

A Performance Comparison of Monofractal and Multifractal Traffic Streams

Ram Balakrishnan Carey Williamson
Department of Computer Science
University of Saskatchewan
E-mail: carey@cs.usask.ca

Abstract

This paper focuses on the differences between monofractal and multifractal network traffic, both in terms of modeling approaches, and in terms of practical impacts on network performance. Empirical traffic traces are used in the parameterization and evaluation of traffic models. Simulation is used to evaluate the performance differences between monofractal and multifractal traffic, both for individual traffic streams and for aggregations of traffic streams. The simulation results indicate that multifractal traffic offers greater potential for multiplexing gains, both within and across sources. The sensitivity of these results to source characteristics and switch buffer size is also explored.

1. Introduction

Accurate traffic modeling is crucial for network performance studies. High speed multimedia networks exhibit complex topologies, heterogeneous transmission speeds, and a wide range of QoS requirements. Computer simulation has become an indispensable tool to support the design, control, and management of integrated services networks such as Asynchronous Transfer Mode (ATM) networks, ultra-fast optical networks, and the “next generation Internet”. Successful modeling and efficient simulation techniques can help in significantly reducing the development costs for such complex, dynamic, and non-linear systems. Also, a good model can enhance the understanding of a system or network by enabling a sensitivity study of the effect of model parameters on the performance of a network through simulation.

In the literature, several models have been suggested for describing network traffic. Early models typically assumed that packet arrivals were Poisson processes [11]. Later studies have shown that the inter-arrival times of packets differ drastically from an exponential distribution; rather, the data packet traffic constitutes a self-similar process [8], with burstiness and correlation across many time scales.

One approach for modeling *monofractal* (i.e., self-similar) traffic is proposed by Norros [10]. This approach uses three parameters – the mean bit rate m (in units of bits/sec), the Norros variance coefficient a (the variance-to-mean ratio of the traffic at the 1.0-second time scale, in units of bit-sec), and the Hurst parameter H (a unit-less value between 0.5 and 1.0, which quantifies the global scaling properties of the traffic across time scales, and hence the degree of long-range dependence) – for characterizing self-similar data traffic [10, 17]. However, one drawback of the Norros traffic model is that it does not characterize short-range correlations [1], which have been shown in the literature to have significant impact on the queueing behavior (and thus the buffer overflow process) in practical networks [6, 7, 13].

More recent network traffic analyses have shown that WAN (Wide Area Network) TCP/IP (Transmission Control Protocol/Internet Protocol) traffic is *multifractal* in nature. That is, the traffic exhibits different scaling characteristics at large time scales (e.g., hundreds of milliseconds or more) than at smaller time scales [3, 5].

To capture the multifractal characteristics of network traffic, several wavelet-based modeling techniques have been proposed [3, 4, 9, 12]. Among the proposed wavelet-based models, the multifractal wavelet model (MWM) is quite promising [1, 12]. MWM provides a unified approach to model both long-range dependence (LRD) and short-range dependence (SRD) in network traffic (e.g., hyperbolic decay of the autocorrelation function for LRD traffic [8], and exponential decay for SRD traffic).

The purpose of this paper is to compare and evaluate approaches for modeling monofractal and multifractal traffic, and to understand the performance differences, if any, between these two types of traffic. Our results confirm the observed multifractal nature of empirical WAN TCP/IP network traffic, and the effectiveness of the MWM approach for modeling such traffic. The results also show perceptible performance differences (in terms of Cell Loss Ratio, CLR) between monofractal and multifractal traffic streams, both individually and in aggregate, in a simulated network. In general, multifractal traffic offers *better* multiplexing gains

than monofractal traffic, in the network scenarios simulated.

The remainder of the paper is organized as follows. Section 2 provides some background on monofractal traffic modeling using the Norros traffic model, and on multifractal traffic modeling using MWM. Section 3 presents a simulation study designed to compare the performance of individual monofractal and multifractal traffic streams. In Section 4, a performance comparison between aggregated monofractal and aggregated multifractal traffic streams is presented. Finally, Section 5 concludes the paper.

2. Traffic modeling overview

To support the research work described in this paper, we developed two interactive traffic modeling toolkits (*synTraff* and *MsynTraff*) for monofractal and multifractal traffic, respectively [1]. The design and operation of these tools is described in a companion paper in the ‘‘Tools’’ session of MASCOTS’2000 [2]. Thus the description provided here is rather brief.

2.1. Modeling monofractal traffic: *synTraff*

synTraff models monofractal (i.e., self-similar) traffic using a three-step traffic modeling technique [17] that was developed to model the three Norros traffic parameters (m , a , and H). In the first step, a Fractional-ARIMA (Auto-Regressive Integrated Moving Average) process called *Hosking’s model* is used to generate a zero-mean LRD time series (X) with a specified Hurst value (H). The values in this time series are real numbers in the range $[-1, 1]$. Each observation in the time series depends on the previous observations (i.e., an observation has a dependency on recent and past history of the process through a weighted autoregressive moving average term, as well as a smaller dependency on a ‘‘random Gaussian noise’’ term). In the second step, the time series is converted into a ‘‘traffic profile’’. The conversion is carried out by linear ($Y = \omega X + c$) or non-linear transformations ($Z = b \frac{1-r}{r+\epsilon} X$). In the linear transformation, the mean (m) and variance coefficient (a) of the time series are adjusted using scaling (ω) and translation (c) parameters, while the Gaussian marginal distribution (i.e., frequency histogram) is preserved. In the non-linear transformation, the mean and variance of the time series are changed, as well as the marginal distribution, using base (b) and density (r) parameters (along with a small positive constant ϵ). In the third step, the short-range correlation characteristics of the traffic are controlled by refining the traffic to finer-grain time scales. This step can be used to provide a better ‘‘fit’’ to empirical traffic characteristics, without affecting m , a , or H .

2.2. Modeling multifractal traffic: *MsynTraff*

MsynTraff implements the multifractal wavelet model (MWM, as described in [12]), along with several improvements suggested in [1]. The wavelet models start with a discrete-time signal (i.e., a time series) $x(t) : 0 \leq t < N$ with N data points, where $N = 2^L$ is a power of two.

In simple terms, wavelet-based models express the traffic characteristics in the wavelet domain rather than the time domain, by integrating the traffic time series with a set of basis functions that capture information about different portions of the time series, at different scales (j) and offsets (k). The models then compute the wavelet coefficients that correspond to each node in a complete binary tree representation of the encoded signal, with $j = 0$ representing the top (root) level of the tree, and $j = L$ corresponding to the bottom (leaf) level, which has $N = 2^L$ observations.

MWM uses the Haar wavelet [12], which has the following prototype scaling (ϕ) and wavelet (ψ) functions:

$$\phi(t) = \begin{cases} 1, & \text{if } 0 \leq t < 1. \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

$$\psi(t) = \begin{cases} 1, & \text{if } 0 \leq t < 1/2. \\ -1, & \text{if } 1/2 \leq t < 1. \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

The scaling coefficients (u_j^k ’s) at the finest resolution ($j = L$) are directly determined from the data points of the signal using:

$$x(k) = 2^{L/2} u_L^k, \quad k = 0, 1, \dots, 2^L - 1. \quad (3)$$

The wavelet coefficients (w_j^k ’s) and scaling coefficients (u_j^k ’s) at coarser resolutions (smaller j values) are determined recursively using:

$$\begin{aligned} w_{j-1}^k(t) &= 2^{-1/2} (u_j^{2k} + u_j^{2k+1}), \\ u_{j-1}^k(t) &= 2^{-1/2} (u_j^{2k} - u_j^{2k+1}). \end{aligned} \quad (4)$$

The wavelet coefficients in the binary tree capture all the essential information about the original signal (time series). To recalculate all the coefficients and reconstruct the data points, only the global mean (u_0^0) of the signal and the ratios

$$a_j^k = \frac{w_j^k}{u_j^k} \quad (\text{where } u_j^k > 0) \quad (5)$$

of the wavelet coefficients to their corresponding scaling coefficients are required. From the a_j^k , the scaling coefficients can be determined by resolving Equation 5 with Equation 4:

$$\begin{aligned} u_j^{2k} &= 2^{-1/2} (1 + a_{j-1}^k) u_{j-1}^k \\ u_j^{2k+1} &= 2^{-1/2} (1 - a_{j-1}^k) u_{j-1}^k \end{aligned} \quad (6)$$

To generate a new time series $y(t)$ synthetically, the a_j^k 's are randomly selected based on their sample variances (calculated from an empirical trace) at each resolution $j < L$. To guarantee non-negative output values [12], these random numbers should be in the range $[-1,1]$. In [12], the symmetric beta distribution, $\beta(p_j, p_j)$, is recommended for selecting the a_j^k 's, with:

$$\text{var}[a_j^k] = \frac{1}{2p_j + 1}. \quad (7)$$

At the top few levels of the binary tree, there are not enough a_j^k 's for a reliable estimation of their sample variance. Hence, Ribeiro *et al.* [12] recommends generating scaling coefficients at level $j = 5$ from the root (*i.e.*, 32 samples) from a Gaussian distribution with mean and variance equal to the sample mean and variance of the scaling coefficients of the empirical trace at this scale. However, [1] recommends an F-ARIMA process for this purpose.

2.3. Discussion and evaluation

Our monofractal and multifractal traffic modeling tools have been tested and evaluated on a set of empirical traces, including LAN, WAN, TCP/IP, and ATM traffic [1].

In general, the results of our studies show that the synthetic multifractal traces generated from the MWM resemble the modeled empirical traces more closely than do the monofractal traces generated using the Norros traffic model. For example, Figure 2 of our companion “Tools” paper [2] compares the profile, autocorrelation, and marginal distribution of a monofractal trace (MON, in the left column of plots), an empirical trace (Bellcore Ethernet LAN traffic [8], in the middle column of plots), and a multifractal trace (MUL, in the right-hand column of plots), at 1.0 second granularity. The MON and MUL traces were synthesized based on the empirical trace.

We also find that the MWM approach captures (with much more fidelity) a wider range of traffic characteristics than the Norros monofractal traffic model. The variances of the a_j^k 's at each time scale (level) of the encoded binary tree reflect the wavelet behavior of a trace, which may differ for several traffic streams, even if their Norros traffic characterizations are identical [1]. Furthermore, the wavelet behaviors of aggregate-level traffic streams can differ significantly from the behaviors of host-level streams. For example, Figure 1 illustrates the wavelet behavior of the aggregate LBL-TCP-3 trace [11] and two individual host-level traces for hosts that contributed substantially to the aggregate traffic. From Figure 1, it can be observed that the wavelet behavior of the host-level traces are strikingly different from that of the aggregate-level trace. Rather than strictly increasing toward finer-grain time scales (larger j values), the variance of the normalized wavelet coefficients

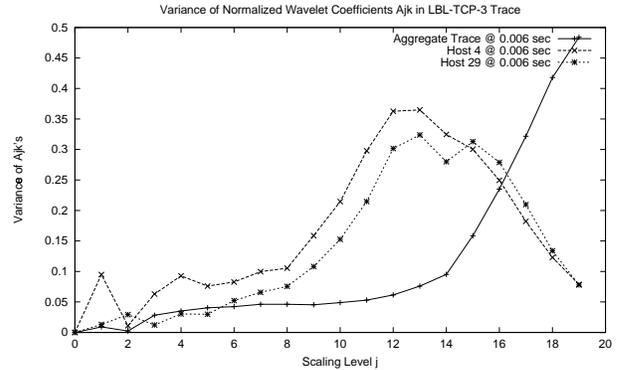


Figure 1. Comparison of wavelet behavior at host-level and aggregate-level in an empirical trace (0.006 sec granularity)

peaks and then drops. The reason for this drop is that network round-trip times, TCP sliding window protocol behaviors, and user think times translate into noticeable ON-OFF periods (particularly OFF periods) at the individual host-level [5]. This behavior is different at the aggregate-level because the OFF periods of one host are superimposed with the ON periods of (many) other hosts.

We thus find that the multifractal wavelet model (MWM) is a powerful, flexible, and efficient ($O(N)$ running time) approach for network traffic modeling. One disadvantage of the MWM approach, however, is that it requires more parameters ($2 + \log_2 N$, where N (a power of 2) is the number of data points in the modeled trace) than the Norros traffic model (3 parameters).

3. Results for individual traffic streams

This section presents a simulation experiment designed to answer the question: Does multifractal traffic behave differently than monofractal traffic, in terms of impact on network performance?

3.1. Experimental methodology

In order to compare the performances of individual monofractal (MON) and multifractal (MUL) traffic streams, five MON and five MUL traffic streams were generated (using *synTraff* and *MsynTraff*) to model an empirical trace of Bellcore Ethernet LAN traffic (BC-pOct89) [8]. The mean bit rate of the empirical trace is 2 Mbps and the source variability is 300,000 bit-sec. The mean and the Norros variance coefficient of the ten generated traffic streams are shown in Table 1. The Hurst value for each trace is approx-

Table 1. Traffic characteristics of individual monofractal and multifractal traffic streams

Stream No.	Traffic Type	m (Mbps)	a (bit-sec)
1	MON	2.037	280,333
	MUL	2.037	294,316
2	MON	1.970	297,955
	MUL	1.967	305,306
3	MON	2.021	325,647
	MUL	2.017	312,286
4	MON	2.024	350,472
	MUL	2.005	349,872
5	MON	2.020	292,958
	MUL	2.020	291,022

imately 0.7. (In [8, 15, 16], the Hurst value of the empirical traffic is reported as 0.8. Experiments in this section considered only the first 1000 seconds of the empirical data (out of 1759.62 seconds), hence the different Hurst value.)

The simulations are conducted using the *ATM-TN* (ATM Traffic and Network) simulator, a cell-level simulator described in [14]. The topology of the simulated network is shown in Figure 2. Traffic flows unidirectionally from the U of S node to the U of R node. The VarLink and Var-Buffer are the bottlenecks of the network. The link speed (in Megabits/sec) of VarLink and the buffer size (in cells) were varied and the performance in terms of CLR was studied. The link speed of VarLink determined how quickly the VarBuffer is emptied. If the buffer is full when a cell arrives, the arriving cell is lost.

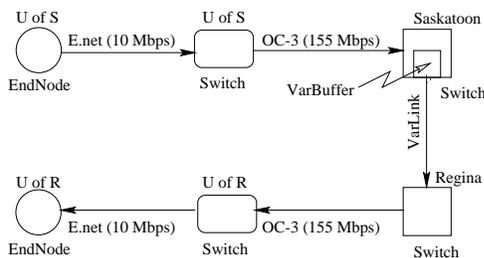


Figure 2. Simulated network topology for experiments with individual traffic streams

Other than the two types of traffic, the main factors that are considered in this study are link capacity and buffer size. Two performance metrics—cell loss ratio (CLR) and link utilization (LU)—were studied. Link utilization, expressing the number of cells delivered relative to the bottleneck link capacity, indicates the load carried over the bottleneck link.

CLR is the ratio of the number of cells lost due to buffer overflow and the number of cells sent. First, the link capacity is fixed at 4 Mbps, and the buffer size is varied from 10 to 4000 cells. Next, the buffer size is fixed at 1000 cells, and the link capacity is varied from 2.5 to 5 Mbps. In each case, the performance of each trace was evaluated individually. The simulation warm-up phase was 100 simulated seconds, followed by observations for 900 simulated seconds [1].

3.2. Simulation results

Figure 3 summarizes the performance comparison for individual MON and MUL traces. The solid lines show the mean performance for the five MUL traces and the error bars show the full range of results observed (*i.e.*, the largest and smallest values observed for the five traces). The dotted lines show the mean performance for the MON traces, and again the error bars on these lines show the range of results observed. Figure 3(a) and (b) show the link utilization of the two traffic types at a constant link speed and at a constant buffer size, respectively. Figure 3(c) and (d) illustrate the corresponding CLR performance results.

From Figure 3(a) and (c), it is clear that when the link speed is fixed and the buffer is small (fewer than 2000 cells), MUL achieves higher utilization and lower CLR. At larger buffer sizes, MUL achieves similar utilization to MON, but still yields lower CLR. When the buffer size exceeds 1000 cells, the difference in CLR between MUL and MON tends to remain constant, at about an order of magnitude.

When the buffer size is held constant, and the link speed increases from 2.5 Mbps to 5 Mbps, the higher link utilization for MUL (initially 5% to 7% higher) gradually decreases to meet that of MON at 5 Mbps. However, the CLR for MUL remained lower at all link speeds. Beyond a link speed of 3 Mbps, the difference in CLR performance between the two types of traces gradually increases, and at 5 Mbps the difference is about an order of magnitude.

3.3. Summary

The performances results for individual traces show that the MUL traces have better multiplexing gains than the MON traces. The difference in CLR can be as large as an order of magnitude, on the network scenarios studied.

Note that multiplexing gains are of two types: within source and across source. Multiplexing gains within a source result from buffering of short-term bursts. Multiplexing gains across sources come from the fact that not all the sources are “on” (or “off”) at the same time (*i.e.*, sources generate bursts independently). The performance results presented in this section account only for multiplexing gains within a source. Multiplexing gains across sources are investigated in the next section.

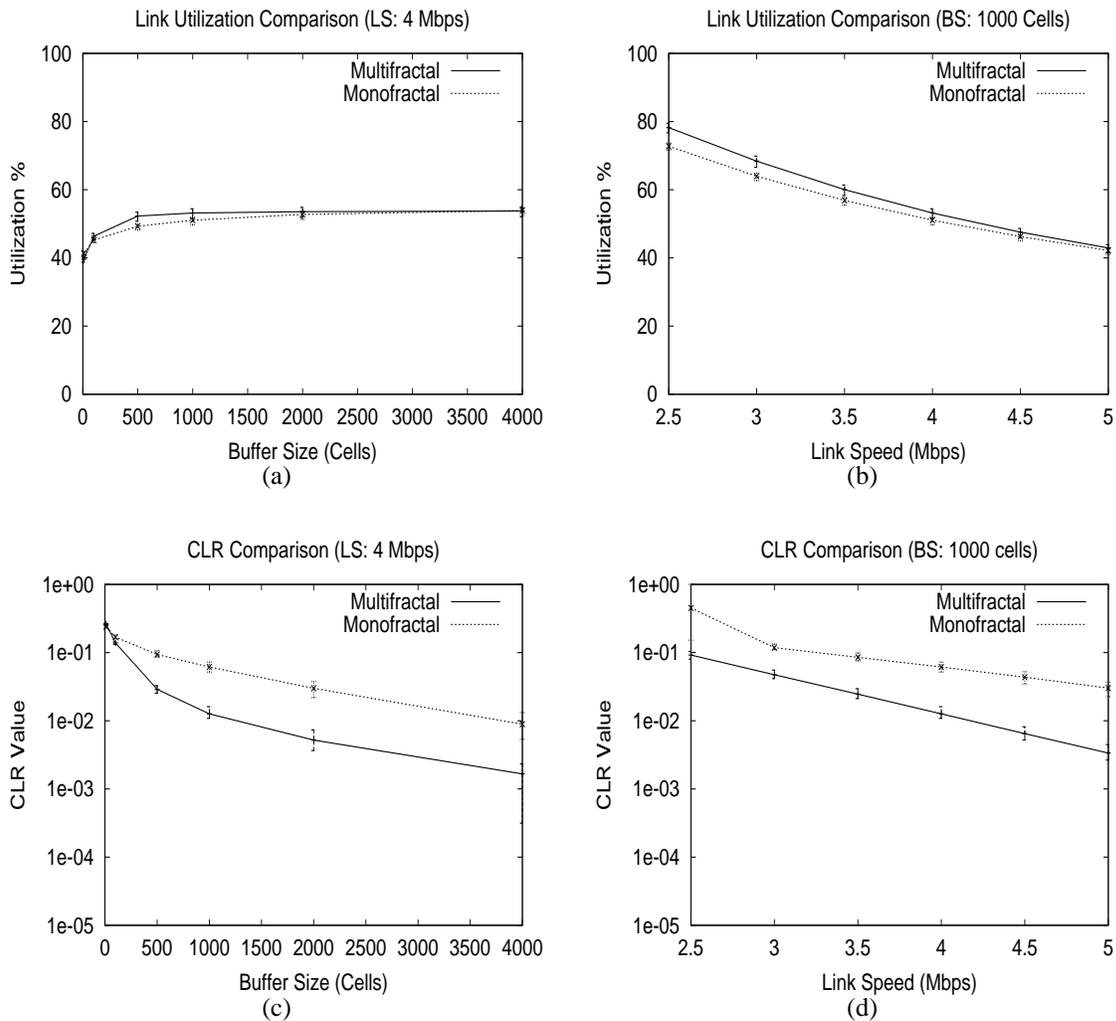


Figure 3. Simulation results for individual traffic streams

4. Results for aggregated traffic streams

This section studies the network behavior of aggregated homogeneous multifractal traffic streams. The results of these studies are presented and compared with the results obtained from similar studies of monofractal traffic by Wang and Williamson [15, 16]. The experimental design closely follows that of Wang [16], to facilitate direct comparisons of results.

4.1. Experimental methodology

Table 2 lists the different factors and levels used in the simulation experiments. The impacts of source granularity (as characterized by the mean bit rate), and source variability (as characterized by the Norros variance coefficient)

were studied, in addition to the impact of buffer size. The impact of each factor was studied for different link capacity (ranging from 10 Mbps to 300 Mbps) for the bottleneck link. Experiments were conducted using a one-factor-at-a-time approach. That is, when impact of a factor is studied by changing its level, the levels of the other factors were fixed at those highlighted in bold in Table 2.

Three performance metrics—Call Acceptance (CA), Link Utilization (LU), and Cell Loss Ratio (CLR)—were used. Call acceptance reflects the number of calls accepted by a *Connection Admission Control* (CAC) algorithm, at a given link capacity. Only the results for the Generic CAC (GCAC) algorithm are presented here; the results for other CAC algorithms are available in [1]. GCAC bases its call admission decision on a traffic descriptor that specifies the Peak Cell

Table 2. Experimental factors and levels

Factors	Levels		
Link Capacity (C)	10 Mbps, 50 Mbps, 100 Mbps, 150 Mbps, . . . , 300 Mbps		
Source Granularity (m)	1 Mbps,	2 Mbps,	3 Mbps
Source Variability (a)	150,000 bit-sec,	300,000 bit-sec,	600,000 bit-sec
Buffer Size (b)	1000 cells,	2000 cells,	4000 cells

Rate (PCR) and Sustained Cell Rate (SCR) of the source, and on switch-level information about available capacity and a target variance factor (VF) [16].

4.2. Previous work

In [15, 16], Wang and Williamson studied the performance of CAC algorithms, when presented with self-similar (monofractal) traffic, using a simulated network. In that study, five CAC algorithms—Peak Cell Rate (PCR) CAC, Sustained Cell Rate (SCR) CAC, Average (AVG: average of PCR and SCR) CAC, Generic CAC (GCAC), and Norros CAC—were considered. Source granularity, source variability, and long-range correlation structure, as characterized by the three Norros traffic parameters, were used to examine the impact of traffic characteristics on CAC performance. The link capacity was varied as an experimental factor. The target CLR was 10^{-6} .

The monofractal traffic results from [15], corresponding to the levels listed in Table 2, are used in this paper for comparing with the results obtained for multifractal traffic, on the same simulated network and under similar settings. What is new in these experiments is the multifractal nature of the generated traffic streams.

4.3. Traffic streams

Six groups of homogeneous multifractal traffic streams, each with 100 independent sources, were generated using *MsynTraff*. Table 3 summarizes the characteristics of each group. Two of the six groups (*m2* and *a2*) correspond to the baseline configuration for the BC-pOct89 empirical trace (see Section 3.1). Two groups (*m1* and *m3*) are for different source granularities, and the other two (*a1* and *a3*) are for different source variabilities. The traffic groups were carefully generated so that within each group, the multifractal (MUL) traces (on average) differ by no more than 3% from their corresponding monofractal (MON) traces, in terms of their Norros traffic characteristics (*m*, *a*, and *H*).

4.4. Network topology and simulation issues

Figure 4 shows the topology of the simulated network used in the experiments. The network has 100 sources con-

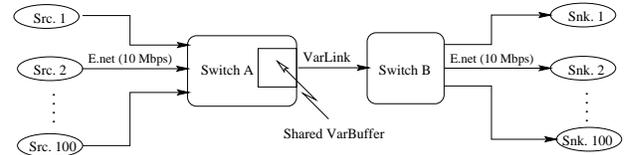


Figure 4. Simulated network topology for experiments with aggregate traffic streams

nected to Switch A through independent 10 Mbps Ethernet links and 100 sinks are similarly connected to Switch B. Switch A is connected to Switch B through a variable bandwidth (bottleneck) link. Each source generated traffic bursts destined to the corresponding sink (*e.g.*, Source 1 transmitted data to Sink 1, Source 2 to Sink 2, and so on). The traffic bursts were segmented into cells before transmitting.

The CAC algorithms were implemented in Switch A and Switch B. These two switches were provided with calculated traffic descriptors for each call. Based on traffic descriptions and available bandwidth, the CAC algorithms made decisions about which calls to accept. Once the call admission decisions are made, Switch A receives cells from all accepted sources and forwards them to Switch B through the bottleneck link, buffering as necessary. The finite capacity output buffer is shared by all traffic streams, and operates in a first-in-first-out fashion, using a tail drop policy. The capacity of bottleneck link determines how quickly the buffer in Switch A is emptied. If a cell arrives when the shared buffer is full, the arriving cell is lost.

As in the previous study [15], the warm-up phase of the experiment was 100 simulated seconds. The simulations ran for the next 900 simulated seconds. In these experiments, CLR values larger than the target CLR value (10^{-6}) can be considered significant [15], and a CLR value of zero should be interpreted as being less than 10^{-6} .

4.5. Simulation results

Figure 5 shows the call acceptance (CA, Figure 5(a)), link utilization (LU, Figure 5(b)), and cell loss ratio (CLR, Figure 5(c)) results for the three levels of source granularity

Table 3. Characteristics of monofractal and multifractal traffic streams

Traffic Characteristic	Source Group	Traffic Type	m (Mbps)			a (bit-sec)		
			Min.	Mean	Max.	Min.	Mean	Max
Source Granularity	m1	MON*	0.698	1.060	1.398	188,482	309,805	445,814
		MUL	0.854	1.084	1.367	263,579	305,667	404,501
	m2	MON*	1.592	2.026	2.519	154,466	309,615	430,495
		MUL	1.729	2.028	2.273	208,270	314,345	409,348
	m3	MON*	2.635	3.081	3.550	169,334	304,778	432,576
		MUL	2.695	3.012	3.261	246,274	304,346	432,487
Source Variability	a1	MON*	1.817	2.088	2.519	76,498	152,179	232,588
		MUL	1.951	2.037	2.087	120,709	154,893	203,653
	a2	MON*	1.673	1.997	2.368	182,211	301,781	486,152
		MUL	1.729	2.028	2.273	208,270	314,345	409,348
	a3	MON*	1.404	2.075	2.722	364,576	599,899	817,761
		MUL	1.668	2.093	2.427	494,701	599,833	828,866

* obtained from [15]

(1 Mbps, 2 Mbps, and 3 Mbps) in the MUL traces. Figure 6 compares the CLR performance of MON (dotted line) and MUL (solid line) traces at the three source granularities. Irrespective of traffic type (MON or MUL) the CA is almost the same [1], since the traffic descriptors are similar for the MON and MUL sources (see Table 1, for example). The CA depends primarily upon the source granularity and link capacity. From Figure 5 it is obvious that the number of calls accepted decreases as the source granularity increases. Both MON and MUL traffic achieved similar link utilization. However, the CLR values for multifractal traffic are consistently lower than those for monofractal traffic.

The impact of source variability is illustrated in Figure 7. The CLR results for MON and MUL traces at the three different source variabilities are compared in Figure 8. At 150,000 bit-sec source variability, two more calls are accepted for MUL than for MON; at 300,000 bit-sec, one more call is accepted for MUL; and at 600,000 bit-sec the CA for MUL and MON are the same at most of the tested link capacities [1]. Consequently, the link utilization for MUL is marginally higher than that for MON at 150,000 bit-sec and 300,000 bit-sec. However, the CLR for MUL is noticeably lower than that for MON at 150,000 bit-sec and 600,000 bit-sec, and slightly lower at 300,000 bit-sec. It can be concluded that at lower source variabilities, MUL has better multiplexing gains than MON, and as the source variability increases the multiplexing gains for MUL and MON are similar.

As illustrated in Figure 9, changing the switch buffer size for MUL traffic affects only the CLR performance, but not the CA and the LU (this is true for MON traffic as well, since the GCAC algorithm ignores buffer size in its call acceptance decision-making [16]). Regardless of the buffer size, the number of calls accepted for MON and MUL were

almost the same. Figure 10 compares the CLR performance of MON and MUL at buffer sizes of 1000, 2000, and 4000 cells (Figures 10(a), (b), and (c), respectively). It is obvious from the figure that increasing the buffer size reduces the CLR for both MON and MUL. However, the reduction in CLR for MUL is more pronounced than for MON. In other words, increasing the buffer size produces a more significant reduction in CLR for MUL than for MON.

4.6. Summary

The performance results for aggregated homogeneous monofractal (MON) traffic streams and aggregated homogeneous multifractal (MUL) traffic streams show that MUL has better across source multiplexing gains, especially when the source granularity is high, source variability is low, or switch buffer size is large.

5. Conclusions

This paper confirms previous work [3, 4, 5, 12] showing that empirical network traffic is multifractal in nature, and shows that the multifractal wavelet model (MWM) [12] can capture the behavior of these traces better than the Norros (monofractal) traffic model. The wavelet behaviors of aggregate-level streams differ from the behaviors of host-level streams, and MWM can capture these differences.

The paper also uses simulation to show that the network performance impacts of monofractal and multifractal traffic streams are different. In particular, the within source multiplexing gains for multifractal traffic (MUL) are considerably larger compared to those for monofractal traffic (MON). MUL has better across source multiplexing gains compared to MON at lower source variabilities, higher

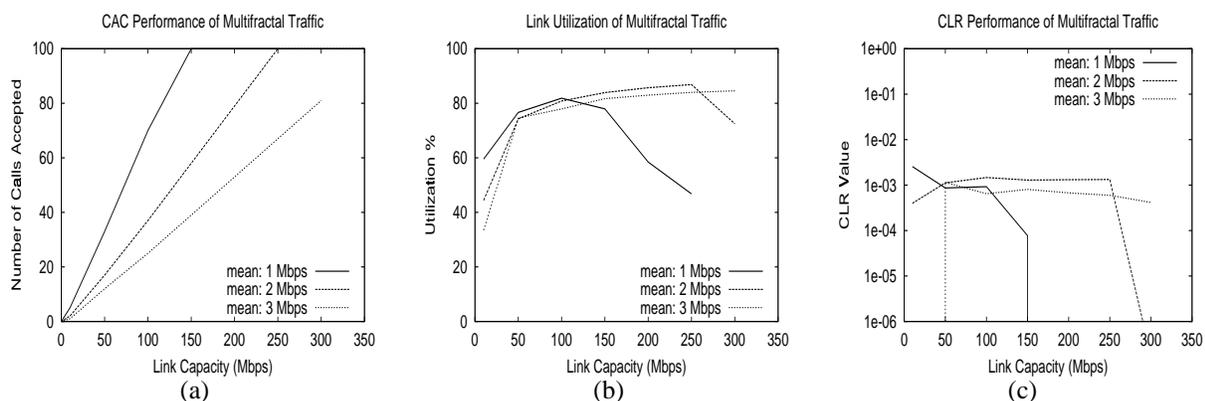


Figure 5. Network performance of multifractal traces for the three different source granularities: (a) Call acceptance; (b) Link utilization; and (c) CLR performance.

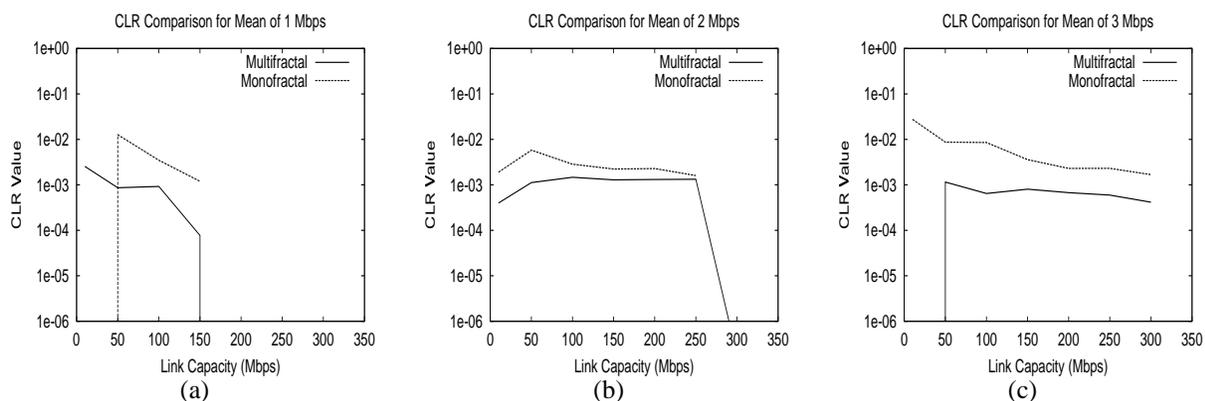


Figure 6. Comparison of CLR for monofractal and multifractal traces for different source granularities: (a) 1 Mbps; (b) 2 Mbps; and (c) 3 Mbps.

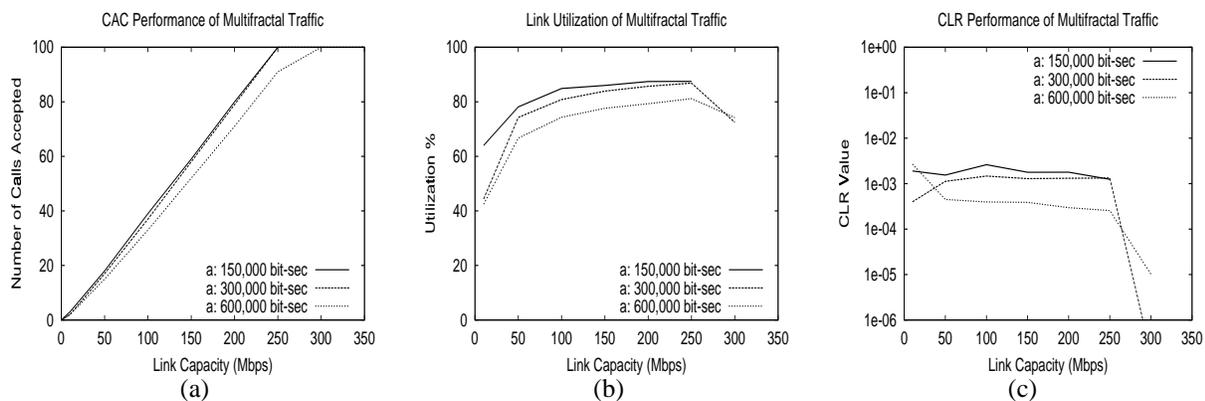


Figure 7. Network performance of multifractal traces for the three different source variabilities: (a) Call acceptance; (b) Link utilization; and (c) CLR performance.

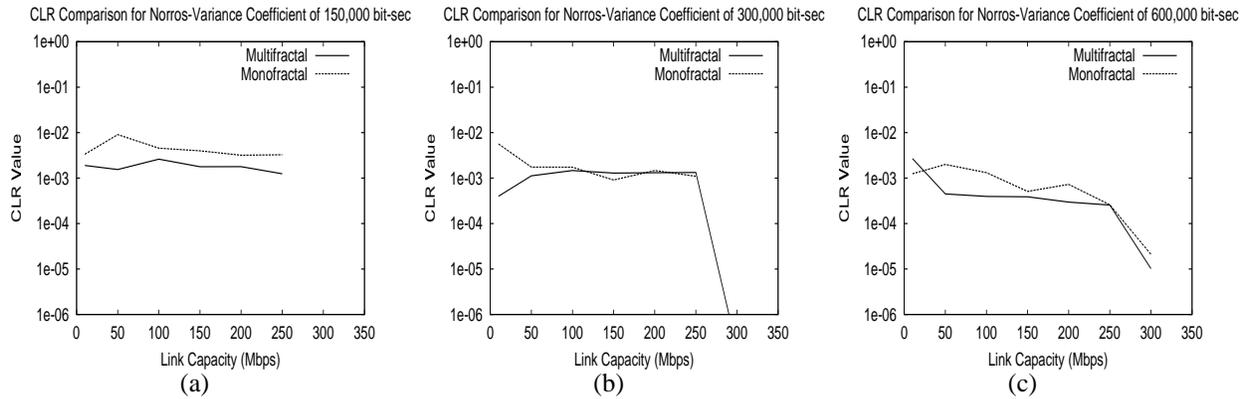


Figure 8. Comparison of CLR for monofractal and multifractal traces for different source variabilities: (a) 150,000 bit-sec; (b) 300,000 bit-sec; and (c) 600,000 bit-sec.

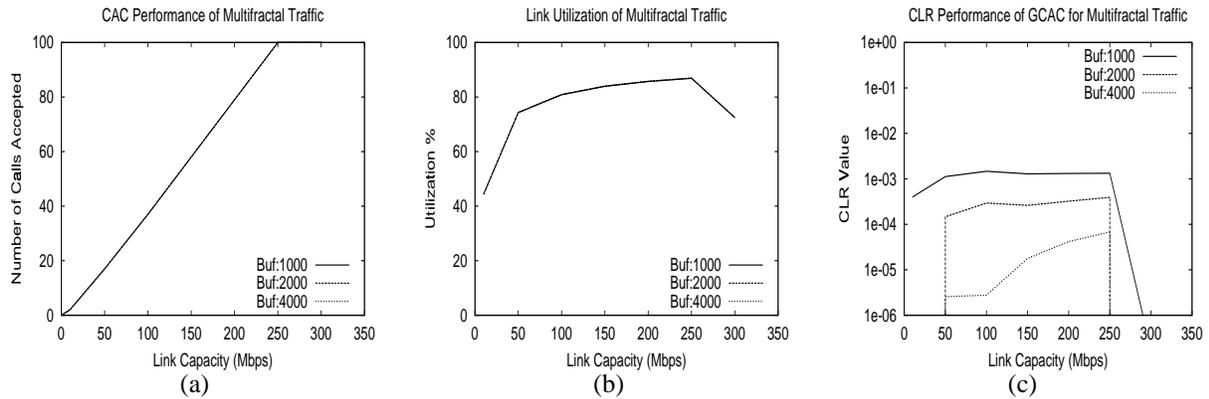


Figure 9. Network performance of multifractal traces for the three different switch buffer sizes: (a) Call acceptance; (b) Link utilization; and (c) CLR performance.

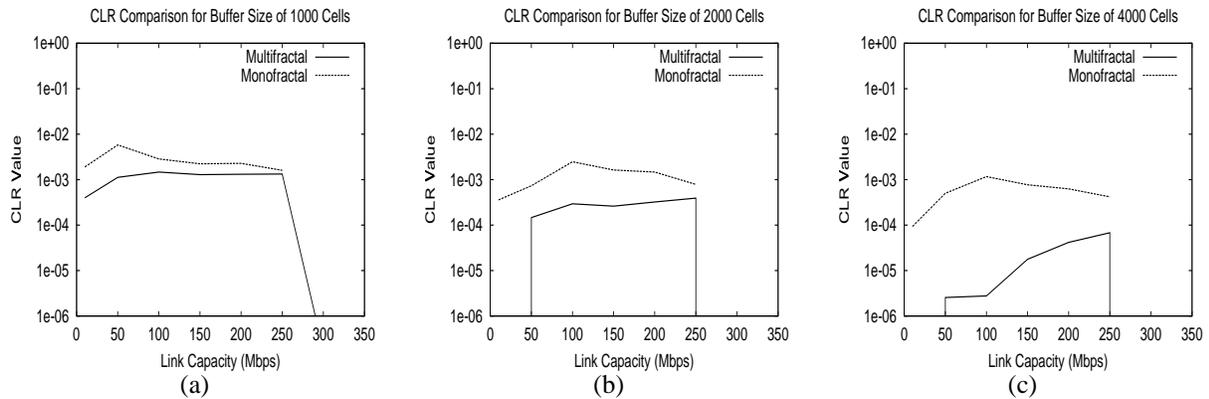


Figure 10. Comparison of CLR for monofractal and multifractal traces for different buffer sizes: (a) 1000 cells; (b) 2000 cells; and (c) 4000 cells.

source granularities, or when the switch buffer size is large. Otherwise, across source multiplexing gains for MUL are similar to those for MON.

Acknowledgements

Financial support for this research was provided by an NSERC Collaborative Research and Development (CRD) grant to the ATM-TN TeleSim project (CRD183839), and by NSERC Research Grant OGP0121969. The authors thank the TeleSim project team for building and maintaining the ATM-TN simulator (on which much of our simulation work depends), and for their constructive feedback on our traffic modeling work. We also value the ongoing interactions with our industrial sponsors, including Newbridge Networks (now Alcatel CID), Nortel Networks, Siemens, and Telus. Comments from the anonymous MASCOTS referees were helpful in improving the clarity of the final paper.

References

- [1] R. Balakrishnan, *Wavelet-Based Network Traffic Modeling*, M.Sc. Thesis, Department of Computer Science, University of Saskatchewan, March 2000.
- [2] R. Balakrishnan and C. Williamson, “The synTraff Suite of Traffic Modeling Toolkits”, Tools Session Paper, *Proceedings of MASCOTS’2000*, San Francisco, CA, August 2000.
- [3] A. Feldmann, A. Gilbert, W. Willinger, and T. Kurtz, “The Changing Nature of Network Traffic: Scaling Phenomena”, *ACM Computer Communication Review*, Vol. 28, No. 2, pp. 5–26, April 1998.
- [4] A. Feldmann, A. Gilbert, and W. Willinger, “Data Networks as Cascades: Investigating the Multifractal Nature of Internet WAN Traffic”, *Proceedings of the 1998 ACM SIGCOMM Conference*, Vancouver, Canada, pp. 42–55, August 1998.
- [5] A. Feldmann, A. Gilbert, P. Huang, and W. Willinger, “Dynamics of IP traffic: A Study of the Role of Variability and the Impact of Control”, *Proceedings of the 1999 ACM SIGCOMM Conference*, Cambridge, MA, pp. 301–313, August 1999.
- [6] M. Grossglauser and J-C. Bolot, “On the Relevance of Long-Range Dependence in Network Traffic”, *Proceedings of the 1996 ACM SIGCOMM Conference*, Stanford, CA, pp. 15–24, August 1996.
- [7] D. Heyman and T. Lakshman, “What are the Implications of Long-Range Dependence for VBR-Video Traffic Engineering?”, *IEEE/ACM Transactions on Networking*, Vol. 4, No. 3, pp. 301–317, June 1996.
- [8] 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, No. 1, pp. 1–15, February 1994.
- [9] S. Ma and C. Ji, “Modeling Video Traffic in the Wavelet Domain”, *Proceedings of 17th Annual IEEE Conference on Computer Communications, INFOCOM*, San Francisco, CA, pp. 201–208, March 1998.
- [10] I. Norros, “On the Use of Fractional Brownian Motion in the Theory of Connectionless Networks”, *IEEE Journal on Selected Areas in Communications*, Vol. 13, No. 6, pp. 953–962, August 1995.
- [11] V. Paxson and S. Floyd, “Wide Area Traffic: The Failure of Poisson Modelling”, *Proceedings of the 1994 ACM SIGCOMM Conference*, London, UK, pp. 257–268, August 1994.
- [12] V. Ribeiro, R. Riedi, M. Crouse and R. Baraniuk, “Simulation of Non-Gaussian Long-Range-Dependent Traffic Using Wavelets”, *Proceedings of the 1999 ACM SIGMETRICS Conference*, Atlanta, GA, pp. 1–12, May 1999.
- [13] B. Ryu and A. Elwalid, “The Importance of Long-Range Dependence of VBR Video Traffic in ATM Traffic Engineering: Myths and Realities”, *Proceedings of the 1996 ACM SIGCOMM Conference*, Stanford, CA, pp. 3–14, August 1996.
- [14] B. Unger, F. Gomes, X. Zhong, P. Gburzynski, T. Ono-Tesfaye, S. Ramaswamy, C. Williamson, and A. Convington, “A High Fidelity ATM Traffic and Network Simulator”, *Proceedings of the 1995 Winter Simulation Conference*, Arlington, VA, pp. 996–1003, December 1995.
- [15] Y. Wang, C. Williamson, and J. Doerksen, “CAC Performance with Self-Similar Traffic: Simulation Study and Performance Results”, *Proceedings of MASCOTS’99*, College Park, MD, pp. 102–111, October 1999.
- [16] Y. Wang, *CAC Performance with Self-Similar Traffic*, M.Sc. Thesis, Department of Computer Science, University of Saskatchewan, August 1999.
- [17] C. Williamson, “Synthetic Traffic Generation Techniques for ATM Network Simulations”, *Simulation Journal*, Vol. 72, No. 5, pp. 305–312, May 1999.