Is Admission-Controlled Traffic Self-Similar?*

Giuseppe Bianchi, Vincenzo Mancuso and Giovanni Neglia

Università di Palermo, Dipartimento di Ingegneria Elettrica
Viale delle Scienze, 90128 Palermo, Italy
bianchi@elet.polimi.it, {vincenzo.mancuso,giovanni.neglia}@tti.unipa.it

Abstract. It is widely recognized that the maximum number of heavy-tailed flows that can be admitted to a network link, while meeting QoS targets, can be much lower than in the case of markovian flows. In fact, the superposition of heavy-tailed flows shows long range dependence (self-similarity), which has a detrimental impact on network performance. In this paper, we show that long range dependence is significantly reduced when traffic is controlled by a Measurement-Based Admission Control (MBAC) algorithm. Our results appear to suggest that MBAC is a value added tool to improve performance in the presence of self-similar traffic, rather than a mere approximation for traditional (parameter-based) admission control schemes.

1 Introduction

The experimental evidence that packet network traffic shows self-similarity[1] was first given in [1], where a thorough statistical study of large Ethernet traffic traces was carried out. This paper stimulated the research community to explore the various taste of self-similarity. This phenomenon has been observed in wide area Internet traffic and many causes that contribute to self-similarity for both TCP and UDP traffic aggregates have been now more fully understood [2–5].

In this paper, we focus our attention on traffic generated by sources non-reactive to network congestion (e.g. real-time multimedia streams). The traffic aggregate offered to a network link results from the superposition of several individual flows. It has been proven [6] that self-similarity or Long Range Dependence (LRD) arises when individual flows have heavy-tailed[2] periods of activity/inactivity. This result is valid asymptotically as the number of sources increases.

We are interested in the practical implications of self-similarity on the design of Call Admission Control (CAC) schemes. In this paper we assume, for convenience, a traffic scenario composed of homogeneous flows. In these conditions, a

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

1 In this paper we use the terms self-similarity and long range dependence in an interchangeable fashion, because we refer to asymptotic second order self-similarity (for details [7]).

2 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.

traditional (parameter-based) CAC rule simply checks that the number of admitted flows never exceeds a maximum threshold N_t. This threshold is selected so that target Quality of Service (QoS) requirements (e.g. loss ratio, delay percentiles, etc.) are met. In what follows we refer to this CAC scheme as MAXC (Maximum number of Calls).

A large amount of work (see [7]) has shown that self-similarity has a detrimental impact on network performance. For the same link capacity and buffer size scenario, the Quality of Service (i.e. loss/delay performance) experienced by LRD traffic results worse than that experienced by Short Range Dependent (SRD) traffic, e.g. modelled as Markov processes. The straightforward interpretation of these results, in terms of traditional CAC, is that self-similarity is a key factor which reduces the maximum number N_t of flows that can be admitted.

We argue that the above interpretation is questionable, as it does not account for recent progress in admission control schemes, and specifically the emergence and increasing popularity of Measurement-Based Admission Control (MBAC) approaches [8–11]. Unlike traditional CAC methods, which rely on 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.

The aim of this paper is to present results which show that MBAC approaches appear capable of smoothing the self-similarity of the accepted traffic aggregate. In this sense, MBAC approaches are not merely "approximations" of ideal CAC schemes in situations where the statistical traffic source characterization is not fully known. On the contrary, this paper shows that MBAC schemes are an effective and important way to cope with the high variability of LRD traffic, and their adoption leads to significant performance advantages with respect to traditional CAC schemes (refer to [11] for an initial insight on the performance advantages of MBAC in an LRD traffic scenario).

The rest of the paper is organized as follows. In section 2 we briefly describe the MBAC principles and we discuss the important role of MBAC in the presence of self-similar traffic. The specific MBAC algorithm adopted and the methods to evaluate self-similarity are described in section 3. Numerical results are presented and discussed in section 4. Finally, concluding remarks are given in section 5.

2 Measurement Based Admission Control

It is frequently assumed that the ultimate MBAC goal is to reach the "ideal" performance of a parameter-based CAC scheme. In fact, MBAC schemes are traditionally meant to approximate the operation of a parameter-based CAC. 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.

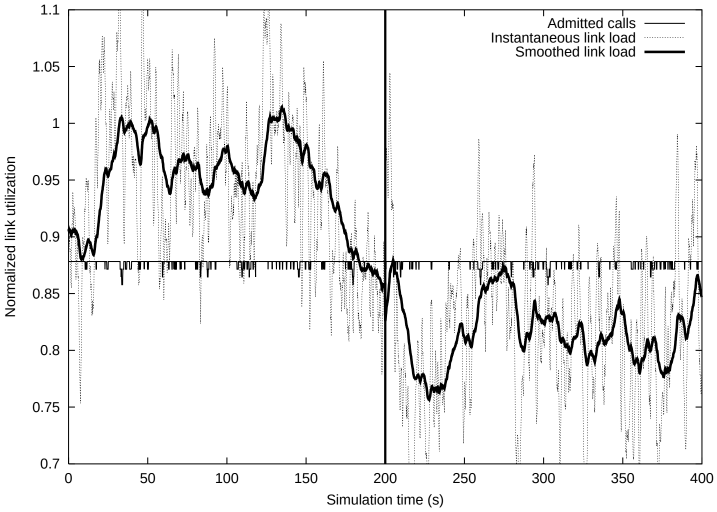

Fig. 1. Traditional Admission Control operation

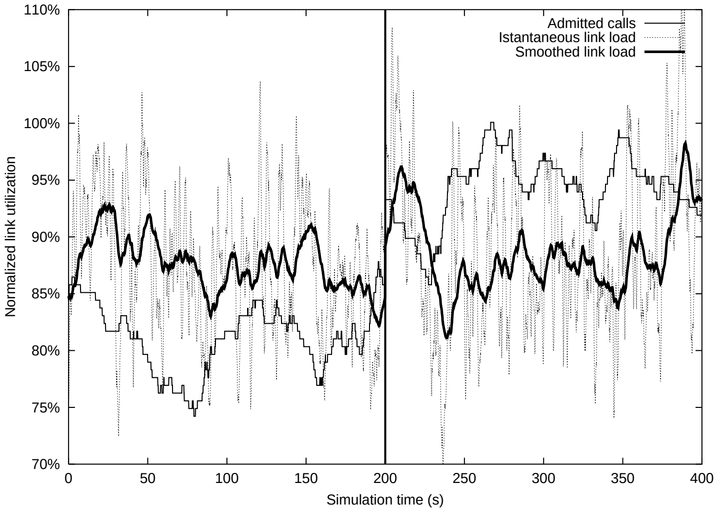

Fig. 2. Measurement-Based Admission Control operation

However, a closer look at the basic principles underlying MBAC suggests that, in particular traffic conditions, these schemes might outperform traditional parameter-based CAC approaches. An initial insight into the performance benefits of MBAC versus parameter-based algorithms in an LRD traffic scenario is given in [11]. In this paper, we present additional results that confirm the superiority of MBAC and we justify them showing that MBAC algorithms are able to reduce the self-similarity of the traffic aggregate generated by the admitted heavy-tailed sources. In other words, we argue that MBAC schemes are not just "approximations" of parameter-based CAC, but they are *in principle superior* to traditional CAC schemes when self-similarity comes into play.

An intuitive justification can be drawn by looking at the simulation traces presented in figures 1 and 2 (simulation details are described in section 3). Each figure reports two selected 200 s simulation samples, which for convenience have been placed adjacently. The y-axis represents the normalized link utilization. The figures report: i) the number of accommodated calls normalized with respect to the link capacity; ii) the instantaneous link load, for graphical convenience averaged over a 1 s time window, and iii) the smoothed link load, as measured by the autoregressive filter adopted in the MBAC, whose time constant is of the order of 10 seconds.

Figure 1 reports results for a parameter-based CAC scheme (MAXC). 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 N_t was set to 129, which corresponds to a target link utilization of about 88%, and a very high offered load (650%) was adopted. As a consequence, the number of flows admitted to the link sticks, in practice, to the upper limit.

The leftmost 200 simulation seconds, represented in Figure 1, show that, owing to LRD 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 under-utilized. The criticality of self-similarity lies in the fact that the described situation occurs at time scales, e.g. the one shown in the figure, which dramatically affect the loss/delay performance.

A very different situation occurs for MBAC schemes. 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 89% (the values 129 in MAXC and 89% in MBAC were selected so that the resulting average throughputs were the same). In this case, 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 of the number of admitted calls (leftmost plot). The opposite situation occurs when the admitted calls persistently emit under their nominal average rate: indeed the rightmost plot shows that in these periods the number of admitted calls significantly increases. This "compensation" capability of MBAC schemes leads us to conclude that MBAC is well-suited to operating 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. A batch simulation approach was adopted. The simulation time is divided into

101 intervals, each lasting 300 simulated minutes, and results collected in the first "warm-up" time interval are discarded.

As in many other admission control works [10, 11], the network model consists of a single bottleneck link. The reason is that the basic performance aspects of MBAC are most easily revealed in this simple network configuration rather than in a multi-link scenario. The link capacity was set equal to 2 Mbps, and an infinite buffer size was considered. Thus, QoS is characterized by the delay (average and 99th delay percentiles) experienced by data packets rather than packet loss as in in [11]. The rationale for using delay instead of loss is threefold. Firstly, loss performance depends on the buffer size adopted in the simulation runs, while delay performance does not require a choice of buffer size (we have actually used infinite 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, which can be well summarized via selected delay percentiles. Thirdly, and most importantly, a limited buffer size acts as a smoothing mechanism for traffic bursts. Large packet losses, occurring during severe and persistent traffic bursts (as that expected for self-similar traffic), have a beneficial congestion control effect on the system performance. Conversely, in a very large buffer scenario, the system is forced to keep memory of non-smoothed traffic bursts and therefore performance is further degraded in the presence of high traffic variability[3].

As our performance figures, we evaluated link utilization (throughput) and delay distribution, summarized, for convenience of presentation, by the average and 99th delay percentile. The 95% confidence intervals have been evaluated. In all cases, throughput results show a confidence interval always 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 MAXC results (this is an obvious consequence of the self-similarity of the MAXC traffic aggregate).

3.1 Traffic Sources

For simplicity, we have considered a scenario composed of homogeneous flows. Each traffic source is modelled as an ON/OFF source. While in the ON state, a source transmits 1000 bit fixed size packets at a Peak Constant Rate (PCR) randomly generated in the small interval 31 to 33 Kbps (to avoid source synchronization effects at the packet level). Conversely, while in the OFF state, it remains idle. The mean value of the ON and OFF periods were set, respectively,

[3] Specifically, this justifies the very different performance results we obtain in high utilization conditions when compared with the loss-utilization performance frontier presented in [11] for LRD sources. In that paper, unlike our results presented in figure 4, it appears that performance of MBAC schemes tend to converge to the performance of traditional CAC schemes - i.e. the MAXC algorithm - as the utilization increases. A theoretical justification for this behavior can be found in [16], where the authors derive a formula to estimate the "correlation horizon" (which results to scale in linear proportion to the buffer size), beyond which the impact on loss performance of the correlation in the arrival process becomes nil.

equal to 1 s and 1.35 s (Brady model for voice traffic). 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 $c = 1.5$ (so they exhibit heavy-tails).

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 admission control algorithm. 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 [12] with mean 300 s and standard deviation 676 s (we adopted unitary variance for the corresponding normal distribution as reported in [12]), 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 $\rho = \lambda \cdot r \cdot T_{hold}/C_{link}$, being r the mean source rate, T_{hold} the average call duration and C_{link} the link capacity.

3.2 Measurement-Based Admission Control Algorithm

Rather than using complex MBAC proposals, we have implemented a very basic MBAC approach. The rationale for the choice of a very simple MBAC scheme is twofold. Firstly, it has been shown [11] that different MBAC schemes behave very similarly in terms of throughput/loss performance. It appears that 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. Secondly, and more importantly, our goal is to show that the introduction of measurement in the admission control decision is the key to obtain performance advantages versus the MAXC 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.

The specific MBAC implementation is described as follows. A discrete time scale is adopted, with sample time $T = 100$ ms. Let $X(k)$ be the load, in bits/sec, entering the link buffer during the time slot k, and let $B(k)$ be a running bandwidth estimate, smoothed by a simple 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.

Consider now a call requesting admission during the slot $k+1$. The call is admitted 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.

An additional well-known issue in MBAC algorithm design [9] is that, when a new flow is admitted, the slow responsiveness of the load estimate will not immediately reflect the presence of the new flow. A solution to prevent this performance-impairing situation is to artificially increase the load estimate to account for the new flow. In our implementation, the actual bandwidth estimate $B(k)$ is updated by adding the average rate of the flow (i.e. $B(k) := B(k) + r$).

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 SRD processes.

In order to evaluate H, we used the well known three methods described below. 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. The methods adopted are:

1. **Aggregate Variance** [13]. 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:

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

2. **Rescaled Adjusted Range (R/S)** [13]. 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:

$$\frac{R}{S}(n) = \frac{1}{S(n)} \left[\max_{0 \le p \le n} \left(Y(p) - \frac{p}{n} Y(n) \right) - \min_{0 \le p \le n} \left(Y(p) - \frac{p}{n} Y(n) \right) \right]$$

Asymptotically:

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

3. **Wavelet Estimator** [14] (see [15] for a freely distributed Matlab implementation). We recall that the spectrum of an LRD process $X(t)$ exhibits power-law divergence at the origin $W_X(f) \sim c_f |f|^{(1-2H)}$. The method 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).

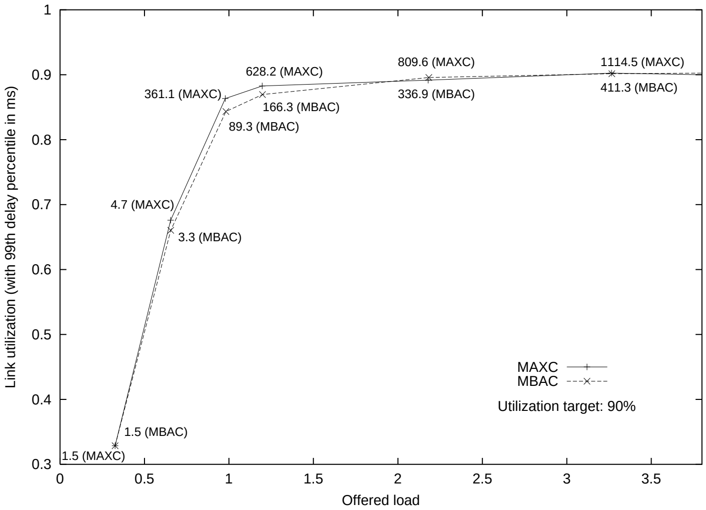

Fig. 3. Link utilization vs offered load

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 [15] 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

A problem arising in the comparison of different CAC schemes is the definition of a throughput/performance operational trade-off. In general, CAC schemes have some tunable parameters that allow the network operator to set a suitable *utilization target* and a consequent QoS provisioning. For example, in the case of the ideal MAXC algorithm, a higher setting of the threshold value results in an increased system throughput, at the expense of delay performance. By adjusting these parameters, 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 MAXC and MBAC tuning parameters so that a target 90% link-utilization performance is achieved in overload conditions. The figure compares the throughput/delay performance (99th delay percentiles, measured in ms, are numerically reported) of MBAC and MAXC, versus the normalized offered load. Minor differences can be noted in the capability of the considered schemes to achieve the performance target (as expected, MAXC converges faster than MBAC to the utilization target). A much more interesting result is the significantly lower MBAC 99th delay performance versus the MAXC one.

It is restrictive to limit the investigation to a single level of performance, but it is preferable to compare different CAC schemes for a wide range of link

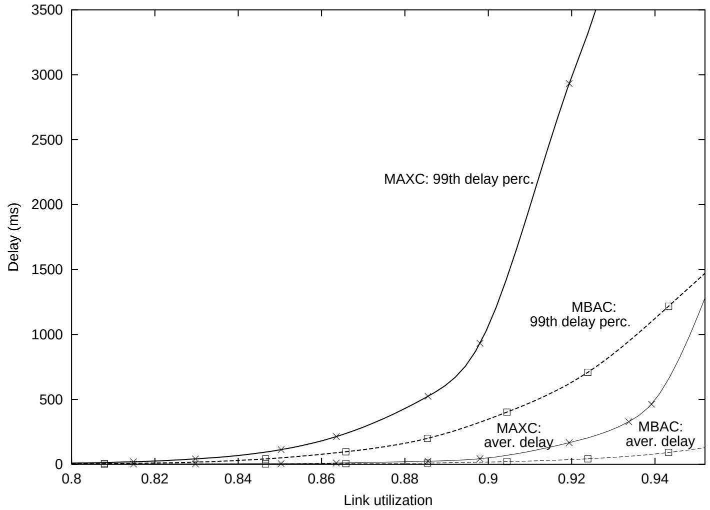

Fig. 4. Delay performance vs link utilization

utilization targets (and, correspondingly, QoS performance), obtained by varying the CAC threshold parameters. Unless otherwise specified, all results presented in what follows are obtained in large overload conditions (650% offered load).

Rather than varying the offered load, figure 4 compares MBAC and MAXC by plotting their QoS performance versus the link utilization (following [11], the QoS versus utilization curve is called *Performance Frontier*). Specifically, the figure reports the delay/utilization performance frontiers of MAXC and MBAC. Both average and 99th delay percentiles are compared. The figure shows that the performance improvement provided by MBAC is remarkable, especially for large link utilization.

We argue that the performance enhancement of MBAC over MAXC is due to the beneficial effect of MBAC in reducing the self-similarity of the accepted traffic aggregate. To quantify this statement, tables 1 and 2 report the Hurst-parameter estimates obtained with the three methods described in section 3.3, along with the corresponding CAC settings (maximum call number for MAXC; link utilization threshold for MBAC), and the achieved link utilization[4]. We see that the methods provide congruent estimates. Results are impressive, and show that the Hurst parameter decreases from about 0.75, in the case of MAXC, to

[4] As we said, these results were obtained with an offered load equal to 650%. It may be remarked that the different results of MBAC and MAXC in shaping traffic reduce in lighter load conditions, and vanish for very low offered loads (when neither MBAC nor MAXC enforce call rejections). By the way, in this situation, traffic self-similarity is irrelevant in terms of performance, as traffic QoS requirements are met. Additional results, not presented here due to space constraints, show that MBAC capability to reduce self-similarity becomes evident as soon as the offered load approaches the target utilization threshold, and Hurst-parameter values reach those presented in table 2 as soon as the offered load becomes 10-20% greater than this target.

Table 1. Hurst-parameter estimate for MAXC controlled traffic

MAXC				
Thresh (calls)	Thrput %	Hurst Variance	Hurst R/S	Hurst Wavelet
105	71.8	0.73	0.79	0.78
110	74.7	0.77	0.78	0.76
115	78.3	0.74	0.78	0.80
120	81.5	0.73	0.79	0.76
125	84.5	0.71	0.79	0.75
127	86.8	0.77	0.77	0.75
130	88.7	0.78	0.76	0.75
132	90.3	0.68	0.73	0.75
135	91.7	0.72	0.72	0.77
137	93.4	0.71	0.77	0.76
140	94.7	0.78	0.80	0.74

Table 2. Hurst-parameter estimate for MBAC controlled traffic

MBAC				
Thresh (util%)	Thrput %	Hurst variance	Hurst R/S	Hurst Wavelet
70	69.1	0.55	0.48	0.55
74	73.0	0.60	0.55	0.57
78	76.9	0.58	0.54	0.58
82	80.8	0.56	0.50	0.58
86	84.6	0.55	0.51	0.60
88	86.6	0.52	0.49	0.53
90	88.5	0.60	0.52	0.57
92	90.4	0.54	0.52	0.56
94	92.4	0.51	0.46	0.56
96	94.3	0.58	0.52	0.58
98	96.2	0.58	0.53	0.57

about 0.5 for MBAC. It is interesting to note that 0.75 is the Hurst-parameter value theoretically calculated in [6] when a flow has heavy-tailed periods of activity/inactivity with a shaping parameter $c = 1.5$ (the formula is $H = (3 - c)/2$). In conclusion, table 2 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 SRD traffic).

To quantify the time behavior of the two MAXC and MBAC traffic-aggregate time series, figure 5 reports a log-log plot of the aggregate variance, computed as described in section 3.3. While the two curves exhibit similar behavior for small values of the aggregation scale, the asymptotic slope of the MAXC plot is very different from the MBAC one. We recall that the asymptotic slope β is related to the Hurst parameter by $\beta = 2H - 2$. The lines corresponding to $H = 0.50$, $H = 0.55$, $H = 0.75$ and $H = 0.80$ are plotted in the figure as reference comparison. Note that the figure 5 appears to suggest that the MBAC-controlled traffic is not self-similar (Hurst parameter close to 0.5).

Similar considerations can be drawn, with greater evidence, by looking at figure 6, which reports a log-log plot of the estimated squared wavelet coefficients $d_x^2(j, l)$ versus the basis-function time scale. The figure shows that, for large time scales, the MBAC-controlled traffic plot tends to lay on a horizontal line (the asymptotic slope γ is related to the Hurst parameter by $\gamma = 1 - 2H$, and thus a horizontal line corresponds to $H = 0.50$, the $H = 0.80$ case is also plotted in the figure as reference comparison).

Finally, figures 5 and 6 show that the MBAC curve departs from the MAXC curve at a time scale of the order of about 100 seconds. Although a thorough 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 [10].

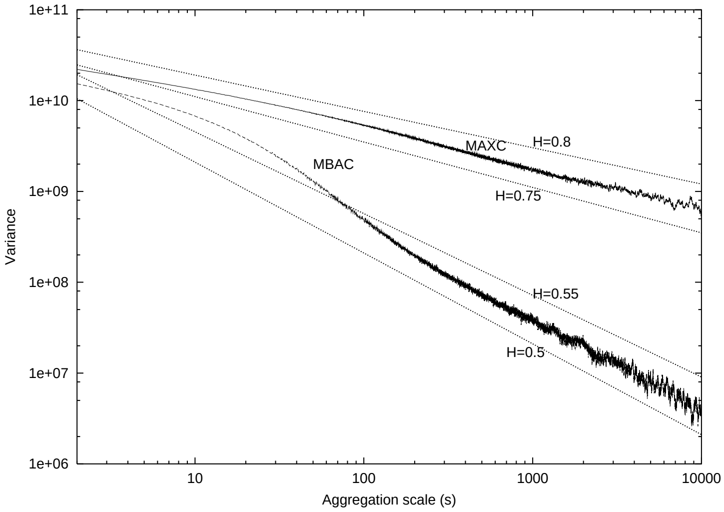

Fig. 5. Aggregated variance plot

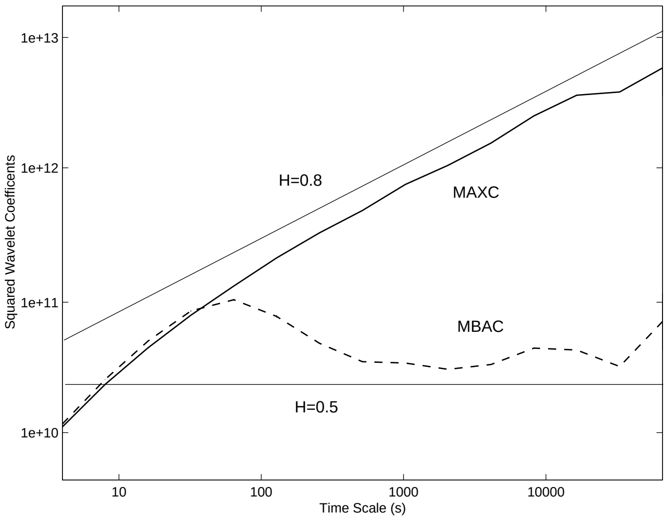

Fig. 6. Wavelet coefficients plot

5 Conclusions

The results presented in this paper appear to suggest that the traffic aggregate resulting from the superposition of Measurement-Based Admission Controlled flows shows a very marginal long range dependence. This is not the case for traffic controlled by a traditional parameter-based admission control scheme.

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.

Acknowledgements

The authors want to thank the graduate student Vito Imburgia for his collaboration in the software development.

References

1. W.Leland, M.Taqqu, W.Willinger, D.Wilson, "On the Self-Similar Nature of Ethernet Traffic", Trans. on Networking, Vol. 2, No. 1, pp. 1-15, Feb. 1994.
2. J.Beran, R.Sherman, W.Willinger, and M.S.Taqqu, "Variable-bit-rate video traffic and long-range dependence", IEEE Transactions on Communications, 1995
3. K.Park, G.T.Kim, M.E.Crovella, "On the Relationship Between File Sizes, Transport Protocols, and Self-Similar Network Traffic", Proceedings of the International Conference on Network Protocols, pp. 171-180, October, 1996.
4. M.E.Crovella and A.Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes", Trans. on Networking, 5(6):835–846, Dec. 1997.
5. A.Feldmann, A.C.Gilbert, W.Willinger, "Data network as cascades: Investigating the multifractal nature of the Internet WAN Traffic", Computer Communications Review 28 (1998)
6. W.Willinger, M.Taqqu, R.Sherman, D.Wilson, "Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level", Trans. on Networking, Vol. 5, No. 1, pp. 71-86, Feb. 1997.
7. [edited by] K.Park and W.Willinger "Self-Similar Network Traffic and Performance Evaluation", Wiley Inter-Science, 2000
8. R.Gibbens, F.P.Kelly, "Measurement-Based Connection Admission Control", Proc. of 15th International Teletraffic Congress, June 1997
9. S.Jamin, P.B.Danzig, S.Shenker, L.A.Zhang, "A measurement-based admission control algorithm for integrated services packet networks", Trans. on Networking Vol. 5, No. 1, Feb 1997, pp. 56-70.
10. M.Grossglauser, D.Tse, "A Framework for Robust Measurement Based Admission Control", Trans. on Networking Vol. 7, No. 3, July 1999.
11. 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.
12. 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.
13. J.Beran, "Statistics for long-memory processes", Chapman & Hall, 1994
14. P.Abry, D.Veitch, "Wavelet Analysis of Long-Range Dependent Traffic", IEEE Transactions on Information Theory, 44(1), pp. 2-15, January 1998.
15. `http://www.emulab.ee.mu.oz.au/~darryl/secondorder_code.html`
16. M.Grossglauser and J.Bolot, "On the Relevance of Long Range Dependence in Network Traffic", Trans. on Networking, 1998.