3.3. Upper bound of the Loss Probability

The filtering of the arrival process $A(t)$ by means of the dropper F allowed the result (13) to be exploited by removing the constraint (12) on the number of flows accessing the multiplexer. In practice, this result allows the definition of a statistical allocation scheme in which the constraint (12) must be replaced by a constraint on the packet loss probability at the multiplexer M.

Now our aim is to give an upper bound for the packet loss probability p_i experienced by the i-th flow through the multiplexer M.

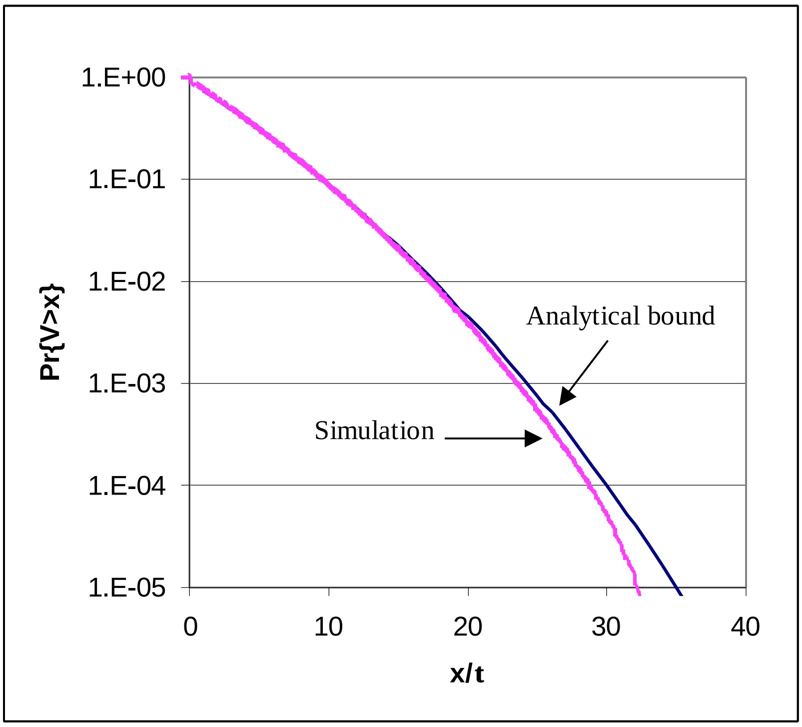

Fig. 4: Tail probability for a single stage: analytical bound vs. simulation results

Let us consider the multiplexing stage shown in Fig. 2 with a finite buffer of size *B* (bits). Let $V_B^F(t)$ be the backlog in such a system, while $V^F(t)$ represents the backlog in a homologous system with infinite buffer. A packet arriving at the multiplexer *M* at instant *t* can be lost:

- in the dropper *F*, if the total number of packets admitted by *F* in the interval [*t*-*D*,*t*) equals N_F, viz. $A^F(t-D,t)=N_F$;
- in the buffer *B*, if the packet is not dropped by *F* and if the unfinished work in the sub-system *S* at the instant *t* exceeds the maximum depletion time of the buffer *B*, viz. $V_B^F(t)>B/C$.

To simplify notations we will denote the event "a packet of the i-th flow arrives at the multiplexer *M* at instant t" with $\Psi(i,t)$. It holds that:

$$p_i = \Pr\{(A^F(t-D,t)=N_F \;\; or \;\; V_B^F(t)>B/C) \mid \Psi(i,t)\}. \qquad (22)$$

By defining a_i and b_i as follows:

$$a_i = \Pr\{A^F(t-D,t)=N_F \mid \Psi(i,t)\},$$

$$b_i = \Pr\{V_B^F(t)>B/C \mid \Psi(i,t)\},$$

From (22) , we can derive a bound for the loss probability

$$p_i \le a_i + b_i \,. \tag{23}$$

Note that a_i represents the packet loss probability at dropper F; b_i is the probability that the buffer of the sub-system S is full when a packet enters the multiplexer M, so b_i can be considered as an upper bound of the loss probability at the buffer B.

In the following upper bounds for a_i and b_i are determined.

Let $P_{D,i}(N)$ be the probability mass function of the total number of arrivals to the system in $[t\text{-}D,t)$ given an arrival to multiplexer M of a packet of the i-th flow at time t, viz. $P_{D,i}(N) = \Pr\{A(t-D,t) = N \mid \Psi(i,t)\}$. $P_{D,i}(N)$ can be easily evaluated by noting that, since packet flows are independent and the interval D has been chosen according to constraint (17), the event $\Psi(i,t)$ excludes any arrival at the multiplexer M from the same i-th flow in the interval $[t\text{-}D,t)$. Thus $P_{D,i}(N)$ is given by the convolution of the probability mass functions $g_m(k)$ relevant to all the flows but the i-th one:

$$P_{D,i}(N) = \underset{\substack{m=1..I \\ m\neq i}}{conv}\{g_m(k)\}. \tag{24}$$

By recalling that $P_D(N)$ is defined as the probability mass function of the total number of arrivals to the multiplexer M in a time interval of length D, we can write:

$$P_D(N) = P_{D,i}(k) \otimes g_i(k). \tag{25}$$

We define $P_{D,i}^{W}(N) = \Pr\{\min(A(t-D,t), N_F) = N \mid \Psi(i,t)\}$. The expression of $P_{D,i}^{W}(N)$ can be derived from (20) by replacing $P_D(N)$ with $P_{D,i}(N)$. From (24) and (25) we can straightforwardly derive:

$$\sum_{m=N}^{N_F} P_{D,i}^{W}(N) < \sum_{m=N}^{N_F} P_D^{W}(N). \tag{26}$$

As the event "N_F packets admitted by F in $[t\text{-}D,t)$" implies that at least N_F packets have arrived to the system in the same interval, the term $\Pr\{A^F(t-D,t) = N_F \mid \Psi(i,t)\}$ can be upper bounded by $\Pr\{A(t-D,t) \ge N_F \mid \Psi(i,t)\}$, so that it holds:

$$a_i \le \sum_{m \ge N_F} P_{D,i}(m) = P_{D,i}^{W}(N_F). \tag{27}$$

Let's now derive an upper bound for b_i.

By observing that, at any instant, whichever the input process is, the backlog for the finite buffer system is not larger than the backlog for the homologous system with infinite buffer, we have:

$$\Pr\left\{V_B^F(t) > x \mid \Psi(i,t)\right\} \le \Pr\left\{V^F(t) > x \mid \Psi(i,t)\right\}. \tag{28}$$

Similarly to equation (18) the right hand term can be written as follows:

$$\Pr\left\{V^F(t) > x \mid \Psi(i,t)\right\} = \sum_{N=0}^{N_F} \Pr\left\{V^F(t) > x \mid A^F(t-D,t) = N, \Psi(i,t)\right\} \cdot \Pr\left\{A^F(t-D,t) = N \mid \Psi(i,t)\right\} \tag{29}$$

As the event $\Psi(i,t)$ does not influence the queue state at instant t, we can write $\Pr\left\{V^F(t) > x \mid A^F(t-D,t) = N, \Psi(i,t)\right\} = Q_D^N(x)$. By following the same procedure adopted to derive the relation (21), it is possible to derive the following relation from (29):

$$\Pr\left\{V^F(t) > x \mid \Psi(i,t)\right\} \le \sum_{N=0}^{N_F} P_{D,i}^W(N) \cdot Q_D^N(x), \tag{30}$$

The (30) conduces to the following upper bound for b_i:

$$b_i = \Pr\left\{V_B^F(t) > B/C \mid \Psi(i,t)\right\} \le \sum_{N=0}^{N_F} P_{D,i}^W(N) \cdot Q_D^N(B/C) \tag{31}$$

By summing the bounds (27) and (31) for a_i and b_i we obtain the following upper bound for the packet drop probability for the generic i-th flow in the multiplexer M:

$$p_i \le P_{D,i}^W(N_F) + \sum_{N=0}^{N_F} P_{D,i}^W(N) \cdot Q_D^N(B/C). \tag{32}$$

Relation (32) gives an upper bound to the loss probability for each input flow. From a practical point of view, it could be convenient to have a single bound for all the different per-flow loss probabilities p_i, in order to have only one value to compute. Such a bound is obtained from (32) by simply replacing $P_{D,i}^W(N)$ with $P_D^W(N)$ and by taking into account (26):

$$p_i \le P_D^W(N_F) + \sum_{N=0}^{N_F} P_D^W(N) \cdot Q_D^N(B/C) \quad \forall i. \tag{33}$$

The consistency of these bounds has been verified by simulations. It has been considered a link of capacity C=5 Mb/s, fed by markovian On/Off flows with peak rate P=32 Kb/s and activity a=0.35. The packet length is L=128 bytes. The dropping window interval is D=31.98 msec and N_F=156. The buffer size has been set to 15 packets (B=15360 bits). The Fig. 5 and Fig. 6 show the packet loss percentage vs. the multiplexer load. In the Fig. 5 the loss ratios at the dropper F and at the buffer B are compared with the analytical bounds (27) and (31), respectively. In the Fig. 6 the total packet loss in the multiplexer M is compared with the bound (33).

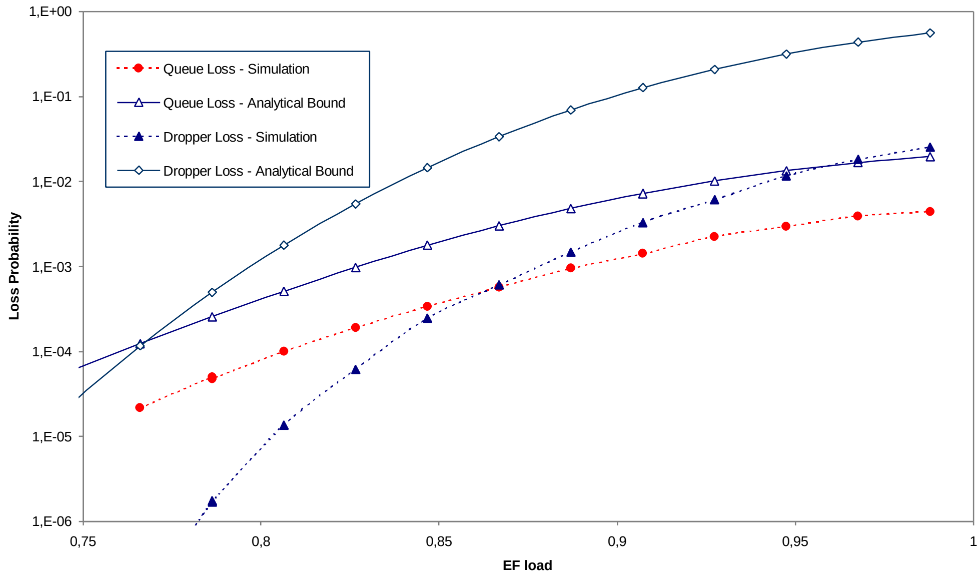

Fig. 5: Dropper and queue Loss probability vs. load: analytical and simulation results for a single stage

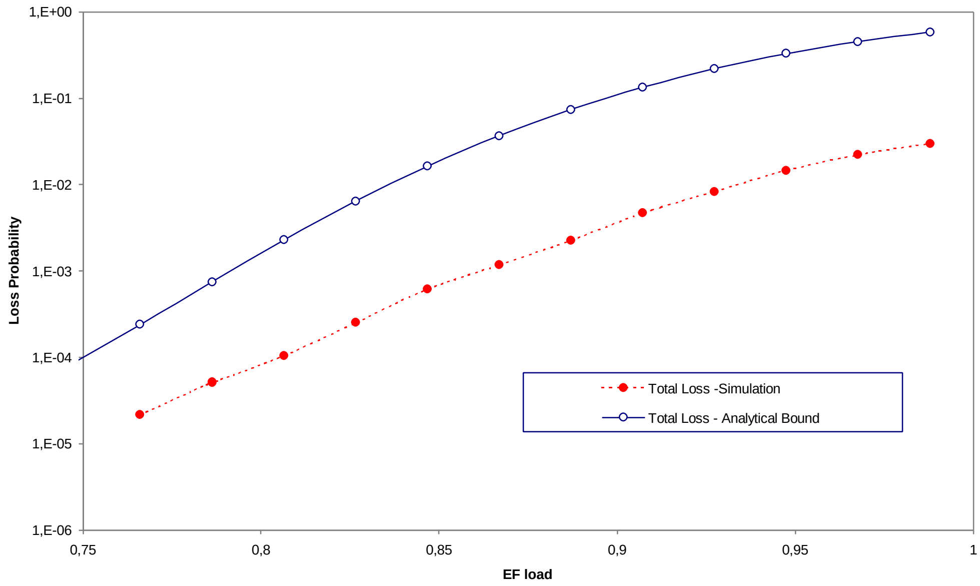

Fig. 6: Total loss probability vs. load: analytical and simulation results for a single stage

The figures highlight how the analytical bounds follow the actual values in the whole range of the considered load. Moreover, note that the loss at the dropper is negligible compared to the loss at the queue for lower loads, while for higher load the effect of the dropper is dominant.

3.4. Interaction between EF traffic and non-EF traffic

The analysis presented in the previous sections has been carried out assuming a network loaded by the EF traffic only. In more realistic scenarios the interaction of EF traffic with the non-EF one (e.g. best effort-TCP traffic) has to be taken into account. By supposing that EF packets are scheduled with non-preemptive simple priority over non-EF packets at every node, the effect of the non-EF traffic is due to the fact that an incoming EF packet, founding the EF queue empty, can not be immediately forwarded because the link is busy for the transmission of a non-EF packet. Obviously, such an interaction leads to a higher queuing delay and a higher loss probability for EF packets; the higher the non-EF packet size is, the stronger these impairments are. These increases of the queuing delay and of the loss probability have to be properly taken into account in the bounds (4) and (33).

Let L_{BE} (bits) be the maximum size of non-EF (Best Effort) packets, i.e. the MTU of the non-EF traffic. The additional delay introduced by the non-EF traffic on the EF packets of the i-th flow in the k-th node is thus bounded by $L_{BE} \big/ C_i^{(k)}$, so the expression (4) can be replaced by

$$d_i^{(k)} = \frac{B_i^{(k)} + L_{BE}}{C_i^{(k)}} \tag{34}$$

If L_{BE} is several times L (the size of a EF packet), such an additional delay can not be neglected; that is expected to be a very likely situation in practice: as a matter of example compare the short packet size of actual Voice over IP application, which is a good candidate to represent EF flows, with the long packet size of other data applications (e.g. FTP, HTTP).

As far as the loss probability is concerned, its increase depends on the fact that the transmission of non-EF packets could imply vacation periods of the server for the EF traffic. The duration of these vacation periods are not higher than L_{BE} / C. A suitable extension of (33) can be straightforwardly derived by considering a virtual multiplexer scheme, in which the server is always dedicated to EF traffic and the buffer size B_{mod} is equal to:

$$B_{mod} = B - L_{BE} \,. \tag{35}$$

It is intuitively that the loss experienced in the virtual multiplexer scheme is not higher than that arising in the actual scheme of Fig. 2 in which the two traffics coexist. On the basis of this observation, the (33) can be replaced by

$$p_i \le P_D^W(N_F) + \sum_{N=0}^{N_F} P_D^W(N) \cdot Q_D^N((B - L_{BE})/C) \quad \forall i \,. \tag{36}$$

The consistency of the bound (36) has been verified by simulations. The same simulation scenario as that in Fig. 6 has been considered, in which lower priority TCP traffic has been added. The size of the non-EF packets is L_{BE}=512 bytes, four times the length of EF packets.

The Fig. 7 shows that the effect of the non-EF traffic is negligible for high loads; in fact, in this region the loss component due to the dropper is dominant over that due to the buffer. Anyway, the bounds provided by (36) follows the actual behavior in the whole range of loads.

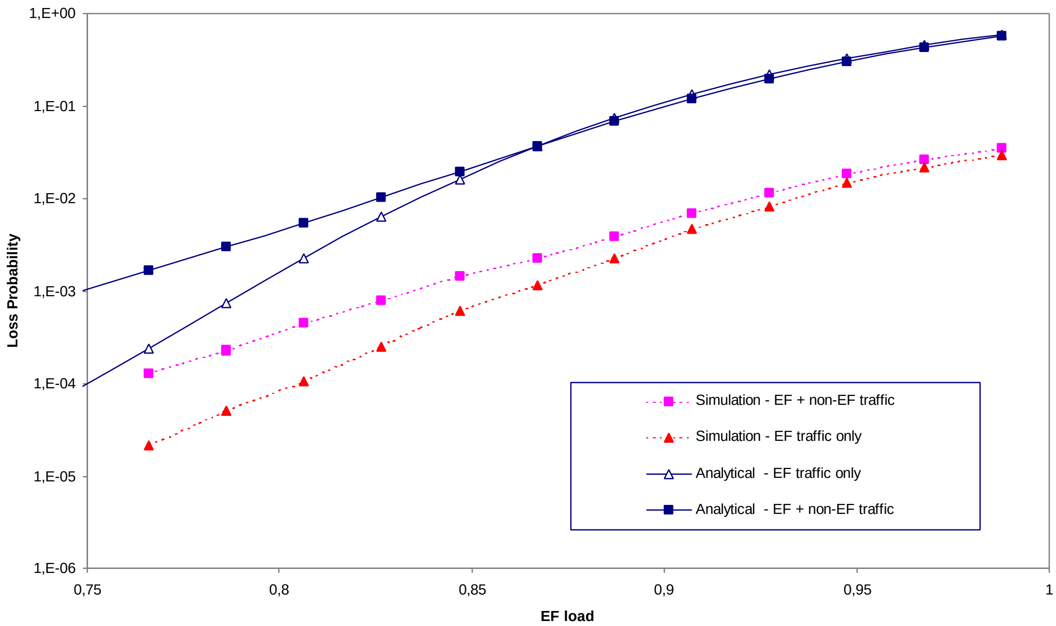

Fig. 7: Loss probability vs. load: comparison of analytical and simulation results for a single stage with and without non-EF traffic

4. THE ADMISSION CONTROL PROCEDURE

The right-hand term of the (33) [or (36)] represents the expression of the function $f(\mathsf{G},\mathsf{F})$ introduced in section 2 On the basis of the results obtained in the previous sections, by considering a generic node Ω, we have $\mathsf{G}_\Omega = \{B_\Omega, C_\Omega, D_\Omega, N_{F\Omega}\}$ and $\mathsf{F}_\Omega = \bigcup_{(i,k)\in \mathsf{H}_\Omega} \{R_i^{(k)}, P_i^{(k)}\}$.

As far as the dropper parameters are concerned, the dropping window duration D and the maximum number N_F of packets admitted in a dropping window have to be chosen accordingly to the input traffic flows set, as they have to meet the constraints (16) and (17). Such constraints define an admissible region for the dropper parameters for any input flows parameters set F_Ω. Inside that region, the values of D_Ω and $N_{F\Omega}$ minimizing the function $f(\mathsf{G}_\Omega, \mathsf{F}_\Omega)$ should be adopted. The aspects concerning the accurate minimization of such function are beyond the scope of the present paper; anyway our experience suggests that, in a wide range of cases, a good choice for the dropper parameters is:

- $N_{F\Omega}$ = the maximum integer such that $N_{F\Omega} \cdot t_\Omega < \min_m \{T_m^{(h)} | (m,h) \in \mathsf{H}_\Omega\}$;
- $D_\Omega = N_{F\Omega} \cdot t_\Omega + e$ (37)

wherein $0 < e << t_\Omega$.

When the peak rate of at least one flow approaches the link capacity, the proposed approach is not applicable. In fact in such a condition, on the basis of the constraints (16) and (17), D_Ω should be set as small as the packet service time t_Ω and the dropper would drop almost all the packets. That constitutes the primary limitation of the method. The effect of such a limitation on the admission control algorithm performance will be discussed in the next section.

It is worth noting that relation (37) is not simply a bound to the overall loss probability at a network node, but it bounds *each* per-flow loss probability independently. So, by applying (37), it is possible to guarantee that any single flow will not suffer packet loss over the admitted level; this is an useful consequence of the dropper action.

Summarizing, when a new flow i requests the EF service, the BB must sequentially perform the following steps:

1. it determines the path of the flow through the domain;

2. it computes the end-to-end deliverable QoS parameters Π_i^{tot} and Δ_i^{tot}, given by (5) and (6), and check if they are compliant with the requested QoS;
3. *For each node along the path*, it checks that

$$T_i^{(k)} > \tau_i^{(k)} \quad \forall k \in \mathrm{M_i} \tag{38}$$

wherein $T_i^{(k)}$ is computed using equations (7); if (38) is not met in every node, the flow is rejected; in fact, (38) is needed for the consistency of constraints (16) and (17).
4. *For each node along the path,* it computes the bound $f\left(\Gamma_i^{(k)},\Phi_i^{(k)}\right)$ from (33) [or (36) in case of presence of non-EF traffic] in the new input flows set and compares it with the node threshold $\Pi_i^{(k)}$. The flow is admitted only if $f\left(\Gamma_i^{(k)},\Phi_i^{(k)}\right) \le \Pi_i^{(k)}$ at any *k*–th stage, otherwise it is rejected.

As regards to step 4, it has to be noted that the computation of the function $f(\Gamma,\Phi)$ in a new input flows set would need a new couple of dropper parameters D and N_F to be determined, accordingly to the constraints (16) and (17). Such parameters should be dynamically updated whenever the input traffic set at a specific node changes, because of a flow admission or a flow termination. That would increase the complexity of the network architecture. Alternatively, the BB can adopt a static scheme, in which the dropper parameters are fixed. In this case the architecture is simplified, at the cost of an higher flow rejecting probability. In case of static scheme, in the control step 2 the test (38) must be substituted by the following:

$$T_i^{(k)} > D_i^{(k)}, \tag{39}$$

We conclude this section discussing some aspects concerning the practical applicability of the proposed allocation scheme.

The dropper aims at limiting the memory of the queue process at the dropping window duration, i.e. the packet time scale, by “filtering” conflicts at a larger time scale. At the burst level, the dropper is equivalent to a zero buffer multiplexer system. This is the reason why we can ignore in our approach the correlation structure of the traffic flows; this is a crucial analytical simplification, especially to take into account the flows degradation inside the network. Moreover, it leads to a practical simplification too, since just average and peak rates have to be known and communicated to the network by the single sources.

The given bound for the loss probability has a “long term” significance, i.e. it rigorously refers to the lost packet percentage over an infinite time scale. In such a context, the burst duration has no impact. In order to satisfy practical applicability constraints, it has to be assured that this bound applies over a finite time scale, comparable with the lifetime of the single flow. In other words, the duration of the burst periods has to be much lower than the duration of the flow. A quantitative evaluation of this aspect is out of the scope of this paper, but we observe that this requirement is widely met by typical real time applications, such as single or aggregated voice flows. However, in case of flows not satisfying this requirement (i.e. with long burst duration) the proposed method can be applied by considering these flows as CBR ones, thus ignoring their activities.

5. PERFORMANCE RESULTS

The analytical results derived in the previous sections have been validated by simulations. The simulations have been performed using the Network Simulator [19]. The reference network topology, depicted in Fig. 7, aims at representing a multistage network with both an access and a core section, characterized by different link speeds and flows aggregation levels. The goal is to verify that the proposed approach takes into account properly the effects of packet clumping and flows aggregation, and that the analytical bound derived for the flow packet loss probabilities are met at the different multiplexing stages.

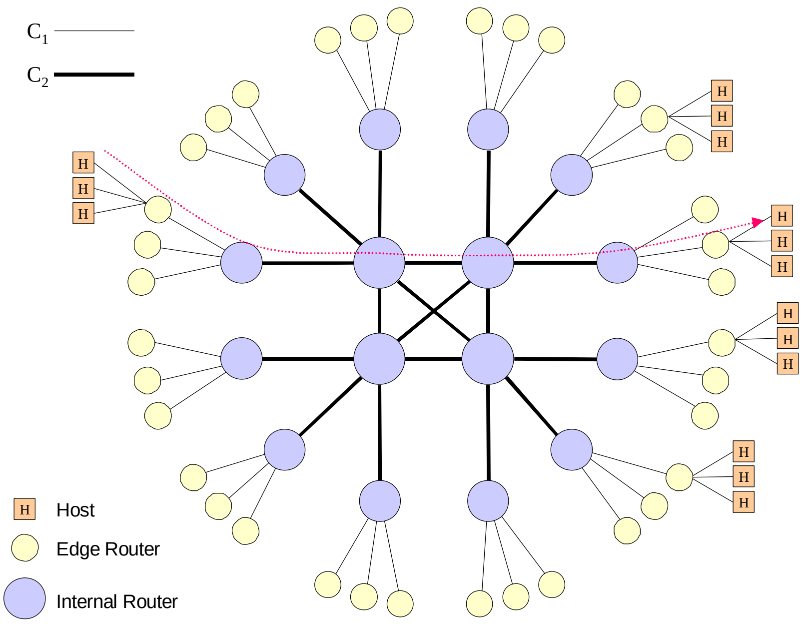

Fig. 8: Simulated network topology

The values of the node parameters for each multiplexing stage are given in Tab. 1. The buffer sizes yield a maximum end-to-end queuing delay of around 10 ms. Each node includes the implementation of the dropper device.

Multiplexing Stage	Link Capacity C (Mb/s)	Buffer Size B (packets)	Maximum Queuing delay	Dropper Parameters D	N_F
1	5	15	3.072ms	31.95ms	156
2	15	20	1.365ms	28.88ms	423
3	15	20	1.365ms	27.51ms	403
4	15	20	1.365ms	26.15ms	383
5	5	15	3.072ms	24.78ms	121

Tab. 1 - Parameters of the network elements.

The packet size for all flows is L=128 bytes. Two kinds of EF traffic sources have been considered:

- ON-OFF markovian sources, with constant bit rate P=32 Kb/s during the ON periods, activity a=0.35, average duration of ON and OFF period 350 ms and 650 ms respectively, packet size 128 bytes; such a model is representative of voice sources;
- CBR sources with rate P=32Kb/s and packet size 128 bytes.

Moreover, non-EF traffic sources with packet size 512 bytes have been modeled by long term TCP connections.

The minimum inter-arrival packet time for EF traffic at the source $T_i^{(1)}$ is 32 ms, therefore the queuing delay is very low, accounting for one third of $T_i^{(1)}$.

Different scenarios with different traffic mixes have been considered:

a) homogeneous scenario with ON-OFF flows only;
b) heterogeneous scenario with both ON-OFF and CBR flows; the number of ON-OFF flows is three times the number of CBR flows, so that 50% of the total amount of traffic is relevant to each traffic type;
c) scenario (a) with additional non-EF traffic;
d) scenario (b) with additional non-EF traffic.

An extensive simulation campaign has been carried out for all the scenarios mentioned above. However, for sake of clarity, only results relevant to the scenario (a) will be presented here. We remark that the results relevant to all of the other scenarios have confirmed the observations and the conclusions given hereafter.

The source and destination end-points are connected to the Edge Routers (ER) numbered from 1 to 36 clockwise. Each flow path is composed of five multiplexing stages from the ingress ER to the egress ER. A regular traffic matrix implying homogeneous load at each link has been used.

The analytical bounds for the per-flow packet loss probability have been compared with the percentage of lost packet obtained by simulations, at different multiplexing stages and for different link load values. In order to have significant results (i.e. measurable loss ratios), the range of load which we have considered is higher than the 50% which is the maximum allocation of the EF service to strictly meet the PHB definition in presence of non-EF traffic [1].

Fig. 9 reports the curves of the analytical bound for the loss probability and the measured loss percentage for two different network stages (3^{rd} and 5^{th}). The figure shows that the loss probability at the 5-th stage is higher of about one order of magnitude than at the 3-rd stage. This is due to the combined effect of the lower link capacity and the higher flow degradations at the last stage; however this is not a general result since further elements have to be taken into account, e.g. the number and the capacities of the input line at each stage. Anyway, fixed a threshold for the loss probability, the calculated bounds allow in both cases an efficiency which is just 10% lower than the measured limits.

As a comprehensive result, highlighting the overall performance of the proposed flow admission procedure, we have evaluated the allocation efficiency achievable by the given admission algorithm. For each multiplexing stage of the considered topology, Fig. 10 shows the maximum achievable load given a per-node maximum admitted loss probability.

Note that a higher load is achievable on the internal links, which have larger capacities and allow for higher statistical gain. For the same link capacity the achievable load diminishes with the multiplexing stage, as the flows degradation augments. Anyway, in the considered scenario such a diminution is quite smooth. This is mainly due to the condition that the involved values for the end-to-end maximum queuing delay ($\cong$10.2 ms) remains quite below the minimum inter-packet interval for the considered traffic flows (32 ms). Therefore the dropping window duration *D* can be chosen sufficiently large to get a good statistical gain even for the fifth stage.

The achievable utilization is compared with the worst case allocation obtained with the methodology described in [11] for the specific topology in Fig. 8. In particular the end-to-end worst case delay is fixed to 10.2 ms and the maximum load that yields deterministic no loss guarantees is evaluated. In this case a uniform load is assumed at each link.

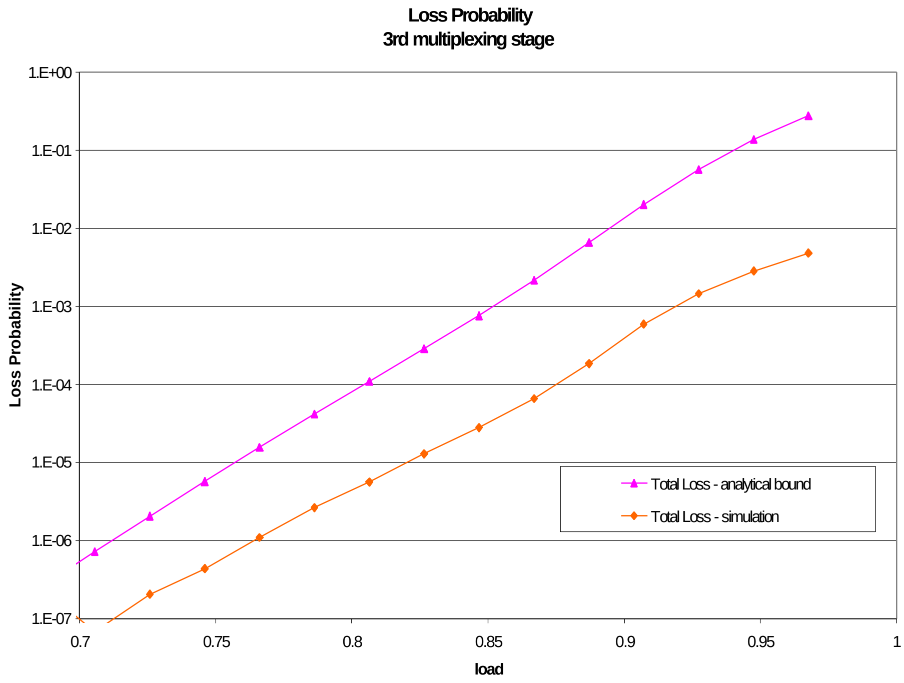

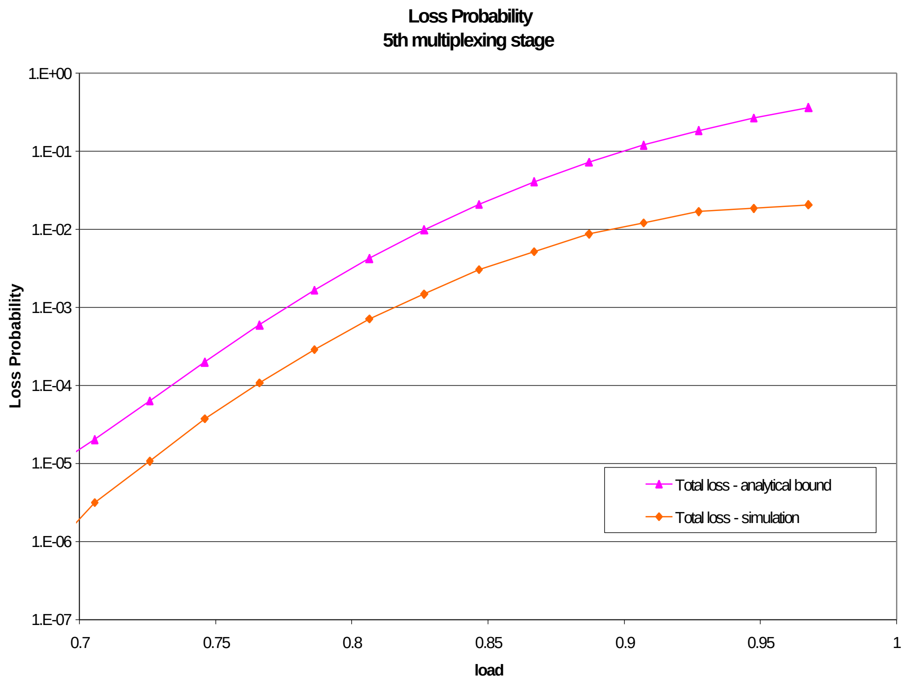

Fig. 9: Analytical bound and simulated loss: a) stage 3; b) stage 5.

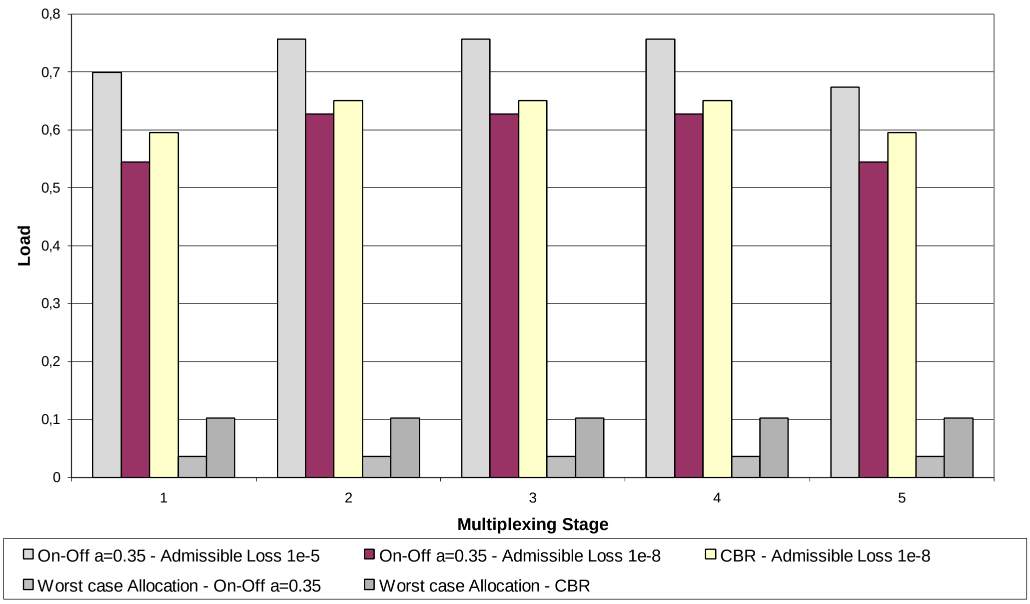

Fig. 10: Maximum achievable load at each multiplexing stage in the considered network.

The primary limitation to the applicability of the proposed method is represented by the constraint (17) on the dropping window duration. In a generic network node, the admission of flows with a high peak rate, comparable to the link speed, constraints the dropping window to be very strict, so that the dropper loss will become unacceptable.

Therefore, it derives that the proposed procedure can not allocate:

1) native flows with peak rate comparable to the link rate;
2) degraded flows for which the accumulated delay along the path approaches the minimum native packet inter-arrival time.

As for the flows of the first type this is not an effective limitation since statistical multiplexing is not applicable to them.

As far as the second flow type, we refer to Fig. 11 in order to analyze the delay accumulation effect. This figure shows the achievable utilization factor vs. the cumulated maximum delay, for ON-OFF and CBR sources with the parameters described above, on a 5 Mb/s link with a buffer size equal to 15 packets. Curves refer to a total loss equal to 10^{-5} and 10^{-8} for ON-OFF flows and 10^{-8} for CBR flows.

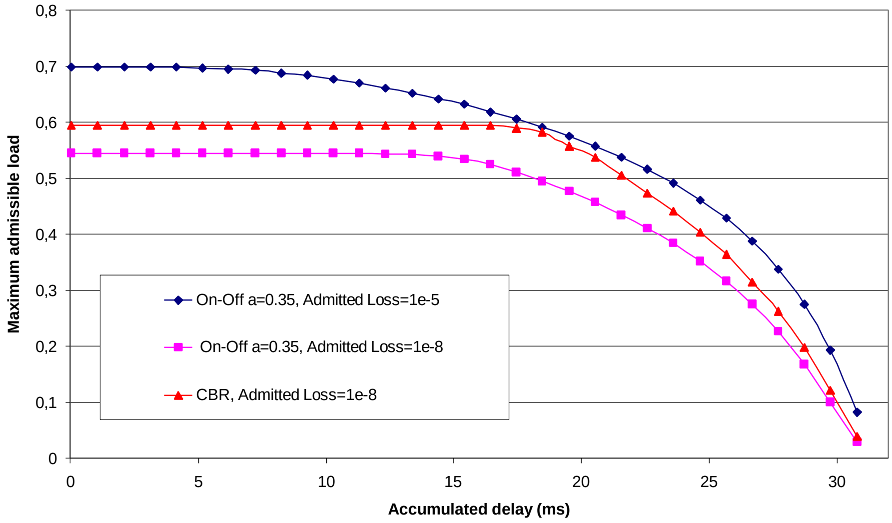

Fig. 11: Maximum admissible load vs. cumulated delay

The plot can be read in the following way: given the maximum delay cumulated along the k-1 previous stages (reported in abscissa) by a flow entering the k-th stage, one can read the achievable utilization at the k-th stage for the given buffer size and loss threshold. When applying the allocation scheme to the successive stages, the delay budget must be increased by the maximum delay introduced by the k-th node. This means that we are moving rightwards on the utilization curves, the horizontal step obviously depends on the node buffer size and on the output line speed.

From the shape of the curves, it can be seen that the allocation efficiency slowly decreases for a number of stages, while the decrease is steeper if the number of stages is higher. The critical value depends on the combination of the system parameters.

6. CONCLUSIONS

In this paper a framework for flow admission control in a Differentiated Services network using the Expedited Forwarding PHB has been proposed. The fundamental assumption is to remove the constraint of deterministic no loss and to consider quantifiable bounds to the loss probability. This approach is especially suitable to deal with variable bit rate flows like voice calls or aggregated voice traffic. The evaluation of the bound to the loss probability is based on an analytical result by Beneš. In order to apply this result in realistic cases where burst scale congestion can appear, a device (dropper) is introduced at the ingress of each multiplexing stage. The dropper device has the purpose to discard packets in order to avoid burst level congestion. An upper bound to the loss introduced by the dropper is evaluated and properly taken into account into the overall loss probability of the node.

The strength of the proposed approach is that it can take into account the degradation of the traffic characteristics in a cascade of nodes. Therefore the upper bound can be evaluated at the different multiplexing stages for a given network topology and end-to-end guarantees can be provided. Simulation results to compare the analytical bound with simulated loss are presented.

Conservative hypothesis are needed to take into account the flow degradation in the end-to-end analysis. Hence the allocation efficiency diminishes for each crossed stage, until it becomes unacceptable. Nevertheless, we have shown that the allocation performance of the algorithm is still satisfactory for network scenario of practical interest, composed of several nodes.

7. REFERENCES

[1] R. Braden, D. Clark, S. Shenker: “ Integrated Services in the Internet Architecture: an Overview”. IETF RFC 1633, June 1994.

[2] S. Shenker, C. Partridge, R. Guerin: “ Specification of Guaranteed Quality of Service”. IETF RFC 2212, September 1997.

[3] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss: “ An Architecture for Differentiated Service”. IETF RFC 2475, December 1998.

[4] K. Nichols, S. Blake, F. Baker, D. Black: “ Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers”. IETF RFC 2474, December 1998.

[5] V.Jacobson , K. Nichols, K.Poduri “An Expedited Forwarding PHB”, IETF RFC 2598, June 1999.

[6] J. Heinanen. F. Baker, W. Weiss, J. Wroclawski “Assured Forwarding PHB group” IETF RFC 2597, June 1999.

[7] X. Xiao, L. M. Ni: “Internet QoS: A Big Picture” IEEE Network, Vol. 13, No. 2, March/April 1999.

[8] S.Carter, T. Hodgkinson, A. O’Neil, D. Moritmore “A Bounded-delay service for the Internet” draft-carter-bounded-delay-00.txt, Internet Draft.

[9] I. Stoica, H. Zhang: “Providing guaranteed services without per flow management”. Technical Report CMU-CS-99-133, May 1999.

[10] ACTS AC310 “ELISA” project - Deliverable “Specification of Traffic Handling” June 1998. http://www-st.inf.tu-dresden.de/elisa/

[11] M. Listanti. L. Veltri: “Worst case analysis of a cascade of nodes”. Coritel Technical Report 14/99.

[12] E.W. Knightly, N.B. Shroff, “Admission Control for Statistical QoS: Theory and Practice”, IEEE Network, Vol. 13, No. 2, March/April 1999.

[13] A. Detti, M. Listanti, S. Salsano, L. Veltri “Supporting RSVP in a Differentiated Service Domain: an Architectural Framework and a Scalability Analysis”, ICC’99, Vancouver (Canada) June 1999, paper S6.3.

[14] Y. Bernet et al. "Integrated Services Operation over Diffserv Networks", Internet Draft, June 1999, http://www.ietf.org/internet-drafts/draft-ietf-issll-diffserv-rsvp-02.txt

[15] U. Mocci, J. Roberts, J. Virtamo eds., “Broadband Network Teletraffic”, Final report of action Cost-242, Springer 1996.

[16] I. Norros, J.W. Roberts, A.Simonian, J.T. Virtamo “The Superimposition of Variable Bit Rate Sources in an ATM Multiplexer”, IEEE J-SAC, Vol. 9, No. 3, April 1991.

[17] V. E. Beneš “General Stochastic Processes in the Theory of Queues”. Addison Wesley 1963.

[18] M.Listanti, F. Ricciato, S. Salsano: “A statistical approach to end-to-end QoS guarantees in packet switched networks”. Coritel Technical Report 15/99.

[19] Network Simulator 2 developed by Lawrence Berkley Network Laboratory and University of California Berkley, *http://www-mash.cs.berkeley.edu/ns*.