Sensitivity of ABR Congestion Control Algorithms to Hurst Parameter Estimates

Sven A. M. Östring[1], Harsha Sirisena[1], and Irene Hudson[2]

[1] Department of Electrical & Electronic Engineering
[2] Department of Mathematics & Statistics
University of Canterbury, New Zealand.
{s.ostring,h.sirisena}@elec.canterbury.ac.nz
i.hudson@math.canterbury.ac.nz

Abstract. Optimal linear predictors can be utilised in ABR control algorithms for the management of self-similar network traffic. However, estimates of the Hurst parameter are required to generate these predictors, and uncertainity in these estimates results in a potential mismatch of the predictors to the traffic. When mismatched congestion control algorithms are used within the network, the impact on the network system is greater queue lengths at buffers and more significant cell losses. The sensitivity of these algorithms to the Hurst parameter estimate is investigated both analytically and using simulations. It is shown that an asymmetry in the sensitivity occurs in the region where the Hurst parameter is significantly underestimated.

1 Introduction

A significant amount of research in the area of teletraffic modeling has been focused on proving the self-similarity of network traffic [1,2]. Methods of accurately estimating the Hurst parameter, the index of self-similarity, is a key issue within this area, and these range from relatively simple methods, such as the variance-time plots and R/S statistic analysis [3], to more sophisticated techniques, such as Whittle's estimator and estimators based on the wavelet transform citeBeran94,Veitch99. Having demonstrated that network traffic is self-similar, the subsequent step is to determine the impact of self-similarity on the network as a system. The performance of a queue is fundamentally different when the input process to the queue is self-similar [5], with the distribution of the queue length now being heavy-tailed. In the case of finite buffers, higher cell losses occur within the system and these losses decrease hyperbolically as the buffer size is increas ed, rather than the exponential decrease which occurs with Poisson processes [6].

These rather serious consequences of self-similar traffic has driven research, though a limited extent, to consider methods of recognising the characteristics of self-similarity within network resource management techniques [7,8,9].

G. Pujolle et al. (Eds.): NETWORKING 2000, LNCS 1815, pp. 36–48, 2000.
© Springer-Verlag Berlin Heidelberg 2000

Our research has focused on incorporating the characteristics of self-similarity into congestion control algorithms, and in particular rate-control mechanisms for the ABR service in ATM networks [10,11]. This work has demonstrate d that the buffer memory requirements and cell losses can be reduced if the control algorithms are developed using the stochastic structure of the self-similar background traffic, and that adaptive algorithms based on on-line Hurst parameter estimators can be implemented which track non-stationarities which occur in the traffic.

An area which has not been addressed by any known research work is that of investigating the effect of poor knowledge of the actual Hurst parameter when resource management algorithms are designed specifically with self-similarity in mind. In this paper, this sensitivity is investigated more thoroughly, both through analysis of the relative change in variance of prediction errors resulting from the mismatched algorithm and through simulations, where data sets are tested with a range of algorithms developed from different Hurst parameter values.

2 Self-similarity of Network Traffic

2.1 Concepts and Models

There is significant statistical evidence that a wide range of classes of network traffic is self-similar in nature. This means that there is no natural length of the bursts in the traffic, and the traffic remains bursty over a range of time scales (hence the term "self-similar"). The Hurst parameter H is the index of self-similarity of a process, and a process is categorized as long-range dependent (LRD) when the parameter lies in the interval $0.5 < H < 1.0$. Long-range dependence means that the correla tions within a process decrease hyperbolically rather than exponentially, so that the autocovariance function for LRD processes is non-summable. While the burstiness of LRD traffic can cause buffer overflows and losses, long-range dependence can be used to one's advantage in terms of prediction [10,12]. These large correlations mean that there is significantly more information within previous states regarding the current state in LRD processes than in short-range dependent (SRD) proc esses, and more accurate predictions can be achieved from appropriately filtering stored measurements of the process in the past.

There are a number of well-known models used for processes which display self-similarity. Fractional Brownian motion is the canonical example of a self-similar process, and its incremental process (called fractional Gaussian noise—fGn), has the autocovariance function:

$$\begin{aligned}\gamma(k) &= \frac{\sigma_0^2}{2}\left(|k+1|^{2H} - 2|k|^{2H} + |k-1|^{2H}\right),\ k \in Z \\ &\sim \sigma_0^2 H(2H-1)k^{2H-2} \text{ as } k \to \infty\end{aligned} \tag{1}$$

where the asymptotic behaviour of the autocovariance function shows that the process is long-range dependent. A self-similar process can be parsimoniously

represented by an fGn model, if the mean, variance and Hurst parameter of the process are known. Another important class of self-similar models is the fractional ARIMA family of models, which are a natural extension of the ARIMA(p, d, q) models where the differencing parameter d is allowed to assume fractional values.

2.2 Estimation of the Hurst Parameter

As stated in Section 2.1, the Hurst parameter plays a key role in characterizing the self-similarity of a process. Thus, the estimation of this parameter from a set of measurements is crucial in determining whether a process is self-similar, and has attracted a significant amount of research. Self-similarity manifests itself within data in three fundamental ways—slowly decaying variances as the data is aggregated, long-range dependence within the covariance structure and a spectral density which obeys a power-law with divergence near the origin. Methods of estimating the Hurst parameter are based on quantifying one of these behaviours, usually by estimating the slope of a graph based on some transformation of the data.

Two heuristic estimation techniques are variance-time analysis and the rescaled adjusted range (R/S) statistic analysis [3]. Variance-time analysis is based on the asymptotic relationship of the variance of sample averages $X^{(m)}$ of non-overlapping blocks of data of size m from the process X, with the relationship given by:

$$Var(X^{(m)}) \sim cm^{-\beta}, \text{ as } m \to \infty \tag{2}$$

with $0 < \beta < 1$ for LRD processes. The other well-known heuristic technique is R/S statistic analysis, where the R/S statistic exhibits the Hurst effect as described by the relationship:

$$\mathbf{E}[R(n)/S(n)] \sim cn^{H}, \text{ as } n \to \infty \tag{3}$$

with the parameter H typically about 0.7.

More refined methods of data analysis are based on maximum likelihood type estimates (MLE) and the use of periodograms which effectively transform the data into the frequency domain, to estimate the power-law behaviour near the origin. The well-known Whittle's estimator is an approximate MLE which is asymptotically normal and efficient. The Abry-Veitch (AV) estimator is a fast estimation technique based on the wavelet transform [4]. The wavelet domain is a natural framework from which to view self-similar processes due to the scaling behaviour which is common to both fields. The wavelet transform generates a set of detail coefficients $d_{\mathbf{x}}(j, k)$ from a data set, and for a LRD process $\mathbf{X}$, the variance of the detail coefficients at each level j is given by the relationship:

$$\mathbf{E}[d_{\mathbf{x}}(j, \cdot)^2] = C \cdot 2^{j(2H-1)} \tag{4}$$

with $C > 0$. An important property of the AV estimator is that it can be reformulated to generate on-line estimates of the Hurst parameter [13]. This is

because it is based on the fast pyramidal filter, which was originally intended for on-line applications. These on-line estimates can be used within adaptive ABR control mechanisms [11].

3 Using Hurst Parameter Estimates in ABR Rate Control

The potential impact of self-similar traffic within networks, such as greater queue lengths at nodes and increased cell losses, makes it necessary to incorporate the characteristics of self-similarity into resource management algorithms. In particular, we have investigated using the properties of long-range dependence to improve the accuracy of predictions of traffic levels in the network to develop congestion control algorithms for the ABR service [10]. Predictors are developed using the stoc hastic structure of the self-similar background traffic, and the estimated Hurst parameter is important in characterising this structure.

3.1 Modeling ABR Control

The concept of the ABR service is to utilise available bandwidth within the network by controlling the sources which agree to the conditions of the ABR service contract. This means that the network returns control information back to the ABR source regarding the allowable rate which the source can transmit at. The approach which has been used in our research is to determine optimal predictions of future traffic levels, and calculate the ABR rates from these predictions. The network model defined here is based on controlling the bandwidth of the outgoing link by aiming to achieve a specified link utilization. This follows the model proposed by Zhao and Li [14]. The congestion avoidance policy can be formulated as follows,

$$U(k) + R_b(k) = \rho\, C \tag{5}$$

where the control aim is to keep the offered load at a proportion ρ $(0 < \rho < 1)$ of the outgoing link capacity C.

The total rate-controlled bandwidth $U(k)$ is made up of the summation of the individual bandwidths used by the N rate-controlled connections. Each connection has its own round-trip delay δ_j through the network. Then the state variable of the system can be defined as deviation from the target utilization, that is

$$x(k) = \rho\, C - \left[\sum_{j=1}^{N} u_j(k - \delta_j) + R_b(k) \right] \tag{6}$$

We can further define the variable $W(k) = \rho\, C - R_b(k)$, thus resulting in the equation

$$x(k) = -\sum_{j=1}^{N} u_j(k - \delta_j) + W(k) \tag{7}$$

Equation (7) is in the form of a multi-input single-output (MISO) control system. The aim of this control system is to determine the inputs $u_j(k)$ so that the variance $E[x^2(k)]$ is minimized. However, a control system for a MISO system is computationally too expensive. To simplify the system, the available bandwidth is equally shared among rate-controlled connections. Thus, we can define:

$$w_j(k) = \frac{W(k)}{N},\ j = 1, \ldots, N \tag{8}$$

The system now becomes a collection of N subsystems, each with their own controller:

$$x_j(k) = -u_j(k - \delta_1) + w_j(k),\ j = 1, \ldots, N \tag{9}$$

where the round-trip delays have been ordered such that $\delta_1 \geq \delta_2 \geq \cdots \geq \delta_N$ without loss of generality.

Our control aim now is to minimize $E[x_j^2(k)]$. This is equivalent (refer to [15]) to requiring $\hat{x}_j(k) = 0$ for all k. Taking the expectation of (9) and setting $\hat{x}_j(k) = 0$, we have the following general control law for each subsystem:

$$u_j(k) = \hat{w}_j(k + \delta_j \mid W(m) :\ m \leq k) \tag{10}$$

Thus, we require the allowed rates of the individual source rates $u_j(k)$ to be equal to the predicted values of the system parameters $w_j(k + \delta_j)$, which are determined by the amount of bandwidth available in the outgoing link. This prediction can be achieved by using the self-similarity of the network traffic.

3.2 Optimal Prediction of Self-similar Processes

As we have defined our system model in Section 3.1, we require the prediction of the background network traffic which is traversing a particular node in the network to determine the desired rates of the controlled sources. This information experiences a delay δ in the network before the effects can be observed at the same node. Hence, we require a δ-step predictor. The long-range dependence property of network traffic can be employed to provide more accurate predictions . The optimal linear predictor $\mathbf{G}_\delta^*$ is of the form:

$$\hat{X}_{k+\delta} = \mathbf{G}^{*T}_{\delta} \mathbf{X}_M \tag{11}$$

where X_k is a covariance stationary stochastic process with zero mean, variance σ^2 and autocovariance function $\gamma^*(k)$ and $\mathbf{X}_M$ is the vector of stored traffic measurements $\{X_k, X_{k-1}, \ldots, X_{k-M+1}\}'$. M is the memory-length of the predictor. The solution is given in [16] and is found by taking the expectations $E(X_{k+\delta} X_{k+\delta-m})$ on both sides of (11) for $m = k - M + 1, \ldots, k$, resulting in the matrix equation:

$$\mathit{\Gamma}^* \mathbf{G}^*{}_\delta = \gamma^*_\delta \tag{12}$$

where $\Gamma^*(i,j) = \gamma^*(i-j)$ is the $M \times M$ covariance matrix, $\gamma^*_\delta = (\gamma^*(\delta), \gamma^*(\delta+1), \ldots, \gamma^*(\delta+M-1))'$ is M-values of the autocovariance function starting at lag δ and $\mathbf{G}_\delta$ is the prediction vector. The variance of the prediction errors is given by

$$v^*_\delta = \gamma^*(0) - {\gamma^*}_\delta^T {\Gamma^*}^{-1} {\gamma^*}_\delta \tag{13}$$

The predicted background traffic levels for traffic with non-zero mean μ_b are then calculated as:

$$\hat{R}^*_b(k+\delta) = \mu_b + {\mathbf{G}^*}_\delta^T (\mathbf{R_b}(k) - \mu_b) \tag{14}$$

3.3 Sensitivity of ABR Control to Hurst Parameter Estimates

Of course, absolute knowledge of the stochastic structure of the background traffic is not possible, and the stochastic structure must be estimated from observations of the traffic. Thus, predictors which are developed for the traffic are inevitably mismatched to some degree to that traffic. This effect of this mismatch can be determined analytical. Consider that the actual predictors are developed from an estimated autocovariance function:

$$\hat{\Gamma}\hat{\mathbf{G}}_\delta = \hat{\gamma}_\delta \tag{15}$$

The variance of the prediction errors using this mismatched predictor now becomes

$$\begin{aligned} \hat{v}_\delta &= \mathbf{E}\left[(R_b(k+\delta) - \hat{R}_b(k+\delta))^2\right] \\ &= \mathbf{E}[R_b(k+\delta)^2] - 2\mathbf{E}[R_b(k+\delta)\hat{R}_b(k+\delta)] + \mathbf{E}[\hat{R}_b(k+\delta)^2] \\ &= \gamma^*(0) - 2{\gamma^*}_\delta^T \hat{\Gamma}^{-1}\hat{\gamma}_\delta + \hat{\gamma}_\delta^T \hat{\Gamma}^{-1}\Gamma^*\hat{\Gamma}^{-1}\hat{\gamma}_\delta \end{aligned} \tag{16}$$

The relative increase in the variance of the prediction errors is

$$K_v = \frac{\hat{v} - v^*}{v^*} \tag{17}$$

which gives Theorem 1.

Theorem 1. *The relative increase in the variance of the prediction errors resulting from a linear predictor which is mismatched to the stochastic structure of the background traffic is given by:*

$$K_v = \frac{\hat{\gamma}_\delta^T \hat{\Gamma}^{-1}\Gamma^*\hat{\Gamma}^{-1}\hat{\gamma}_\delta + {\gamma^*}_\delta^T {\Gamma^*}^{-1}{\gamma^*}_\delta - 2{\gamma^*}_\delta^T \hat{\Gamma}^{-1}\hat{\gamma}_\delta}{\gamma^*(0) - {\gamma^*}_\delta^T {\Gamma^*}^{-1}{\gamma^*}_\delta} \tag{18}$$

Using Theorem 1, the sensitivity of congestion control algorithms for the ABR service to the Hurst parameter estimates can be investigated. The fractional Gaussian noise model is used to calculate the relative increase in the variance of the prediction errors when an estimate of the Hurst parameter $\hat{H}$ is used.

The autocorrelation function for fGn is given in (1), and the relative change K_v is calculated for pairs of $(H^*, \hat{H})$. The resulting surface is shown in Fig. 1. The figure reveals an asymmetry in the sensitivity of the predictors to the Hurst parameter estimates, with an asymptote situated at the point $(H^*, \hat{H}) = (1.0, 0.5)$. Thus, the relative increase in variance in the prediction errors rapidly grows when the burstiness of the data becomes more significant, ie. as $H^* \nearrow 1.0$, and yet the Hurst parameter is estimated to be close to SRD, ie. $\hat{H} \searrow 0.5$. This result emphasizes the importance of recogn ising the presence of self-similarity within network traffic if it exists. By incorporating the characteristics of self-similarity into resource management algorithms, the impact of the burstiness of self-similar traffic can be avoided.

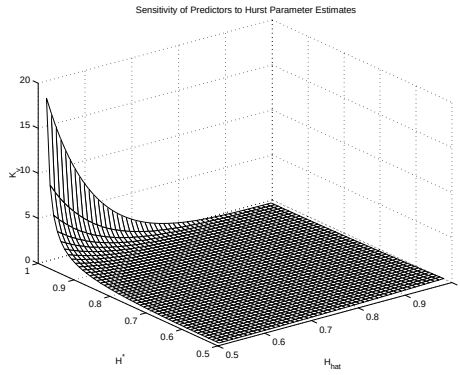

Fig. 1. The sensitivity of linear predictors to Hurst parameter estimates using the fractional Gaussian model.

Figure 1 also indicates that another conclusion can be drawn regarding the Hurst parameter sensitivity. The asymmetry at the point diagonally opposite from the asymptote, $(H^*, \hat{H}) = (0.5, 1.0)$ suggests that it is actually beneficial to err on the side of over-estimating the Hurst parameter and to assume a greater burstiness than may actually exist within the traffic. While this may be intuitively appealing, it does demand an explanation.

Predictors for the boundary values for the Hurst parameter are shown in 2. From the predictors, it is clear why the asymmetry occurs and the reason for its orientation. When $H^* = 0.99$, the optimal predictor is shown in Fig. 2(a) where significant weight is given to previous samples because of the LRD effect. However, if the predictor in Fig. 2(b) is used, no weight is given to the previous samples (even though there is a significant amount of information in these samples). This results in significant prediction errors. Consider the other situation, where $H^* = 0.5$ represents in normally distributed white Gaussian noise, and the mismatched predictor in Fig. 2(a) is used rather than the one in Fig. 2(b). In this case, there is no information about the future sample in the stored data samples and the optimal predictor gives no weight to any of the

samples. When the mismatched predictor in Fig. 2(a) is used, weight is given to these samples, but since the burstiness of the samples is mi nimal ($H^* = 0.5$), each individual sample does not impact the weighted sum as significantly and overall the prediction errors are comparitively smaller.

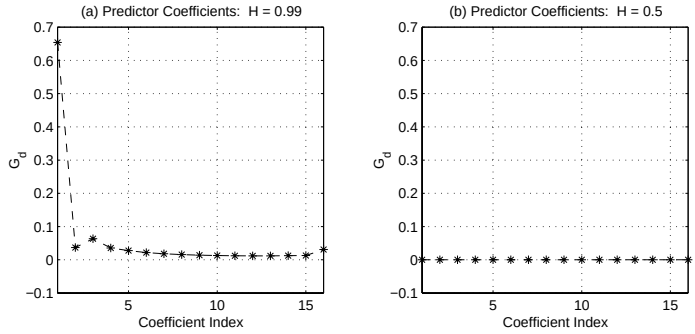

Fig. 2. Predictors for the boundary Hurst parameter values: (a) $\hat{H} = 0.99$, (b) $\hat{H} = 0.5$.

4 Simulation Results

The asymmetry of the sensitivity of ABR controllers with respect to the Hurst parameter estimates is investigated in this Section using data sets, giving simulation results which we can compare with the analytical results. The simulation model of the ABR system is shown in Fig. 3, and it consists of six ABR sources competing for the use of a congested link. Each source has its own round-trip delay, and the sources are ordered according to the delay $\delta_j = \{10, 8, 6, 4, 2, 1\}$. T he congested link is also carrying non-ABR background traffic, which uses a significant proportion of the link capacity, in this case 50% of the capacity on average.

The background traffic is modelled by data sets which consist of ten aggregated VBR data sets [17,18] which have been randomly shifted in time and have been filtered to remove GOP correlation structure of the VBR data. The sources of these aggregated data sets and the Hurst parameter estimates for the data sets are summarised in Table 1. Two additional data sets (the *Combination VBR Data* and *White Gaussian Noise* sets) were used for comparison. The *Com bination VBR Data* was produced by aggregating all the previous VBR data sets sources in the table, which had been randomly shifted in time. Finally the *White Gaussian Noise* data set was generated from a normally distributed random number generator, and transformed so that it has an equivalent mean and variance to the *Star Wars* data set.

Comparing the estimates from the different estimation techniques, it is observed that while there is sufficient agreement between the techniques that self-similarity is evident in each of the aggregated VBR data sets and that, for most

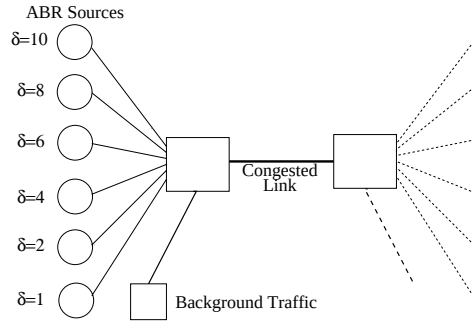

Fig. 3. The network model for the simulation study.

cases, the confidence 95% generated from the AV estimator contains all three estimates, there is also sufficient discrepancy between the values for the Hurst parameter to leave uncertainity in the choice of the value when an ABR congestion control algorithm is being develope d. To determine the effect of this uncertainity on the performance of the ABR system, predictors for the simulation model were developed across the entire interval of possible estimates, $\hat{H} \in [0.5, 1.0)$ and the system was simulated using each predictor.

Buffers at network nodes are designed to handle traffic overflows which occur in the network. In our case, these overflows represent positive prediction errors. The maximum queue length is used as measures of the performance of the network, and an indicator of the sensitivity of the control algorithms to the Hurst parameter estimate under the assumption of infinite network buffers. For

Table 1. Comparison between the Hurst parameter estimates for the variance-time (V-T), R/S statistic (R/S) and Abry-Veitch (AV) estimation methods.

Data Set Source	**V-T Estimate**	**R/S Estimate**	**AV Estimate** with 95% CI
Star Wars	0.84	0.90	0.835 [0.774,0.897]
Mr. Bean	0.76	0.86	0.817 [0.634, 1.000]
James Bond: Goldfinger	0.89	0.86	0.851 [0.669, 1.034]
Jurassic Park	0.79	0.81	0.850 [0.6667, 1.033]
Terminator II	0.77	0.83	0.838 [0.732,0.943]
ATP Tennis Final 94: Becker - Sampras	0.77	0.85	0.884 [0.778, 0.989]
Soccer World Cup Final 94: Brazil - Italy	0.57	0.79	0.791 [0.685, 0.896]
Combination VBR Data	0.71	0.85	0.811 [0.629, 0.994]
White Gaussian Noise	0.47	0.51	0.495 [0.469, 0.522]

finite buffers, the cell loss ratio (CLR) is used as the performance metric. These simulation results are shown in Figs . 4 and 5.

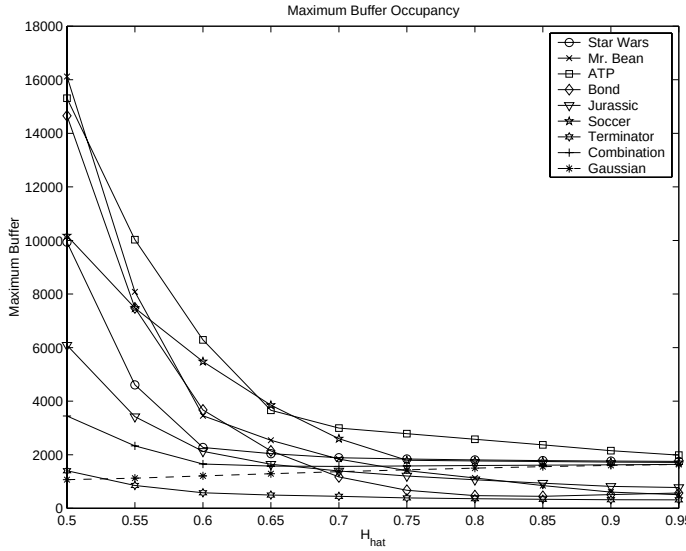

Fig. 4. The maximum queue length at the node against the Hurst parameter estimate.

It is observed from these figures that the performance of the ABR system is highly sensitive to the Hurst parameter estimate with self-similar background traffic. This occurs in the region of the interval $\hat{H} = [0.5, 1.0)$ where $\hat{H}$ is close to the 0.5 value, which means that the design has assumed short-range dependence within the network traffic. This sensitivity is revealed in longer queues in the nodes, and greater cell losses. The sensitivity of the system is significantly reduced when the estimates for the Hurst parameter are in the interval $\hat{H} \in (0.75, 1.0)$. Comparing these simulation results with the surface for predictor sensitivity derived analytically (Fig. 1), we have further confirmation that the actual values of the Hurst parameters for this type of data is in the interval $\hat{H} \in (0.75, 1.0)$.

The *Combination VBR Data* and *White Gaussian Noise* data sets provide an interesting comparison. In both cases the overall sensitivity to $\hat{H}$ is significantly less than in the other cases. In the case of the *White Gaussian Noise* data set, this reduction in sensitivity agrees with the conclusions from our analytical work, that there is insignificant change in network performance when a mismatched filter is used with SRD traffic. The slight increase in the maximum queue lengths and CL R as the Hurst parameter is increased across the interval $[0.5, 1.0)$ confirms that the true value for the Hurst parameter is close to 0.5. In the case of the *Combination VBR Data* set, while the sensitivity is not as significant as the other VBR data sets, it still appears that the assumption of self-similarity and

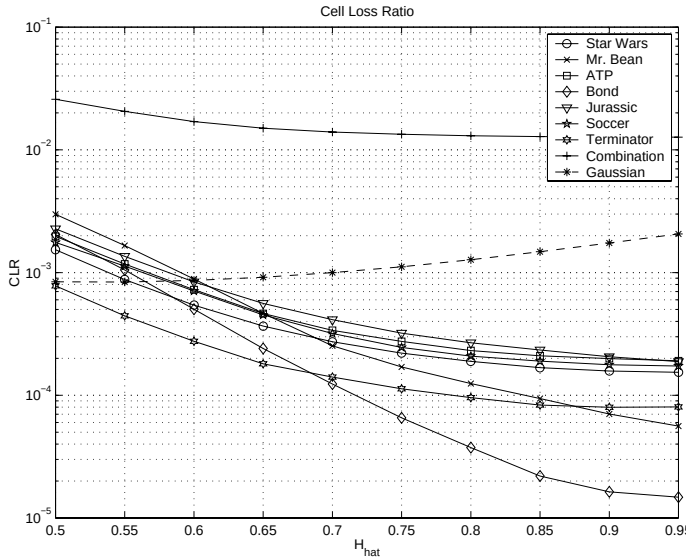

Fig. 5. The Cell Loss Ratio, for the system with a finite buffer of maximum capacity 100 cells, versus the Hurst parameter estimate.

a higher estimate for the Hurst parameter does result in marginally improved network performance.

While our simulation results do not prove self-similarity for the aggregated VBR data sets (which is not the purpose of this research), two main conclusions can be drawn from these results. Firstly, significant losses will occur within the ABR control system if the self-similarity of network traffic is not accounted for in the design of the control algorithms. These losses in the system can be reduced by incorporating the characteristics of self-similarity into the algorithms. Secondly, if two estimates $\hat{H}_1$ and $\hat{H}_2$ have been obtained for the Hurst parameter of the network traffic, where $\hat{H}_1 \leq \hat{H}_2$, then it is more judicious to choose the greater estimate $\hat{H}_2$ when designing the ABR control algorithm, provided that $\hat{H}_2 < 1.0$.

5 Conclusions

There is an inevitable uncertainity in the specific value which is assigned to the Hurst parameter when research moves from attempting to prove self-similarity to utilising its characteristics within network architecture. This issue of the impact of the Hurst parameter value has been addressed here both analytically and through simulations. The analytical study revealed an asymmetry in the sensitivity of the predictors to the estimate, especially in the region where the Hurst parameter was significantly u nderestimated and the background traffic was assumed to be SRD when in reality it was significantly LRD in character. The sensitivity was also addressed using simulations of an ABR control system

using aggregated VBR data sets to model background network traffic. Previous research has statistically proven the self-similarity of the data used, and our own Hurst parameter estimates agreed with that conclusion. The simulations investigated the performance of network buffers when a set of congestion control algorithms were used which were derived from Hurst parameter values that spanned the interval $\hat{H} \in [0.5, 1.0)$. The results of our simulations confirmed the analytical study.

This research work has confirmed previous conclusions regarding self-similarity in network traffic, specifically that there is significant impact on the network performance if the self-similar nature of traffic is ignored. Thus, it is important to continue to develop and improve Hurst parameter estimators, and to investigate methods of incorporating the characteristics of self-similarity into network management techniques. In addition, this research demonstrated that, when given a set of estimates of th e Hurst parameters, the most judicial choice of the Hurst parameter value is the highest one as this ensures the best network performance.

Acknowledgements

The authors wish to acknowledge the financial support of Telecom New Zealand (Research Grant E4567).

References

1. J. Beran, R. Sherman, M. Taqqu, and W. Willinger, "Long-range dependence in variable-bit-rate video traffic," *IEEE Transactions on Communications*, vol. 43, pp. 1566–1579, Feb/Mar/April 1995.
2. W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of ethernet traffic (extended version)," *IEEE/ACM Transactions on Networking*, vol. 2, pp. 1–15, Feb. 1994.
3. J. Beran, *Statistics for Long-memory Processes*. Chapman & Hall, New York, 1994.
4. D. Veitch and P. Abry, "A wavelet based joint estimator of the parameters of long-range dependence," *IEEE Transactions on Information Theory*, vol. 45, pp. 878–898, Apr. 1999.
5. I. Norros, "A storage model with self-similar input," *Queueing Systems*, vol. 16, pp. 387–396, 1994.
6. D. Heyman and T. Lakshmann, "What are the implications of long-range dependence for VBR-video traffic engineering?," *IEEE/ACM Transactions on Networking*, vol. 4, pp. 301–317, June 1996.
7. S. Giordano, M. Pagano, R. Pannocchia, and F. Russo, "A new call admission control scheme based on the self similar nature of multimedia traffic," *Proceedings of ICC '96*, Mar. 1996.
8. A. Adas and A. Mukherjee, "On resource management and QoS guarantees for long range dependent traffic," *Proceedings of INFOCOM '95*, pp. 779–787, 1995.
9. G. Chiruvolu and R. Sankar, "An approach towards resource management and transportation of VBR video traffic," *Proceedings of ICC '97*, pp. 550–554, 1997.
10. S. Östring, H. Sirisena, and I. Hudson, "Dual dimensional ABR control scheme using predictive filtering of self-similar traffic," *Proceedings of ICC '99*, June 1999.

11. S. Östring, H. Sirisena, and I. Hudson, "Adaptive ABR congestion control algorithm for non-stationary self-similar traffic," *Submitted to IEEE ATM Workshop '00*, June 2000.
12. G. Gripenberg and I. Norros, "On the prediction of fractional brownian motion," *Journal of Applied Probability*, vol. 33, pp. 400–410, 1996.
13. M. Roughan, D. Veitch, and P. Abry, "On-line estimation of the parameters of long-range dependence," *Proceedings of GLOBECOM '98*, Nov. 1998.
14. Y. Zhao, S. Q. Li, and S. Sigarto, "A linear dynamic model for design of stable explicit-rate ABR control scheme," *Proceedings of INFOCOM '97*, pp. 283–292, Apr. 1997.
15. K. Åström and B. Wittenmark, *Computer-Controlled Systems Theory and Design*. Prentice Hall, New Jersey, 3rd ed., 1997.
16. P. Brockwell and R. Davis, *Time Series: Theory and Methods*. Springer-Verlag, New York, 2nd ed., 1991.
17. O. Rose, "Statistical properties of MPEG video traffic and their impact on traffic modelling in ATM systems," *Proceedings of 20th Conference on Local Computer Networks 1995*, pp. 397–406, Oct. 1995.
18. M. Garrett and W. Willinger, "Analysis, modelling and generation of self-similar VBR video traffic," *Computer Communications Review Proceedings of ACM SIGCOMM*, vol. 24, pp. 269–280, Aug. 1994.