Performance Improvements on Self-Similar Traffic Using Measurement-Based Admission Control

Vincenzo Mancuso, Giovanni Neglia and Giuseppe Bianchi
Department of Electrical Engineering, University of Palermo, Italy
{vincenzo.mancuso, giovanni.neglia}@tti.unipa.it
bianchi@elet.polimi.it

Abstract

In this paper we consider an admission-controlled traffic scenario, in which non-adaptive connections are offered to the network. We show the superiority of Measurement Based Admission Control (MBAC) algorithms over Parameter Based Admission Control (PBAC) algorithms, in particular when Long Range Dependence (LRD) or asymptotic second order self-similarity arises in the offered traffic. We show that a pure MBAC approach is capable of significantly reduce the Long Range Dependence of the accepted traffic aggregate, thus yielding a considerable performance improvement, in terms of Quality of Service (QoS) experienced by each flow sharing the media. Moreover the MBAC approach results in a powerful tool, robust to traffic correlations properties and burstiness, and able to overcome the horizon of PBAC scheme, despite of the traditional meaning of MBAC role, commonly intended as an even coarse approximation for PBAC.

1 Introduction

As observed in wide area Internet traffic ([1], [2]), self-similarity is shown by the traffic aggregate, both considering TCP or UDP flows. In particular this paper focuses on traffic generated by non-reactive sources, i.e. sources that do not react when congestion arises. Even if this kind of flows is not numerically predominant, it relates to a class of marcket strategically relevant applications, as real-time multimedia streams, that is speedily growing, and a great effort has been recently expressed in order to offer to these applications a suitable QoS level.

Since it has been shown that Long Range Dependence, also known as asymptotic second order self-similarity, arises when flows has heavy-tailed periods of activity/inactivity (see [3]), we assume, for convenience, a traffic aggregate scenario resulting from the superposition of homogeneous heavy-tailed flows. Since a widespread evidence that humans as well as computer sources tend to behave as heavy-tailed on/off sources ([2], [4], [5]), the result should be considered a physical explanation of the traffic self-similarity, independently of network or protocol traits, rather than a mere way to generate self-similar traces.

Moreover, self-similarity has a negative impact on network performance ([5], [6], [7], [8], [9]). In fact, for a given link capacity, a buffer size, and a fixed number N of superposed offered flows, the QoS (e.g. loss ratio, delays, etc) experienced by heavy-tailed flows results much worse than that experienced by flows whose activity/inactivity periods are drawn from exponential distributions. The PBAC scheme consists in checking that the number of flows admitted to the considered link never exceeds a threshold N_t, computed as the maximum number of calls that can be admitted while still satisfying predetermined QoS requirements. The "parameter-based" stays in the fact that the threshold N_t depends on the flow statistic parameters. As a consequence of self-similarity, the maximum number N_t of heavy-tailed flows that can be admitted to a link may be much lower than in the case of markovian (MRK) flows.

When a PBAC rule is replaced by a Measurement Based Admission Control (MBAC) scheme ([10], [11], [12], [13]), a number of quite unexpected beneficial effects can be experienced. Firstly, MBAC schemes provide superior performance than PBACs when LRD flows are considered. Secondly, the traffic aggregate resulting from the accepted flows shows very little Long Range Dependence - in other words, unlike PBACs, MBAC appears capable of smoothing the self-similarity of the accepted traffic aggregate. Finally, we argue that MBAC approaches are not mere "approximations" of ideal CAC schemes, useful in situations where the statistical traffic source characterization is not fully known. On the contrary, they appear to be a promising, powerful and practical way to compensate the high variability of LRD traffic, and therefore improve the network efficiency.

The paper is organized as follows. In section 2

* This research is supported by European Community and MIUR in the frame of the Pollens project (ITEA, if00011a).

we describe the MBAC principles and we discuss and intuitively justify the important role of MBAC in the presence of self-similar traffic. Section 3 describes the specific MBAC algorithm that we have adopted, and the methods to evaluate self-similarity. Simulation results are given and discussed in section 4. Finally, concluding remarks are reported in section 5.

2 MBAC vs PBAC implementations

While PBAC methods rely on the a-priori knowledge of the statistical characterization of the offered traffic, MBAC algorithms base the decision whether to accept or reject an incoming call on run-time measurements on the traffic aggregate process. A large number of Measurement Based Admission Control algorithms have been proposed in literature (see for example [10], [11], [12]), but it appears that the measurement process, and in particular the length of the averaging periods and the way in which new flows are taken into account, are much more important than the specific admission criteria (either heuristic or theoretically founded) in determining the MBAC performance, in terms of throughput and packet loss (see [13]).

Moreover, it is frequently considered "obvious" that the ultimate goal of any MBAC scheme is to reach the "ideal" performance of a PBAC scheme. In fact, MBAC schemes are traditionally meant to approximate the operation of a parameter-based CAC (i.e estimate the status of the system, in terms of number $n(t)$ of admitted flows at time t). On the opposite, in this section we want to highlight that MBAC algorithms should have a broader theoretical target.

The well known problem of MBAC schemes relates to the fact that they cannot rely on the detailed a-priori knowledge of the statistical traffic characteristics, as this information is not easy supplied in an appropriate and useful form by the network customer. Therefore, their admission control decisions are based on an estimate of the network load obtained via a measurement process that runs on the accepted traffic aggregate. However, a closer look at the basic principles underlying MBAC suggests that, in particular traffic conditions, these schemes might outperform traditional PBAC approaches. In particular we support the thesis that MBAC schemes are not just "approximations" of PBAC, but they are *in principle superior* to traditional CAC schemes when self-similarity comes into play.

This superiority can be intuitively justified by comparing the simulation traces presented in figures 1 and 2, in which two typical traffic behaviors, respectively related to a PBAC scenario and to an MBAC one, are depicted. The figures report both the normalized

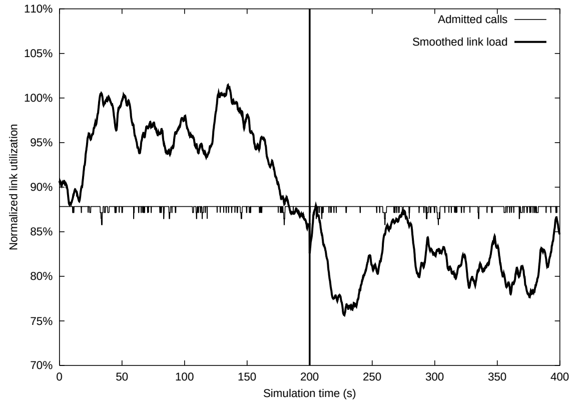

Fig. 1. Traditional Admission Control operation

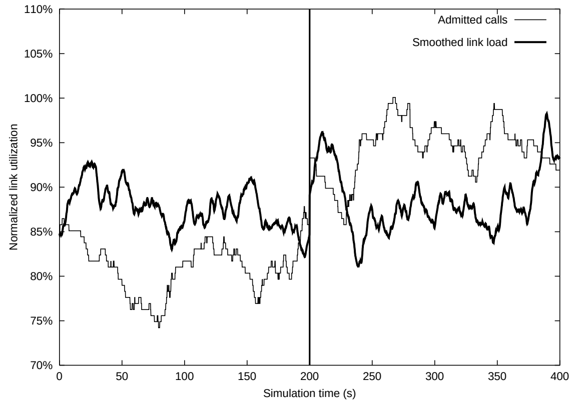

Fig. 2. Measurement-Based Admission Control operation

number of accommodated calls, and the smoothed link load, as measured by the autoregressive filter adopted in the MBAC, whose time constant is of the order of 10 seconds (see section 3).

Figure 1 reports results for an ideal PBAC. According to this scheme, a new flow is accepted only if the number of already admitted flows is lower than a maximum threshold N_t. In the simulation run a very high offered load ($\sim 600\%$) was adopted, so that the number of flows admitted to the link sticks, in practice, to the upper limit (i.e. 129 flows, in the example, corresponding to about 88% in link utilization).

The leftmost 200 simulation seconds, represented in Figure 1, show that, owing to Long Range Dependence of the accepted traffic, the load offered by the admitted sources is consistently well above the nominal average load. Traffic bursts even greater than the link capacity are very frequent. On the other hand, as shown by the rightmost 200 seconds, there are long periods of time in which the system remains underutilized. The criticality of self-similarity lies in the fact that the described situation occurs at time scales which dramatically affect the loss/delay performance.

Aggregate traffic behaves very differently when an MBAC scheme is odopted. Figure 2 reports results for the simple MBAC scheme described in section 3.2. In this case, new calls are blocked as long as the offered-load measurement is higher than a given threshold [1]. Specifically, we see from both leftmost and rightmost plots that the offered-load measurement fluctuates slightly around the threshold. However, long term traffic bursts are dynamically compensated by a significant decrease in the number of admitted calls (leftmost plot). Instead, if admitted calls continually emit below their nominal average rate, the number of admitted calls significantly increases. This "compensation" capability of MBAC schemes leads us to conclude that MBAC is very suited to operate in LRD traffic conditions: the quantitative analysis carried out in section 4, in fact, confirms this insight.

3 The simulation scenario

To obtain simulation results, we have developed a C++ event-driven simulator, based on a batch simulation approach. The simulation time was divided into 101 intervals, each lasting 300 simulated minutes, and results collected in the first "warm-up" time interval were discarded.

As in many other admission control works ([12], [13]), the network model consisted of a single bottleneck link. In fact, the basic performance aspects of MBAC are most easily revealed in this simple network configuration rather than in a multi-link scenario. Unless otherwise specified, the link capacity was set equal to 2 Mbps, and an infinite buffer size was considered. Thus, QoS is characterized by the delay experienced by data packets rather than packet loss as in [11]. The rationale for using delay instead of loss is twofold. Firstly, loss performance depends on the buffer size adopted in the simulation, while delay performance do not require a choice of buffer size. Secondly, the loss performance magnitude may be easily inferred, for a given buffer size, from the analysis of the distribution of the delay. Furthermore, in a very large buffer scenario, the system is forced to keep memory of non-smoothed traffic bursts and therefore performance are strongly degraded in the presence of high traffic variability. However, a comparison with results dealing with a finite buffer scenario is also presented.

As performance figures, we evaluated link utilization and delay distribution, summarized by the 99th delay percentile. The 95% confidence intervals were evaluated. In all cases, throughput results show a confidence interval consistently lower than 0.3%. Instead, despite the very long simulation time, higher confidence intervals occur for 99th delay percentile results: less than 5% for MBAC results, and as much as 25% for PBAC results (this is an obvious consequence of the self-similarity of the PBAC traffic aggregate). Nonetheless, even accounting for such uncertainty in the results, the PBAC and MBAC delay performance are clearly very different (see figure 4).

3.1 Traffic sources

For simplicity, we have considered a scenario composed of homogeneous ON/OFF flows. While in the ON state, a source transmits 1000 bit fixed size packets at a Peak Constant Rate (PCR) of 32 Kbps [2]. Conversely, while in the OFF state, it remains idle. The mean value of the ON and OFF periods have been set, respectively, equal to 1 s and 1.35 s (Brady model for voice traffic [14]). This results in an average source rate $r = 0.4255 \cdot E[PCR] \approx 13.6$ Kbps. ON and OFF periods were drawn from two Pareto distributions with the same shaping parameter $c = 1.5$ (infinite variance), which exhibit heavy tails[3], hence the traffic aggregate is self-similar ([4]).

Simulation experiments were obtained in a dynamic scenario consisting of randomly arriving flows. Each flow requests service from the network, and the decision whether to admit or reject the flow is taken by the specific simulated CAC. A rejected flow departs from the network without sending any data, and does not retry its service request again. The duration of an accepted flow is taken from a lognormal distribution [15] with mean 300 s and standard deviation 676 s (we adopted unitary variance for the corresponding normal distribution as reported in [15]), but call duration is extended to the end of the last ON or OFF period. Because of this, the real call-lifetime exhibits longer mean (320 s) and infinite variance. If the last burst were cut off, the process variance would become finite.

The flow arrival process is Poisson with arrival rate λ calls per second. For convenience, we refer to the normalized offered load ρ:

$$\rho = \lambda \frac{r\, T_{hold}}{C_{link}}$$

[1] in correspondance with the choise of N_t = 129 in PBAC, we select 89% in MBAC, since these selected value results in the same average throughputs in both PBAC and MBAC scenarios

[2] more precisely, the PCR is randomly generated in the small interval 31 to 33 Kbps, in order to avoid source synchronization effects at the packet level.

[3] A random variable is said to be "Heavy-Tailed" when its cumulative distribution function converges to $F(t) \sim 1 - at^{-c}$, as $t \to \infty$ with $1 < c < 2$, being a a constant. The cumulative distribution function of a Pareto Random Variable is $F(t) = 1 - \left(\frac{t+s}{s}\right)^{-c}$ for $t \geq 0$, where s is a scale parameter.

being r the mean source rate, T_{hold} the average call duration and C_{link} the link capacity. Depending on the simulation experiment, the arrival rate ranges from underload conditions (less than 50% of the link capacity) to severe overload conditions (up to 650%).

3.2 Adopted MBAC algorithm

We have adopted and tested a very basic MBAC approach. In fact, The results in [13] show that different MBAC schemes present similar performance, and, more importantly, our goal is to show that the introduction of measurement in the admission control decision is the key to obtaining performance advantages in comparison to the PBAC approach, rather than the careful design of the MBAC algorithm. In this perspective the simpler the MBAC scheme is, the more general the conclusions are. Our MBAC implementation can be described as follows: a discrete time scale is adopted, with sample time $T = 100$ ms. Let $X(k)$ be the load, entering the link buffer during the time slot k, and $B(k)$ its smoothed version, the bandwidth estimate, obtained by using a first order autoregressive filter:

$$B(k) = \alpha B(k-1) + (1-\alpha)X(k)$$

We chose $\alpha = 0.99$, corresponding to about 10 s time constant in the filter memory.

If a call requests admission during the slot $k+1$, the call is admitted only if the estimated bandwidth $B(k)$ is less than a predetermined percentage of the link bandwidth. By tuning this percentage, performance figures can be obtained for various accepted load conditions.

Moreover, when a new flow is admitted, a step input is offered to the system, and a transient phase occurs, in which the load is underestimated, thus not reflecting the presence of a new flow ([11], [16]). A solution to prevent this performance-impairing situation is to artificially increase the load estimate to account for the new flow. Specifically, in our implementation, the actual bandwidth estimate $B(k)$ is updated by adding the average rate of the flow:

$$B(k) := B(k) + r$$

3.3 Statistical analysis of self-similarity

The Hurst parameter H is able to quantify the self-similarity of the accepted traffic aggregate. For a wide range of stochastic processes $H = 0.5$ corresponds to uncorrelated observations, $H > 0.5$ to LRD processes and $H < 0.5$ to Short Range Dependence processes.

In order to evaluate H, we used three different methods. All methods receive in input a realization $X(i)$ of the discrete-time stochastic process representing the load offered, during a 100 ms time window, to the link buffer by the accepted traffic aggregate.

Aggregate Variance. The original series $X(i)$ is divided into blocks of size m and the aggregated series $X^{(m)}(k)$ is calculated as:

$$X^{(m)}(k) = \frac{1}{m} \sum_{i=(k-1)m+1}^{km} X(i) \qquad k = 1, 2, \ldots$$

The sample variance of $X^{(m)}(k)$ is an estimator of $Var\left(X^{(m)}\right)$; asymptotically:

$$Var\left(X^{(m)}\right) \sim \frac{Var(X)}{m^{2(1-H)}}$$

R/S. For a time series $X(i)$, with partial sum $Y(n) = \sum_{i=1}^{n} X(i)$, and sample variance $S^2(n)$, the R/S statistics or the rescaled adjusted range, is given by:

$$\begin{aligned} \frac{R}{S}(n) &= \frac{1}{S(n)} \max_{0 \le p \le n} \left(Y(p) - \frac{p}{n}Y(n)\right) + \\ &\quad - \frac{1}{S(n)} \min_{0 \le p \le n} \left(Y(p) - \frac{p}{n}Y(n)\right) \end{aligned}$$

Asymptotically:

$$E\left\{\frac{R}{S}(n)\right\} \sim Cn^H$$

Wavelet Estimator. The spectrum of a LRD process $X(t)$ exhibits power-law divergence at the origin:

$$W_X(f) \sim c_f |f|^{(1-2H)}$$

The method, proposed in [17], recovers the power-law exponent $1 - 2H$ and the coefficient c_f turning to account the following relation

$$E\left\{d_X^2(j, l)\right\} = 2^{j(1-2H)} c_f C$$

where $d_X(j, l) =< X, \psi_{l,j} >$ are the coefficients of the discrete wavelet transform of the signal $X(t)$, i.e. its projections on the basis functions $\psi_{l,j}$, constructed by the mother wavelet through scaling and translation (2^j and l are respectively the scaling and the translation factor). A MatLab implementation is freely distributed ([18]).

The three methods described allow to compute the Hurst parameter through a simple linear regression in a log-log diagram.

A comprehensive overview of LRD statistical-parameters estimation, including the first two methods, can be found in [19]. The first two methods have the

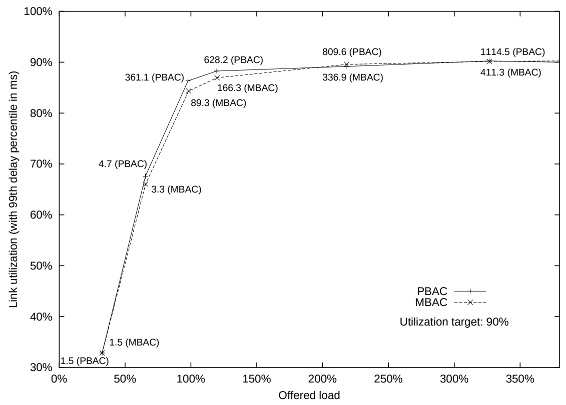

Fig. 3. Link utilization vs offered load

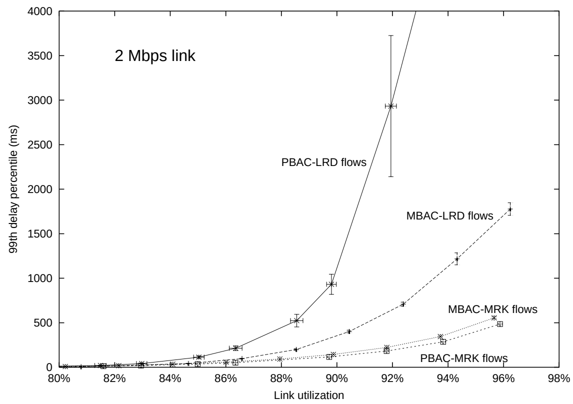

Fig. 4. Delay performance vs link utilization (2 Mbps link)

advantage of being simple and practical to implement, but often exhibit poor statistical properties ([20]). The wavelet-based joint estimator, according to [17], displays reliable and robust statistical performance. Besides, if the process $X(i)$ is gaussian, this estimator is able to provide confidence intervals for the Hurst parameter.

A common problem is to determine over which scales LRD property exists, or equivalently the alignment region in the logscale diagrams. Using the fit test of the matlab tool [18] we determined for our traces the range from 2000 s -11th octave- to 250000 s -18th octave- (the two last octaves were discarded because there were too few values). All the three methods were applied over this scale.

4 Performance evaluation

Every CAC scheme has some adjustable parameters that allow the network operator to set a suitable *utilization target* and a consequent QoS provisioning. In the present case, both for ideal PBAC and MBAC algorithms, a higher setting of the threshold value results in an increased system throughput, at the expense of delay performance. By adjusting this parameter, the proposed CAC rules can be designed to be more aggressive or conservative with regard to the number of flows admitted.

Results presented in figure 3 were obtained by setting the PBAC and MBAC tuning parameters so that a target 90% link-utilization performance is achieved in offered traffic overload conditions. The figure compares the throughput/delay performance (99th delay percentiles, measured in ms, are numerically reported) of MBAC and PBAC, versus the normalized offered load. Minor differences can be noted in the capability of the considered schemes to achieve the performance target. A much more interesting result is the significantly lower MBAC delay versus the PBAC one.

Rather than varying the offered load, figure 4 compares MBAC and PBAC by plotting their QoS performance versus the link utilization (following [13], the QoS versus utilization curve is called *Performance Frontier*). Specifically, the figure reports the delay/utilization performance frontiers of PBAC and MBAC in terms of 99th delay percentiles. The figure depicts results obtained for both LRD and Markovian flows, when CAC thresholds range from low to full utilization, while the offered load is very high ($\sim$ 600%). It is shown that better performance are obtained using a Markovian traffic model, but it is also emphasized the remarkable performance improvement provided by MBAC with respect to PBAC in LRD assumptions, especially for large link utilization. Under Markovian assumptions PBAC acts slightly better than MBAC, in fact no memory arises in offered traffic process, thus the best bandwidth control simply consists in monitoring the number of admitted connections. Instead, with LRD flows, MBAC performance frontiers assume intermediate values, better than PBAC-LRD performance frontiers but worst than the curves obtained under Markovian assumptions. Thus, MBAC appears to be more robust than PBAC to the traffic statistical properties.

We argue that the impressive performance enhancement of MBAC over PBAC is due to the beneficial effect of MBAC in reducing the self-similarity of the accepted traffic aggregate.

Further results obtained by considering a finite buffer scenario are depicted in figure 5 and 6, where buffers having respectively 100 and 2000 packets length, are considered. Note that in the former case, in which a short buffer of 100 packets is considered, PBAC performance is comparable to the MBAC one only in extreme underload condition or conversely in a harsh over-utilized scenario. In the latter case, with a 2000 packets buffer, MBAC performance does not saturate at all, while in the PBAC scenario a greater

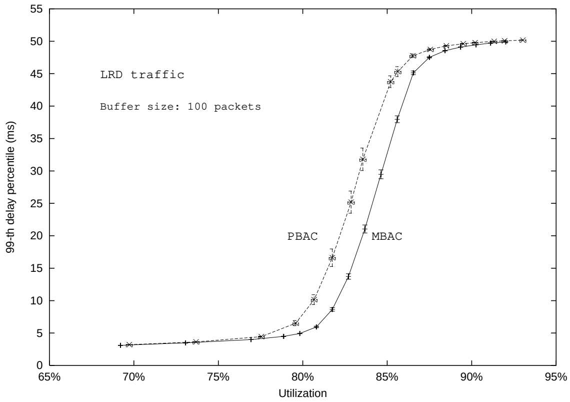

Fig. 5. Delay performance vs link utilization, using a 100 packets buffer

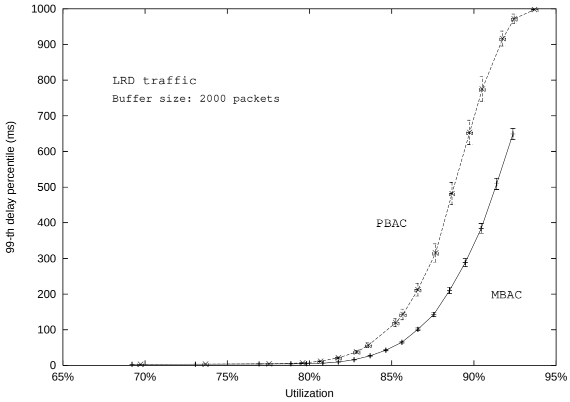

Fig. 6. Delay performance vs link utilization, using a 2000 packets buffer

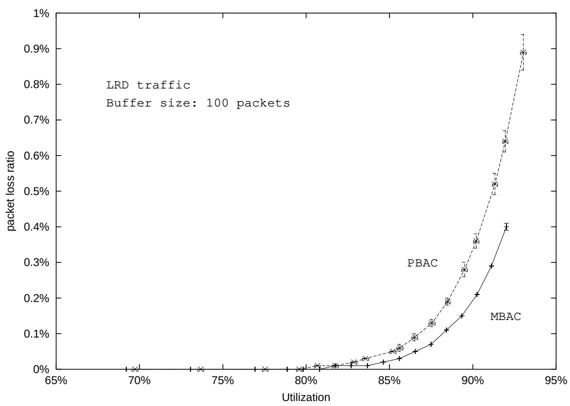

Fig. 7. Packet loss ratio in high offered load condition, using a 100 packet buffer

buffer length should be needed in order to avoid losses occourring when the buffer is entirely filled. These figures show, in the same conditions as before, i.e in very high offerd load conditions, that the performance gain resulting from MBAC scheme adoption is not impaired by a finite buffer application. Moreover figure

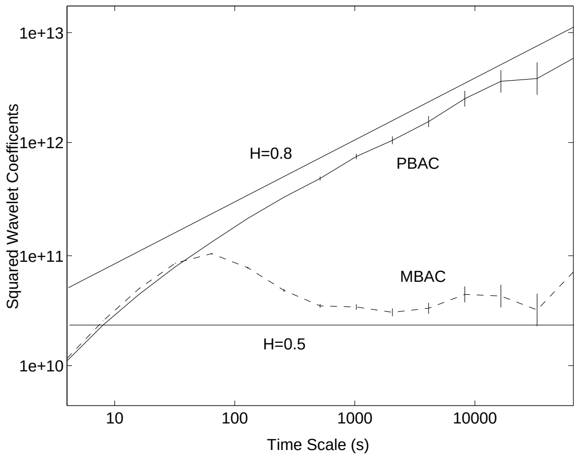

Fig. 8. Wavelet coefficients plot

7 shows how MBAC approach allows a significantly improvement in packet loss ratio, which is controlled well under 1% even in a high offered load and a high link utilization scenario.

To quantify the time behavior of the two PBAC and MBAC traffic aggregate time series, figure 8 plots the estimated squared wavelet coefficients $d_x^2(j,l)$ versus the basis-function time scale. 95% confidence interval under gaussian assumption are depicted. While the two curves exhibit similar behavior for small values of the aggregation scale, the asymptotic slope of the PBAC plot is very different from the MBAC one. For reference purposes, the lines corresponding to $H = 0.50$, and $H = 0.80$ are also plotted in the figure. Note that the figure 8 appears to suggest that the MBAC-controlled traffic is not self-similar (Hurst parameter close to 0.5). An interesting consideration is that in figure 8 the MBAC curve departs from the PBAC curve at a time scale of the order of about 100 seconds. Although a complete understanding of the emergence of such a specific time scale is outside the scope of the present paper, we suggest that it might have a close relationship with the concept of "critical time scale" outlined in [12].

The Hurst-parameter estimates are reported in tables I and II, with the corresponding CAC settings (the maximum call number for PBAC and the maximum link utilization for MBAC), and the achieved link utilization. For the wavelet estimates 95% confidence interval are also indicated (see section 3.3). The three methods described in section 3.3, provide congruent estimates. Results are impressive, and show that the Hurst parameter decreases from about 0.75, in the case of PBAC, to about 0.5 for MBAC. It is interesting to note that 0.75 is the Hurst-parameter value theoretically calculated in [3], [4] and [21] under different as-

PBAC				
Thresh (calls)	Thrput %	Hurst Variance	Hurst R/S	Hurst Wavelet
105	71.8	0.73	0.79	0.78 [0.74,0.82]
115	78.3	0.74	0.78	0.80 [0.76,0.84]
125	84.5	0.71	0.79	0.75 [0.71,0.79]
130	88.7	0.78	0.76	0.75 [0.71,0.79]
135	91.7	0.72	0.72	0.77 [0.74,0.81]
140	94.7	0.78	0.80	0.74 [0.70,0.78]

TABLE I

HURST-PARAMETER ESTIMATE FOR PBAC CONTROLLED TRAFFIC (INFINITE BUFFER SIZE)

MBAC				
Thresh (util%)	Thrput %	Hurst variance	Hurst R/S	Hurst Wavelet
70	69.1	0.55	0.48	0.55 [0.51,0.58]
78	76.9	0.58	0.54	0.58 [0.54,0.62]
86	84.6	0.55	0.51	0.60 [0.56,0.64]
90	88.5	0.60	0.52	0.57 [0.53,0.60]
94	92.4	0.51	0.46	0.56 [0.52,0.60]
96	94.3	0.58	0.52	0.58 [0.54,0.62]

TABLE II

HURST-PARAMETER ESTIMATE FOR MBAC CONTROLLED TRAFFIC (INFINITE BUFFER SIZE)

Thresh (util%)	Hurst PBAC	Hurst MBAC
70	0.75 [0.72,0.80]	0.51 [0.47,0.55]
80	0.74 [0.70,0.78]	0.60 [0.56,0.63]
85	0.75 [0.71,0.79]	0.51 [0.48,0.55]
90	0.75 [0.71,0.79]	0.51 [0.48,0.55]
94	0.79 [0.75,0.82]	0.51 [0.48,0.55]

TABLE III

HURST-PARAMETER WAVELET ESTIMATE FOR PBAC AND MBAC, USING A 100 PACKETS BUFFER

sumptions, when a flow has heavy-tailed periods of activity/inactivity with a shaping parameter $c = 1.5$ (the formula is $H = (3-c)/2$). We note that, as expected, the Hurst parameter does not depend on the link utilization. Finally, table III reports the Hurst parameter estimate obtained via the Wavelet method, when a finite and short buffer is employed. Even if these results relate to a short buffer scenario (100 packets instead of infinite), the values reported in table III for H are very similar to those of tables I and II, where an infinite buffer size was adopted. In conclusion tables II and III quantitatively supports our thesis that self-similarity is a marginal phenomenon for MBAC controlled traffic (the achieved Hurst parameter is very close to 0.5, which represents Short Range Dependent traffic).

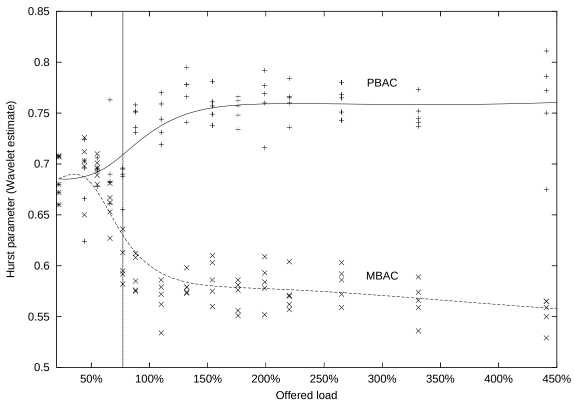

Fig. 9. Hurst parameter vs offered load

Most of the previous results were obtained under constant overload conditions. The aim of the figure 9 is to show the behavior of the two CAC schemes in a wide range of offered load conditions, with a fixed target link-utilization (the thresholds chosen, 110 connection for PBAC and 77% for MBAC, give very similar throughput performance under the same offered load). The vertical dashed line corresponds to this target. Hurst-parameter was estimated by the wavelet estimator. When the offered load is below the target, the Hurst-parameter estimates[4] are quite similar because MBAC and PBAC do not enforce any rejection. By the way, in this situation, no need of access control arises and delay/loss performance copes with high QoS requirements. Instead, the effect of CAC rules becomes evident when the offered load exceeds the target utilization: the PBAC curve approaches to $H = 0.75$, while the MBAC one decays and approaches to non LRD values. Moreover, the uncertainty of statistical results is shown by plotting several points for each simulated scenario, obtained with different seeds for the random generator.

5 Conclusions

Unlike traditional PBAC, the degree of self-similarity in the MBAC controlled traffic, resulting from the superposition of heavy-tailed Measured Based Admission Controlled flows seems to be very marginal. For this reason MBAC allows to minimize the performance impaiments due to the bursty nature of self-similar flows, thus making the system robust to statistical traffic properties. A greater QoS can be achieved using MBAC than PBAC schemes in order to control LRD flows. Moreover, even if traffic statistics

[4] In accordance with the fit test of the matlab tool [18] the LRD hypothesis on the range from 2000 s to 250000 s (see section 3.3) should not be rejected.

are almost Markovian, no significative differences can be noted using MBAC or PBAC.

We feel that there are two important practical implications of our study. Firstly, our study support the thesis that MBAC is not just an approximation of traditional CAC schemes, useful when the statistical pattern of the offered traffic is uncertain. On the contrary, we view MBAC as a value-added traffic engineering tool that allows a significant increase in network performance when offered traffic shows long range dependence. Secondly, provided that the network is ultimately expected to offer an admission control function, which we recommend should be implemented via MBAC, our results seem to question the practical significance of Long Range Dependence, the widespread usage of self-similar models in traffic engineering, and the consequent network oversizing.

References

[1] V.Paxson and S.Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling", IEEE/ACM Transactions on Networking, Vol. 3 No. 3, pp. 226-244, June 1995.

[2] M.E.Crovella and A.Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes", IEEE/ACM Transactions on Networking, 5(6):835–846, December 1997.

[3] M.Greiner, M.Jobmann, and C.Klüppelberg Telecommunication Traffic, Queueing Models, and Subexponential Distributions. Blauer Bericht TUM M9805, Technische Universitẗ München, May 1998

[4] W.Willinger, M.Taqqu, R.Sherman, D.Wilson, "Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level", IEEE/ACM Transactions on Networking, Vol. 5, No. 1, pp. 71-86, February 1997.

[5] A.Erramilli, O.Narayan, W.Willinger, "Experimental Queueing Analysis with Long Range Dependent Packet Traffic", IEEE/ACM Transactions on Networking, April 1996.

[6] K.Park, G.Kim, M.Crovella, "On the Effect of Traffic Self-similarity On Network Performance", Proceedings of SPIE's International Symposium and Education Program on Voice, Video, and Data Communications, November 1997, Dallas, Texas.

[7] P.R.Morin, "The Impact of Self-Similarity on Network Performance Analysis", Ph.D. Dissertation, Carleton University, Dec. 1995

[8] M.Grossglauser and J.Bolot, "On the Relevance of Long Range Dependence in Network Traffic", IEEE/ACM Transactions on Networking, 1998.

[9] Z.Sahinoglu, S.Tekinay, "On Multimedia Networks: Self-Similar Traffic and Network Performance", IEEE Communications Magazine, January 1999, pp. 48-52.

[10] R.Gibbens, F.P.Kelly, "Measurement-Based Connection Admission Control", Proc. of 15th International Teletraffic Congress, June 1997

[11] S.Jamin, P.B.Danzig, S.Shenker, L.A.Zhang, "A measurement-based admission control algorithm for integrated services packet networks", IEEE/ACM Transactions on Networking Vol. 5, No. 1, Feb 1997, pp. 56-70.

[12] M.Grossglauser, D.Tse, "A Framework for Robust Measurement Based Admission Control", IEEE/ACM Transactions on Networking Vol. 7, No. 3, July 1999.

[13] L.Breslau, S.Jamin, S.Shenker, "Comments on the Performance of Measurement Based Admission Control Algorithms", Proc. of IEEE Infocom 2000, Tel Aviv, Israel, March 2000.

[14] P.T.Brady, "A statistical analysis of on-off patterns in 16 conversations", Bell System Technical Journal, vol. 47, pp. 73–91, Jan. 1968.

[15] V.Bolotin, "Modeling Call Holding Time Distributions for CCS Network Design and Performance Analysis", IEEE Journal on Selected Areas in Communications 12, 3 (Apr. 1994), 433–438.

[16] S.Floyd, "Comments on measurements based admission control for controlled-load services", Technical report, July 1996.

[17] P.Abry, D.Veitch, "Wavelet Analysis of Long-Range Dependent Traffic", IEEE Transactions on Information Theory, 44(1), pp. 2-15, January 1998.

[18] http://www.emulab.ee.mu.oz.au/~darryl/secondorder_code.html

[19] J.Beran, "Statistics for long-memory processes" Chapman & Hall, 1994

[20] M.S.Taqqu, V.Teverosky, W.Willinger, "Estimators for long-range dependence: an empirical study", Fractals 1996

[21] N.Likhanov, B.Tsybakov, and N.Georganas, "Analysis of an ATM buffer with Self-Similar ("fractal") Input Traffic", Proc. IEEE INFOCOM '95, pp.985-992, 1995