Evaluation and Comparison of MBAC Solutions

Doreid Ammar*, Thomas Begin*, Isabelle Guérin-Lassous* and Ludovic Noirie[†]
*Université Lyon 1 / LIP (UMR ENS Lyon - INRIA - CNRS - UCBL)
Email: {doreid.ammar, thomas.begin, isabelle.guerin-lassous}@ens-lyon.fr
[†]Alcatel-Lucent Bell Labs, Nozay, France
Email: ludovic.noirie@alcatel-lucent.com

Abstract—**Admission control is a mechanism used to restrict access to a computer network to some flows based on the current utilization level of the network resource. By regulating the number of on-going flows, admission control aims at preventing overloading, congestion and performance collapses, so that, accepted flows receive a sufficient level of Quality of Service (QoS). In this paper, we evaluate three existing measurement-based admission control (MBAC) solutions, and we compare their efficiency in the context of semantic networks.** *Semantic networks* **refer to networks that autonomously acquire a knowledge on the on-going traffic as well as on any new incoming flow requesting admission. In this framework, we configure the three MBAC solutions in a way they have an identical target in terms of maximum tolerable packet loss rate or maximum tolerable packet queueing delay. We evaluate the solutions performance analytically or by simulation, and compare them to the "ideal" admission control. The results show that one solution, outperforms the others in meeting the target performance.[1]**

I. Introduction

Admission control is a mechanism used to restrict access to a computer network to some flows based on the current utilization level of the network resource. By regulating the number of on-going flows, admission control aims at preventing overloading, congestion and performance collapses, so that, accepted flows receive a sufficient level of Quality of Service (QoS). Despite the large number of studies carried out to date, very few solutions have been deployed on operational networks. It is common for operators to consider that their networks are over-dimensioned as compared to current needs in communication, *i.e.*, workload, and therefore, that implementing an admission policy is out of the picture.

Recently, network operators have observed profound changes in the daily traffic. The number of applications that generate traffic has tremendously grown up. Numerous delay-sensitive applications (*e.g.*, Telephony over IP) and resource-intensive applications (*e.g.*, streaming video) are constantly emerging. This steady increase of applications, combined with intensive use, has significantly affected the utilization level of networking resources, and it might, ultimately, cause significant network congestions and performance disruptions. A case in the point is the traffic collapse that occurred during summer 2010 on AT&T wireless access networks, in which available bandwidth is known to be a scarce resource. Furthermore, given the foreseeable increasing demand for network services

(*e.g.*, Internet TV, Video on Demand, high definition 3D video), it is the authors point of view that congestion problems may also occur in wired networks. We believe that admission control is still an active and relevant research area.

In this paper, we focus our study on admission control in the context of semantic networks [7]. *Semantic networks* refer to computer networks that autonomously acquire a knowledge on the on-going traffic. They analyze the features of the transmitted traffic at the flow granularity and exploit this knowledge to dynamically adjust their behavior. In the context of admission control, not only does the network acquire knowledge on the characteristics of on-going traffic, as it would be the case for any measurement-based admission control (MBAC), but it also gets knowledge on any new incoming flow requesting admission (thanks to the inspection of its first packets). It is this paper goal to evaluate and to compare the performance of three existing MBAC solutions in the context of semantic networks.

The originality of this work is twofold. First, as opposed to previous comparison studies [5], [2], [6], we do not assume any explicit knowledge, neither on incoming flows nor on on-going traffic. To this end, we introduce a method to estimate the peak rate of an incoming flow based on its first transmitted packets. Second, we parameterize three existing MBAC solutions in a way they have an identical target in terms of performance and we compare their efficiency. The selected target is alternatively the packet loss rate or the packet queueing delay.

The remainder of this paper is organized as follows. In Section II, we detail the admission control solutions, which are investigated in our study. Our experimental framework is presented in Section III. Section IV is devoted to the numerical results. Finally, Section V concludes this paper.

II. Investigated MBAC Solutions

In our study, we investigate three MBAC solutions. Note that all these solutions assume to know the peak rate of each new flow requesting admission. We denote by r the peak rate of an incoming flow and by C the nominal capacity (transmission speed) of a communication link. In Section III-C, we detail a simple technique to evaluate r.

A. Measured Sum (M.S.)

Jamin et al. present in [5] a MBAC solution based on the measured load of existing traffic over the link, denoted by R.

This solution admits a new flow requesting admission, with a peak rate r, if and only if:

$$R + r \leq \nu C, \tag{1}$$

where ν is a parameter that defines the targeted link utilization. The measured load of existing traffic is updated every measurement window of length T. This time window is split into smaller sampling periods of equal length. The average rate of existing flows is then computed on every sampling period. At the end of a measurement window, R is defined as the highest average rate seen in the sampling periods constituting this time window. However, the value of the measured link load may be updated within a measurement window for two reasons: whenever an individual average rate on a sampling period exceeds the current link load of the measurement window or whenever a new flow is admitted, the value of the measured network load is then updated with the value of the average rate of the sampling period or with the peak rate of the new admitted flow added to the current load respectively. Note that the measured load on the sampling periods are always stored and used to compute the average load at the end of a measurement window.

Jamin et al. introduce a delay test to their admission control solution. The measured delay, denoted by $\widehat{D}$, tracks the maximum queueing delay of every packet computed over a time window of length T. The solution rejects an incoming flow requesting admission if admitting this new flow violates the following constraint:

$$D > \widehat{D} + \frac{b_i}{C}, \tag{2}$$

where D is the delay bound and b_i is the burstiness of the flow (see details in [5]). The value of $\widehat{D}$ is updated at the end of each measurement window. Whenever an individual delay measurement exceeds the estimated maximum queueing delay, the value of $\widehat{D}$ is also updated to be λ times this sampled delay. Finally, we update the measured delay to the right side of (2), whenever a new flow is admitted.

B. Equivalent Capacity (E.C.)

In [3], Floyd presents an admission control solution based on the estimation of the *Equivalent Capacity* of the link for a set of aggregated flows. A new flow is accepted if the sum of the peak rate r, requested by this flow, and the *Equivalent Capacity* of the link, $C(\epsilon)$, is less than or equal to the capacity of the link C. More formally, this condition is expressed as:

$$C(\epsilon) + r \leq C \tag{3}$$

The critical point of this method relies on the estimation of the *Equivalent Capacity*, $C(\epsilon)$. In our case study, we chose the formula given in [4] because it is easier to use in the context of *semantic networks*. The *Equivalent Capacity* proposed in [4] is a linear function of the average rate of aggregate traffic and its standard-deviation, denoted by $\widehat{r}$ and σ, respectively. This function is given by:

$$C(\epsilon) = \widehat{r} + \alpha.\sigma, \text{ with } \alpha = \sqrt{2\ln\frac{1}{\epsilon} + \ln\frac{1}{2\pi}}, \tag{4}$$

where ϵ is the probability that the arrival rate exceeds the expected *Equivalent Capacity*.

In order to compute the average arrival rate of aggregated traffic, $\widehat{r}$, *Floyd* suggests to define it as an exponential-weighted moving average with a weight ω updated after each measurement window T. The average arrival rate could then be calculated using: $\widehat{r} = (1 - \omega).\widehat{r} + \omega.R$, where R is the average rate of the aggregated traffic measured every measurement window T and ω is a real number between 0 and 1. Since nothing was recommended by the authors about the computation of the standard-deviation σ, we chose to compute the value of σ from the M previous measured values of R.

C. Aggregate Traffic Envelopes (Env.)

Qiu and Knightly present in [8] a MBAC solution that aims to characterize the aggregate traffic rate by the maximal rate envelope. To do this, they consider a time window of length T divided into t sampling periods of equal length. Within a time window, maximal rate measurements are done on different time scales. R_k^m represents the maximal observed rate in the time scale k. This time scale is equal to k sampling periods in the m^{th} measurement window. The rate of the aggregate traffic and its standard-deviation are estimated over the last M measurement windows as follows:

$$\overline{R}_k = \sum_{m=1}^{M} \frac{R_k^m}{M} \text{ and } \sigma_k^2 = \frac{1}{M-1} \sum_{m=1}^{M} (R_k^m - \overline{R}_k)^2. \tag{5}$$

This measurement-based admission control solution consists of two parts: a short time scale test that ensures that no packet is too long delayed, and a long time scale test that checks that the flow requesting admission does not exceed the link capacity. Note that envelopes are used only to check the first condition. A new flow requesting admission with a peak rate r is accepted if and only if:

$$\max_{k=1,...,t} \{k\tau(\overline{R}_k + r + \alpha_E \sigma_k - C)\} \leq C \times D \tag{6}$$

and

$$\overline{R}_t + r + \alpha_E \sigma_t \leq C \tag{7}$$

where D is the maximum delay requirement and α_E is a constant specifying the confidence level, $\Phi(\alpha_E)$, that on-going flows do not experience any packet loss. $\Phi(\alpha_E)$ is defined as:

$$\Phi(\alpha_E) \approx \frac{1}{\sqrt{2\pi}\sigma_k} \int_{-\infty}^{\overline{R}_k + \alpha_E \sigma_k} \exp\left(-\frac{(r - \overline{R}_k)^2}{2\sigma_k^2}\right) dr. \tag{8}$$

III. METHODOLOGY

A. Objective

In this paper, we aim at highlighting the ability for each of the three MBAC presented above to achieve the maximum level of utilization of the link, while respecting a given target in terms of performance. In this work, we choose alternatively the loss rate and the queueing delay as target. We consider the following values for the target loss rate, Pr, and for the target queueing delay, D, namely $Pr = 10^{-2}$ and $D = 10$ ms.

In this context, the *ideal admission control* is clearly defined: it accepts the maximum number of flows over a communication link, thus achieving the maximum utilization rate, while meeting the QoS target. For the sake of comparison, we always include next to the MBAC results those provided by the ideal admission control.

B. Scenarios

We consider a single communication link of capacity, C, set to 10 Mbps. The queue size is set to 20 ms when we evaluate the loss rate, and to 60 ms when we study the queueing delay. The queueing discipline is FIFO (*First In First Out*) and the queue management algorithm is Drop-Tail.

Incoming flows requesting admission are modeled by constant bit rate (CBR) flows. They can be viewed as audio flows with a sending rate of 64 Kbps and a constant packet size equal to 190 bytes, which corresponds to the use of the codec G.711. To take into account possible variations in network conditions, instead of modeling the inter-packets arrival time as a simple constant, which is exactly the case for a CBR flow, the inter-packets arrival time is modeled by a constant added to a truncated Normal distributed random variable.

We consider that the aggregate traffic on the link, denoted by *background traffic*, consists of an initial background traffic modeled by a Poisson process, to which the aggregation of CBR flows already accepted by the admission control is added. This initial background traffic generates packets of length 190 bytes with an average rate set between 1 and 7 Mbps.

C. Estimating the peak rate of incoming flows

As said previously, we focus our studies in the context of *semantic networks*. In such networks, the network acquires knowledge on flows by itself via an analysis of the on-going traffic. In this section, we detail the procedure we implement to let the network estimate the peak rate of a new flow requesting admission.

To estimate the peak rate of a new incoming flow, we track the first n packets of this flow. We use a sliding window of length equal to k packets. For every possible window on the first n packets, we compute the average rate. Finally, the peak rate corresponds to the highest value among the $(n - k + 1)$ windows. In this work, the estimated peak rate of an injected flow is computed based on the 20 first packets ($n = 20$) with a sliding window of length equal to 5 packets ($k = 5$).

D. Calibration of the admission control algorithms according to a target loss rate, Pr

We now detail the configuration of the investigated admission controls. As said before, we calibrate their parameters such that all of them have an identical target in terms of maximum tolerable packet loss rate. Remind that Pr denote this target loss rate.

1) Measured Sum: The authors of the Measured Sum algorithm did not provide specific guidelines for selecting the value of ν. Obeying to the analysis principle proposed in [5], we choose the value of ν as equal to the ratio of the average packets arrival rate to the average transmission (service) rate so that the link modeled by a corresponding *M/M/1/K* queue, with K set to 131 packets, leads to a packet loss rate equal to Pr. Therefore, we set the value of ν to 0.95 which is associated to $Pr = 10^{-2}$.

2) Equivalent Capacity: The authors denote by ϵ the probability that the instantaneous arrival rate of the background traffic, modeled by a Normal distribution, exceeds the *Equivalent Capacity* of the communication link, $C(\epsilon)$. To link the value of ϵ to the value of Pr, we proceed as follows. Assuming that the probability ϵ also represents the steady-state probability of having the buffer full (which would be the case for a buffer length of 1), and assuming that the steady-state probabilities are the same as the probability of the state seen by an arriving packet (which would be the case if the incoming flows were Poisson, see the PASTA property [1]), then, ϵ would also be the probability for an incoming packet to be rejected, namely Pr. Based on this rationale, we select ϵ equal to Pr, and thereby computing the value of α, namely $\alpha = 2.7$.

3) Aggregate Traffic Envelopes: The selected value for the confidence level, α_E, determines the expected probability that on-going flows do not experience any packet loss, $\phi(\alpha_E)$. To choose the value of α_E, we simply associate the value of $\phi(\alpha_E)$ to the target packet loss rate for accepted flows, namely Pr. Setting the value of Pr to 10^{-2} leads to an α_E value equal to 3.6.

E. Calibration of the admission control algorithms according to a target queueing delay, D

We describe here how we parameterize the admission controls according to a target queueing delay. Note that *Equivalent Capacity* is obviously out of this section as it does not provide a control on the packet delay. Recall that D denote the target maximum tolerable queueing delay over a single communication link (see formulas (2) and (6)).

1) Measured Sum: The value of λ aims at tuning the stringency level of the admission control. The greater λ, the more conservative the admission control is, and the less accepted flows. As no specific guidelines are given by the authors of *Measured Sum* for setting the value of λ, we let λ be equal to the most favorable value that we observed in several experiments, namely $\lambda = 2$.

2) Aggregates traffic envelopes: There is no clear recommendation from the authors on the choice for α_E. Therefore, we set it to the more favorable value among couple of experimented values, namely $\alpha_E = 3.6$.

IV. PERFORMANCE COMPARISON

In this section, we evaluate the performance of the three MBAC solutions and we compare them to the ideal admission control. We performed solutions using the following settings. *1) Measured Sum*: we set the value of T to be equal to 4 s, and we use sampling periods of 200 ms. *2) Equivalent Capacity*: we fixed the measurement window length to be equal to 200 ms. *3) Aggregate Traffic Envelopes*: we set the value of T to be equal to 4 s. This time window is split into 20 sampling periods of 10 ms. Our approach proceeds as follows: four
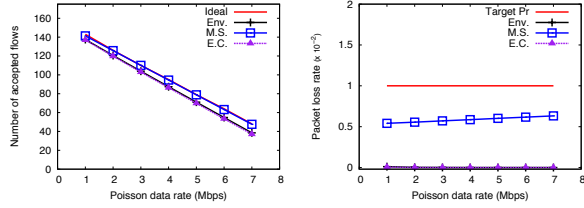
(a) Number of accepted CBR flows   (b) Achieved loss rate

Fig. 1: Solutions performance for a target packet loss rate of 10^{-2}

(a) Number of accepted CBR flows   (b) Achieved queueing delay

Fig. 2: Solutions performance for a target queueing delay of 10 ms

CBR flows arrive per second according to a Poisson process and attempt to access the communication link as described in Section III-B. Whenever a new flow is accepted, it keeps transmitting packets until the end of the simulation time. An incoming flow requesting admission is rejected by the admission control if the MBAC solution considers that the target loss rate Pr (the target queueing delay D, respectively) has been reached. In this case, we consider that the maximum level of link utilization for this MBAC has also been reached. To compute the number of incoming flows accepted by the ideal admission control, denoted by *Ideal* in the following figures, we iteratively look for the maximum number of flows that can enter in the considered communication link while sucessfully meeting the performance target. To do this, we rely on an analytical approach when we evaluate the packet loss rate, and on a simulation, performed with the NS-3 simulator, when we evaluate the packet queueing delay.

A. Loss rate

It is straightforward to assess analytically the maximum number of CBR flows that each MBAC can admit. We simply rely on the formulas (1) (3) and (7), respectively, given that variance values for CBR flows and a Poisson process are known. Then, we approximate the actual loss rate of CBR flows to the overall packet loss rate obtained from the corresponding *G/G/1/K* queue. To compute the ideal number of admitted CBR flows, we iteratively look for the maximum number of CBR flows that can enter in the *G/G/1/K* queue while keeping the loss rate below Pr. Note that we choose the parameters of the inter-arrival and service time distributions so that the mean and the variance values of inter-arrival and service times reproduce those derived from the combination of a Poisson process and CBR flows.

Figure 1a represents the total number of accepted CBR flows for each of the three MBAC solutions as compared to the ideal admission control for different rates of the initial background traffic. We observe that each MBAC is able to closely meet the ideal number of accepted CBR flows. The packet loss rate associated to this utilization level is shown in Figure 1b. This latter shows that each considered MBAC leads to a packet loss rate value lower than 10^{-2} whatever the intensity of the background traffic is.

B. Queueing delay

We now consider the packet queueing delay as target instead of packet loss rate. In Figure 2a, we represent the total number
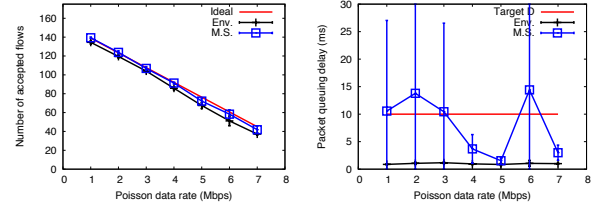
of accepted CBR flows for each of the two MBAC solutions as compared to the ideal admission control for different rates of the initial background traffic. Figure 2b shows the corresponding average packet queueing delay with regards to the target delay D. The latter results were obtained using simulations. To begin with, we observe that only the *Aggregate Traffic Envelopes* solution achieves to always guarantee the target queueing delay D. This clearly implies a decrease in the utilization level, but at a moderate degree. On the other hand, the *Measured Sum* solution exhibits erratic results. For some of the simulated replications, the number of accepted CBR flows is too large, and therefore the target queueing delay is getting violated.

V. CONCLUSION

In the context of semantic networks, we compared three MBAC solutions without assuming any explicit knowledge on the incoming flows requesting admission. We parameterized them in a way they share an identical target in terms of maximum loss rate or in terms of maximum queueing delay. We evaluated the performance of the three solutions considering CBR incoming flows and an initial background traffic, modeled by a Poisson process.

The results tend to show that, in this framework, all investigated solutions are able to fulfill the target loss rate. On the other hand, only the *Aggregate Traffic Envelopes* is always able to meet the target queueing delay unlike the *Measured Sum* solution that may violate this target in some cases.

REFERENCES

[1] A. Allen, "Probability, statistics and queueing theory, second edition," in *Academic Press, London*, 1990.
[2] L. Breslau, S. Jamin, and S. Shenker, "Comments on the performance of measurement-based admission control algorithms," in *Infocom*, 2000.
[3] S. Floyd, "Comments on measurement-based admissions control for controlled-load services," Tech. Rep., 1996.
[4] R. Guerin, H. Ahmadi, and M. Naghshineh, "Equivalent capacity and its application to bandwidth allocation in high-speed networks," *IEEE JSAC*, vol. 9, no. 7, pp. 968–981, Sep 1991.
[5] S. Jamin, S. Shenker, and P. B. Danzig, "Comparison of measurement-based admission control algorithms for controlled-load service," in *Infocom*, 1997.
[6] A. Nevin, Y. Jiang, and P. J. Emstad, "Robustness study of mbac algorithms," in *ISCC*, 2008.
[7] L. Noirie, E. Dotaro, G. Carofiglio, A. Dupas, P. Pecci, D. Popa, and G. Post, "Semantic networking: Flow-based, traffic-aware, and self-managed networking," *Bell Labs Technical Journal*, 2009.
[8] J. Qiu and E. W. Knightly, "Measurement-based admission control with aggregate traffic envelopes," *IEEE/ACM Transactions on Networking*, vol. 9, no. 2, pp. 199–210, Apr 2001.