On the Self-Similarity of Measurement-Based Admission Controlled Traffic*

Giuseppe Bianchi, Vincenzo Mancuso, Giovanni Neglia
Università di Palermo, Dipartimento di Ingegneria Elettrica
Viale delle Scienze, 90128 Palermo, Italy

***Abstract*— In this paper, we focus on an admission controlled traffic scenario. Flows, characterized by heavy-tailed ON/OFF periods, are admitted to a network link according to a Measurement Based Admission Control algorithm. Our simulation results show that the Long Range Dependence of the accepted traffic aggregate is marginal, particularly when compared with that resulting from a traffic aggregate accepted by a parameter-based admission control scheme. Our results appear to suggest that Measurement Based Admission Control is a value added tool to dramatically improve performance in the presence of self-similar traffic, rather than being a mere approximation for traditional (parameter-based) admission control schemes.**

I. INTRODUCTION

The experimental evidence that packet network traffic shows self-similarity 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 also observed in wide area Internet traffic [2], [3], and many of the causes that contribute to self-similarity for both TCP [4], [5] and UDP [6] traffic aggregates have been now more fully understood.

In this paper, we focus our attention on traffic generated by sources non-reactive to network congestion (e.g. real-time multimedia streams). We assume, for convenience, a traffic aggregate scenario resulting from the superposition of homogeneous flows. Long Range Dependence (LRD) or (asymptotic second order) self-similarity arises when a flow has Heavy Tailed (HT) periods of activity/inactivity [7]. In particular, [8], [9] give insights into the relation between LRD of aggregate traffic and heavy-tailedness. There is evidence [3], [8] that human as well as computer sources behave as HT ON/OFF sources, so this result should be considered a physical explanation of the traffic self-similarity - independent of network/protocol characteristics [3] - rather than a mere way to generate self-similar traces.

Many works [10], [11], [12], [13] show that self-similarity has a severe detrimental impact on network performance. Consider a link capacity, a buffer size, and a given (constant) number N of superposed offered flows. The QoS (e.g. loss, delay percentiles, etc) experienced by HT flows results much worse than that experienced by flows whose activity/inactivity periods are drawn from exponential distributions. By repeating this study for various values of N, it is straightforward to design a "traditional" (parameter-based) Connection Admission Control (CAC) scheme, in what follows referred to as MAXC (Maximum number of Calls). The MAXC 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 significantly depends on the flow statistic parameters. As a consequence of LRD, the maximum number N_t of HT flows that can be admitted to a link may be much lower than in the case of markovian (MRK) flows.

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

The aim of this paper is to understand what happens when a parameter-based CAC rule is replaced by a Measurement Based Admission Control (MBAC) scheme. Several MBAC schemes have been proposed in [14], [15], [16], [17], [18]. Our simulation results allow to draw several conclusions. Firstly, MBAC schemes provide superior performance than parameter-based CACs, when LRD flows are considered. Secondly, unlike parameter-based CACs, 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 rest of the paper is organized as follows. Section II briefly describes MBAC principles and its importance in the presence of self-similar traffic. The specific MBAC scheme adopted and the methods to evaluate self-similarity are described in section III. Numerical results are presented and discussed in section IV. Conclusions are drawn in section V.

II. MEASUREMENT BASED ADMISSION CONTROL

While traditional CAC methods rely on the a-priori knowledge of the statistical characterization of the offered traffic, MBAC schemes base the decision whether to accept or reject an incoming call on run-time measurements on the traffic aggregate. It has been shown [18] that different MBAC schemes behave very similarly in terms of throughput/loss performance. Following [18], 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 theoretical) in determining how close MBAC schemes approach ideal CAC performance.

It is frequently considered "obvious" that the ultimate goal of any MBAC scheme 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 (i.e. by estimating the status of the system). They cannot rely on the detailed a-priori knowledge of the statistical traffic characteristics, as this information is not easy supplied by the network customer. Therefore, their admission control (AC) 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 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 [18]. In this paper, we present additional

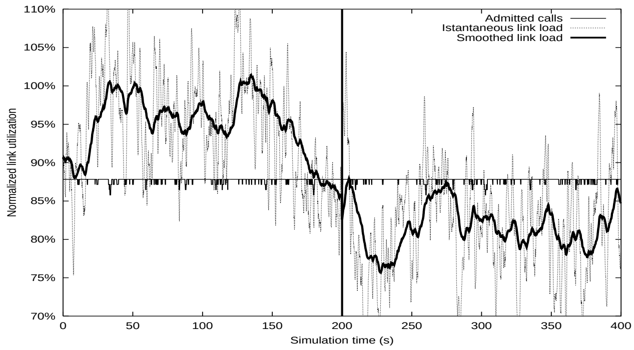

Fig. 1. Traditional Admission Control operation

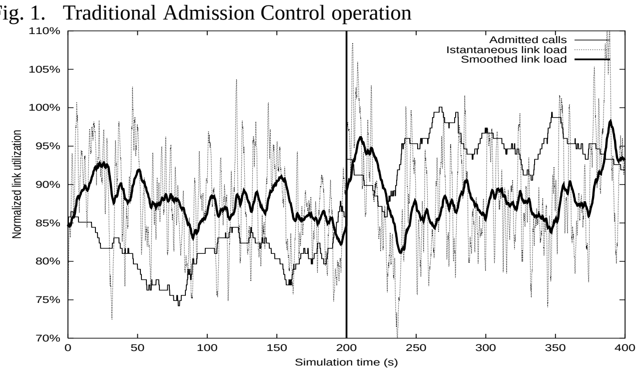

Fig. 2. Measurement-Based Admission Control operation

results that confirm the superiority of MBAC and, in addition, we justify them showing that MBAC algorithms are able to reduce the self-similarity of the traffic aggregate generated by the admitted HT sources. In other words, we support the thesis 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 simulations presented in figures 1 and 2 (the simulation model is described in section III). Each figure shows 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 normalized number of accommodated calls; ii) the 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 plots results for the MAXC scheme. In this scheme, a new flow is accepted only if the number of already admitted flows is lower than a threshold N_t. In the simulation run N_t has been set to 129, which corresponds to a target link-utilization of about 88%. 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 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 which dramatically affect the loss/delay performance.

MBAC schemes behave very differently, as reported in figure 2 for the simple scheme described in section III-B. In this case, new calls are blocked when the measured offered-load is higher than 89% [1]. We see that the accepted load 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). The opposite situation occurs when the admitted calls continually emit below their nominal average rate (rightmost plot): in these periods 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, as quantitatively confirmed in section IV.

III. The Simulation Scenario

A batch event-driven C++ simulation approach was adopted. Traces has been divided into 101 intervals, each lasting 300 minutes. Results collected in the first "warm-up" interval were discarded. As in many other admission control works [17], [18], the network model consists of a bottleneck link, as 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 packets rather than loss as in [16]. The rationale for using delay instead of loss is threefold. Firstly, the loss depends on the buffer size adopted, while delay performance do not require a choice of buffer size. Secondly, the loss may be easily inferred, for a given buffer size, from the analysis of the delay distribution. Thirdly, and most importantly, a limited buffer size smooths traffic bursts. Large packet losses, occurring during severe and persistent traffic bursts (as that expected for LRD traffic), have a beneficial congestion control effect on the system performance [12]. Conversely, with a very large buffer, the system is forced to keep memory of non-smoothed traffic bursts and so performance is further degraded in the presence of high traffic variability.

We evaluated link utilization (throughput) and 99th delay percentile with 95% confidence intervals. In all cases, throughput results show an uncertainty well below 0.3%. Instead, despite the very long simulation time, higher confidence intervals occur for delay results: less than 5% for MBAC results, and as much as 25% for MAXC results (as a consequence of the LRD of the MAXC traffic aggregate). However, even accounting for such uncertainty in the results, the MAXC and MBAC delay performance are clearly very different (see figure 4 and 5).

A. Traffic Sources

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) 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. ON and OFF periods have a mean, respectively, equal to 1 s and 1.35 s (Brady's voice model). 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 shaping parameter $c = 1.5$ (infinite variance), which exhibit heavy tails[2], hence

[1]the values 129 in MAXC and 89% in MBAC were selected so that the resulting average throughputs were the same.

[2]i.e. its cumulative distribution function (cdf) approaches $F(t) \sim 1 - at^{-c}$, as $t \rightarrow \infty$ with $1<c<2$. The cdf 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.

the traffic aggregate is self-similar [8].

A dynamic scenario, consisting of randomly arriving flows, was simulated. Each flow requests service from the network, and the admission decision is taken by the specific simulated CAC. Rejected flows does not retry their service request again. The duration of an accepted flow is taken from a lognormal distribution [19] with mean 300 s and standard deviation 676 s, but call duration is extended to the end of the last ON or OFF period. Thus 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 rate λ calls/s. For convenience, we refer to the normalized offered load $\rho = \lambda \frac{r\, T_{hold}}{C_{link}}$, 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 C_{link}) to severe overload (up to 650%).

B. Adopted MBAC Algorithm

We decided to implement a very basic MBAC proposal. In fact, the results in [18] show that different MBAC schemes present similar performance; moreover, and more importantly, our goal is to show that the introduction of measurement in the AC decision is the key to obtaining performance advantages in comparison to the MAXC approach, rather than the careful design of the MBAC scheme: the simpler the MBAC scheme is, the more general the conclusions are.

In our simple MBAC implementation, a discrete time scale is adopted, with sample time $T = 100$ ms. Let $X(k)$ be the load, in bps, 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 with a time constant of about 10 seconds:

$B(k) = \alpha B(k-1) + (1-\alpha)X(k)$, with $\alpha = 0.99$.

A call requesting admission during the slot $k+1$ is admitted if the estimated bandwidth $B(k)$ is below a threshold. By tuning this threshold, performance figures can be obtained for various accepted load conditions. An additional well-known issue in MBAC algorithm design [14], [16] 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. This performance-impairing situation can be prevented by artificially increasing 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:

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

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

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:

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

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

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 [20], 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 $d_X(j,l)$ are the wavelet coefficients for $X(t)$.

With the three methods described the H can be calculated through linear regression in a log-log diagram. The first two methods have the advantage of being simple and practical to implement, but often exhibit poor statistical properties [21]. Maximum Likelihood Estimator (MLE) techniques, not used in this work, have better statistical properties, but involve minimization procedures which are complex and slow and need parametric assumptions. The wavelet-based joint estimator is faster than MLE techniques and, according to [20], displays statistical performance comparable to MLE techniques, when their parametric assumptions are satisfied, and greater robustness under departures from them. Besides, if the process $X(i)$ is gaussian, this estimator provides confidence intervals for H.

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 [22] we determined for our traces the range from about 2000 s to 250000 s. All the three methods were applied over this scale.

IV. Performance Evaluation

CAC schemes have some tunable parameters that allow the network operator to set a suitable *utilization target* and a consequent QoS provisioning even in overload conditions. With the exception of figure 3, unless otherwise specified, we have evaluated the performance in the presence of large overload conditions (650% offered load). Figure 3 shows that performance tend to stabilize as the offered load grows. Results presented in figure 3 were obtained by setting the MAXC and MBAC tuning parameters to achieve an asynptotic 90% target link-utilization performance. The figure compares the throughput/delay performance (99th delay percentiles 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. A much more interesting result is the significantly lower MBAC 99th delays versus the MAXC ones.

It is restrictive to limit the investigation to a single level of performance, but it is preferable to compare different CAC

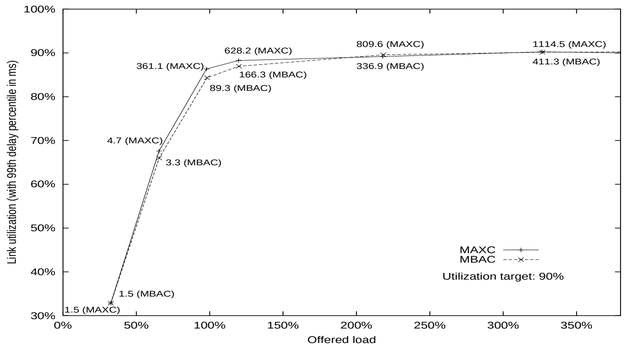

Fig. 3. Link utilization vs offered load

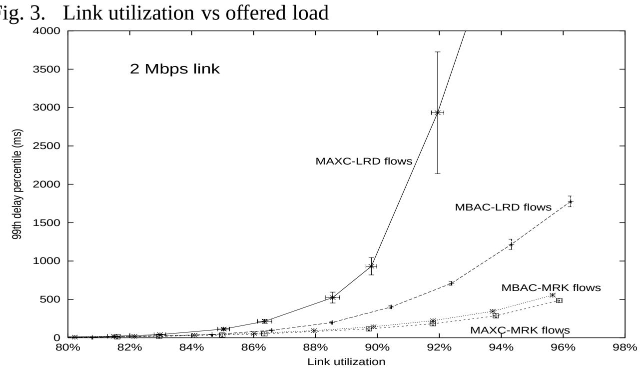

Fig. 4. Delay performance vs link utilization (2 Mbps link)

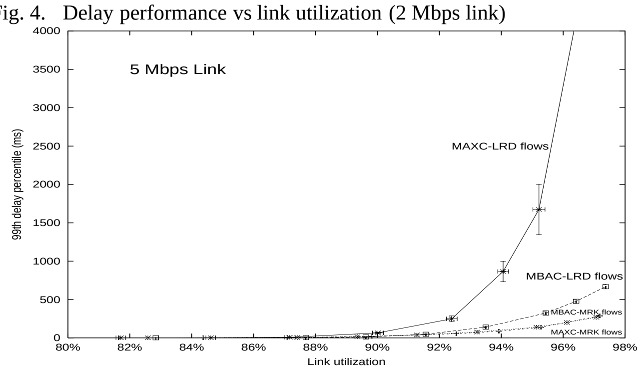

Fig. 5. Delay performance vs link utilization (5 Mbps link)

schemes for a wide range of link utilization targets (and, correspondingly, QoS performance), obtained by varying the CAC threshold parameters. Figure 4 compares MBAC and MAXC by plotting their QoS performance versus the link utilization (following [18], the QoS versus utilization curve is called *Performance Frontier*). The figure reports the delay/utilization performance frontiers of MAXC and MBAC in terms of 99th delay percentiles. The figure reports the results obtained for both LRD and Markovian flows. It is shown that better performance are obtained using a Markovian traffic model, but it is also enlightened the remarkable performance improvement provided by MBAC with respect to MAXC in LRD assumptions, especially for large link utilization. Considering Markovian flows, MAXC 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, between MAXC-LRD and Markovian curves. Thus, MBAC appears to be more robust than MAXC to the traffic statistical properties. Moreover, in figure 5 the performance frontiers are plotted for a 5 Mbps link. Beside the general performance improvement in comparison to the 2 Mbps link scenario shown in figure 4, one can see that MBAC behavior, with Markovian traffic, is closer

	Thresh (calls)	Thrput%	H-Variance	H-R/S	H-Wavelet
	105	71.8	0.73	0.79	0.78 [0.74,0.82]
M	115	78.3	0.74	0.78	0.80 [0.76,0.84]
A	125	84.5	0.71	0.79	0.75 [0.71,0.79]
X	130	88.7	0.78	0.76	0.75 [0.71,0.79]
C	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 MAXC CONTROLLED TRAFFIC

	Thresh (util%)	Thrput%	H-Variance	H-R/S	H-Wavelet
	70	69.1	0.55	0.48	0.55 [0.51,0.58]
M	78	76.9	0.58	0.54	0.58 [0.54,0.62]
B	86	84.6	0.55	0.51	0.60 [0.56,0.64]
A	90	88.5	0.60	0.52	0.57 [0.53,0.60]
C	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

to the MAXC behavior, since the traffic granularity is reduced and the impact of a flow erroneously admitted is less significant. We argue that the impressive 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 the time behavior of the two MAXC and MBAC traffic aggregate time series, figure 6 reports a log-log plot of the aggregate variance, computed as described in section III-C. 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, suggesting that the MBAC-controlled traffic is not self-similar ($H \sim 0.5$). We recall that the asymptotic slope β is related to H 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.

Similar considerations can be drawn by looking at figure 7, which 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. For reference purposes, the lines corresponding to $H = 0.50$, and $H = 0.80$ are also plotted in the figure. An interesting consideration is that in both figures 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 [17].

The H estimates are reported in tables I and II, with the corresponding CAC settings (the maximum call number for MAXC and the maximum link utilization for MBAC), and the achieved link utilization. For the wavelet estimates 95% confidence interval are also indicated. The three methods described in section III-C, provide congruent estimates. Results are impressive, and show that H decreases from about 0.75, in the case of MAXC, to about 0.5 for MBAC. It is interesting to note that 0.75 is the H value theoretically calculated in [7], [8] and [9] under different assumptions, when a flow has HT periods of activity/inactivity with a shaping parameter $c = 1.5$ (the formula is $H = (3 - c)/2$). We note that, as expected, H does not depend on the link utilization. In conclusion, table II quantitatively supports our thesis that self-similarity is a marginal phenomenon for MBAC controlled traffic (H close to 0.5).

Figure 8 was obtained by adjusting the offered load, with a fixed target link-utilization. H is a wavelet estimate, and the vertical dashed line corresponds to the target. When the offered load is below the target, the H estimates are quite similar because MBAC and MAXC do not enforce any rejection.

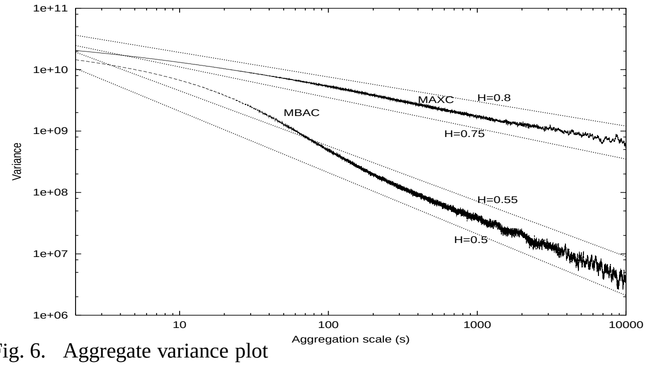

Fig. 6. Aggregate variance plot

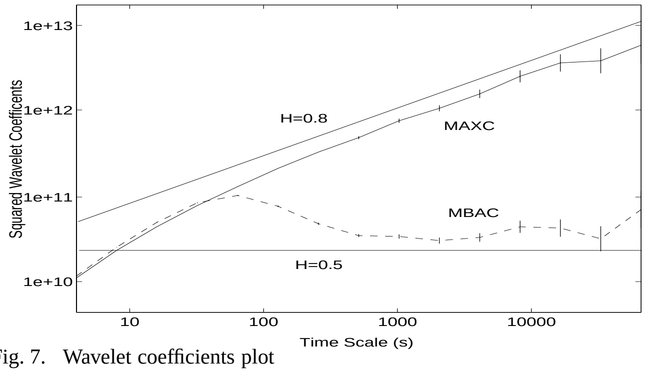

Fig. 7. Wavelet coefficients plot

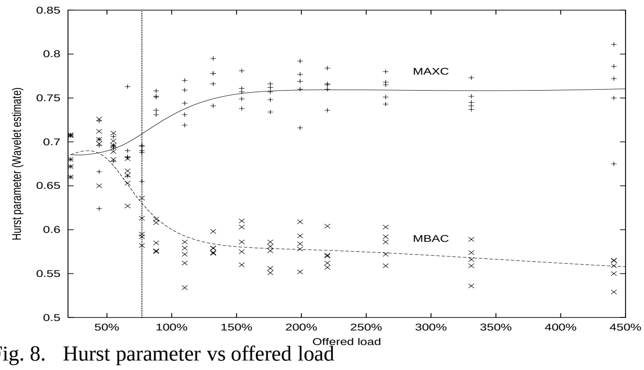

Fig. 8. Hurst parameter vs offered load

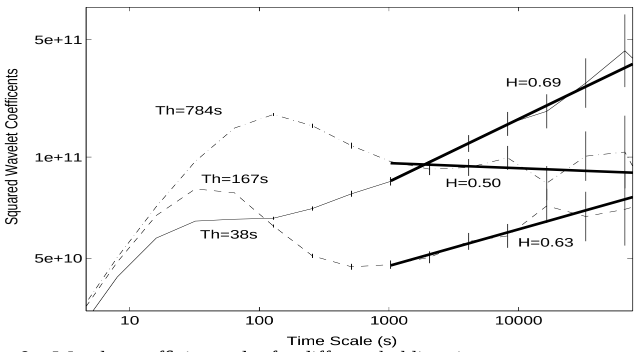

Fig. 9. Wavelet coefficients plot for different holding time

By the way, in this situation, no need of access control arises and performance copes with high QoS requirements. The effect of CAC rules becomes evident when the offered load exceeds the target utilization: the MAXC curve approaches to $H = 0.75$, while the MBAC one decays and approaches to non-LRD values. The uncertainty of statistics is shown by plotting several points for each simulated scenario, obtained with different seeds for the random generator.

The impact of the connection duration is drawn in figure 9. The ability of reducing traffic LRD is more effective as the duration increases. In fact, MBAC measurements are not able to efficiently track traffic variability when the holding time is too short in comparison to the filter memory.

V. Conclusions

The results presented in this paper appear to suggest that the traffic aggregate resulting from the superposition of MBAC flows shows a very marginal self-similarity. This is not the case for traffic controlled by a traditional parameter-based CAC. 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 the network performance when offered traffic shows LRD. 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] W.Leland, M.Taqqu, W.Willinger, D.Wilson, "On the Self-Similar Nature of Ethernet Traffic", IEEE/ACM Trans. on Networking, 2(1), Feb '94.

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

[3] M.Crovella, A.Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes", IEEE/ACM Trans. on Networking, Dec '97.

[4] K.Park, G.T.Kim, M.E.Crovella, "On the Relationship Between File Sizes, Transport Protocols, and Self-Similar Network Traffic", Proc. of the International Conference on Network Protocols, pp. 171-180, Oct 1996.

[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] J.Beran, R.Sherman, W.Willinger, M.S.Taqqu, "Variable-bit-rate video traffic and long-range dependence", IEEE Trans. on Communications, '95

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

[8] 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 Trans. on Networking, Vol. 5 No. 1, Feb '97.

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

[10] K.Park, G.Kim, M.Crovella, "On the Effect of Traffic Self-similarity On Network Performance", Proc. of SPIE's Voice, Video, and Data Communication Conference, Nov '97, Dallas, Texas.

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

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

[13] Z.Sahinoglu, S.Tekinay, "On Multimedia Networks: Self-Similar Traffic and Network Performance", IEEE Communication Magazine, Jan '99.

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

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

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

[17] M.Grossglauser, D.Tse, "A Framework for Robust Measurement Based Admission Control", ACM Trans. on Networking, Vol. 7 No. 3, Jul '99.

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

[19] 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.

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

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

[22] www.emulab.ee.mu.oz.au/˜darryl/secondorder_ code.html