Simulating internet traffic with Markov-modulated processes

Richard G. Clegg (richard@richardclegg.org)
Dept of Elec Eng,
University College London,
London.

June 21, 2007

Abstract

This paper considers various Markov-modulated processes which have been used in the literature to model internet traffic. In particular, such processes have been used to replicate the long-range dependent nature of internet traffic and the justification for doing so is that long-range dependence (LRD) has important effects on delays and buffer occupancy in traffic models.

However, in combination with an extremely simple queuing model, the traffic generation models have their parameters tuned to match real network traffic traces. The queuing performance is then measured both in terms of the probability of the buffer exceeding a given size and in terms of mean queue length (or mean delay)

1 Introduction

Markov chains (MC) and Markov-modulated processes (MMP) are well-known modelling techniques which are successful in a wide variety of fields. They are also a traditional tool for queuing theory and for investigating networks and queues of networks. In the last ten years several models have been introduced for the purposes of modelling internet traffic based on MMP. These models are often motivated by the idea of capturing the long-range dependence (LRD) which is seen in real internet traffic and replicating the Hurst parameter H which characterises long-range dependence. The models have a common form, they produce a process which is one or zero (*on* or *off*) and work in discrete time. The *on/off* process can then be seen as a time series of packets and inter-packet gaps. Obviously such a simplistic model of network traffic cannot be expected to capture all the behaviour of a real traffic network. However, in exchange for this loss of realism, the analytic simplicity of the models means that mathematical insights into their performance can be gained which might not be seen in simpler models. Additionally, the computational performance of such simple models is typically good allowing for fast modelling of large systems or large numbers of packets in small systems.

There is a considerable body of work on MMP to analyse the queuing of internet traffic ([15], produces some useful results for analysing queuing of MMP type models). This paper concentrates on those MMP models using *on/off* processes with the intention of capturing LRD. Four different MMP based models which were introduced with the aim of modelling LRD are described in this paper. Certainly many other models are possible. In the case of one model (the pseudo-self-similar traffic model) some new results (both theoretical and computational) are presented which question whether this model can produce traffic with a known Hurst parameter.

The models are then compared in simulation with 100,000 packets from two real traffic traces. The traces are both easily available from the internet for research. The data sets chosen represent an older but extremely well-studied packet trace (from 1989) and a more modern data set likely to be more representative of modern internet traffic (from 2003).

Previous work by Huebner et al [11] attempts to use a different set of models to replicate the performance of internet traffic traces. The authors tested a Poisson model, a Weibull model, a two state MMP, an autoregressive (AR(1)) model, a Pareto model and a Fractional Brownian Motion model (the latter captures LRD). The authors concluded that none of the models tested produced a good match for queuing performance in all circumstances but that the model with LRD was the nearest when a large buffer size was considerd.

An unpublished expanded version of this paper can be found in [6].

1.1 Long-Range dependence in internet traffic

The introduction to LRD given here is, by necessity, brief. For a fuller introduction see Beran [4] and for an introduction in the context of internet traffic, see Clegg [5, Chapter one]. For a summary of work on LRD in internet traffic see Willinger et al[22].

Let X_t be a weakly-stationary time series $\{X_t : t \in \mathbb{N}\}$ with mean μ and variance σ^2. The autocorrelation function (ACF) as a function of lag k is given by

$$\rho(k) = \frac{\mathrm{E}\left[(X_t - \mu)(X_{t+k} - \mu)\right]}{\sigma^2}.$$

Definition 1 *A weakly-stationary time series is said to be* long-range dependent *(LRD) if the sum of its ACF is not convergent. That is, the sum $\sum_{k=0}^{\infty} \rho(k)$ diverges. Note that sometimes the weaker condition that $\sum_{k=0}^{\infty} |\rho(k)|$ diverges is given.*

Often a specific asymptotic form for the ACF is assumed,

$$\rho(k) \sim c_\rho k^{-\alpha}, \tag{1}$$

where $k, c_\rho > 0$ are constants, $\alpha \in (0, 1)$, and $\sim$ here and throughout this paper means asymptotically equal to as $k \to \infty$. This form is used to define the Hurst parameter which is given by $H = 1 - \alpha/2$. Note that not all LRD processes necessarily have a definable Hurst parameter but the Hurst parameter where $H \in (1/2, 1)$ is usually considered the standard measure of LRD. For a discussion of measuring the Hurst parameter in the context of telecommunications traffic see [7]. Long-Range dependence is also sometimes expressed in terms of asymptotic second order self-similarity although the concepts are subtly different.

A related topic is that of heavy-tailed distributions.

Definition 2 *A random variable X is heavy-tailed if, for all $\varepsilon > 0$ it satisfies*

$$\mathbb{P}\left[X > x\right] e^{\varepsilon x} \to \infty \quad \text{as} \quad x \to \infty. \tag{2}$$

Again, often a specific form is assumed

$$\mathbb{P}\left[X > x\right] \sim C x^{-\beta}, \tag{3}$$

where $C > 0$ is a constant and $\beta > 1$.

From Heath et al [10, Theorem 4.3], heavy tails and long-range dependence are related. An *on*/*off* process with heavy-tailed *on* periods of the form given in (3) and *off* periods which fall off faster is a long-range dependent process. Note that if a process has heavy-tailed *off* periods and *on* periods which fall off faster then this, then theorem can still be applied since the ACF of an *on*/*off* process is unchanged if *on* and *off* are reversed.

The area of long-range dependence became of interest to internet researchers when LRD was discovered in measurements of packets per unit time on an Ethernet segment [13]. It was later shown that the data sources for the traffic exhibited heavy tails in their on periods [23]. These heavy tails are speculated to be the cause of the LRD in internet traffic. These measurements have been repeated many times since. The reason this is important is that LRD can have severe implications for queuing performance. Traffic exhibiting LRD can have much longer delays although the relationship is not a simple one [16, 20].

2 Traffic generation models considered

Four models have been found in the literature of which two use the same topology but different parameters. In order to simplify the explanations given in this paper, the models will be described using the same notation even where this will differ from the notation given by the authors in the papers cited. While the Wang model was not actually suggested as a model for internet traffic, it is included here as it is the oldest model the author has found in the literature which associates on/off MMP with LRD.

Each of the models which attempt to replicate the mean μ traffic level for the real traffic. This can be thought of as the proportion of the time that the network is occupied with transmitting traffic or the occupancy. In the case of the MMP simulation models, it is the proportion of time that the model spends in on states. The value of μ is critically important since it controls the amount of traffic the model will produce. The models will also, where appropriate, attempt to replicate the Hurst parameter.

2.1 Wang model

The Wang model [21] grew out of the problems associated with calculating the invariant density of certain non-linear maps. In particular, it arises from a piecewise linearisation of the Manneville–Pomeau map [17] which is itself used in internet traffic modelling [9]. The topology of the MC is given in Figure 1. The parameters f_i are transition probabilities with f_i being the probability that the zero state (off) is then followed by a run of exactly i on periods and f_0 the probability that another off period follows.

The Wang model [21] uses two parameters to generate the infinite family of parameters f_i. One of the parameters corresponds to the Hurst parameter and the other corresponds to the mean. While the Hurst parameter can be explicitly set, no closed form solution is known for how to obtain a given mean from the model but an iterative procedure can be used to work out the parameter setting for a given mean.

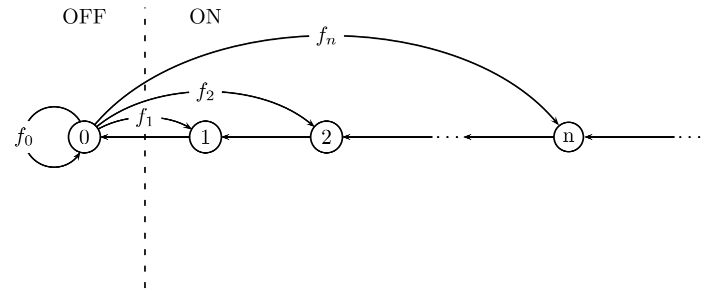

Figure 1: A topology used by models which generate LRD for particular choices of f_i.

2.2 The Pseudo-Self-Similar Traffic (PSST) model

The PSST model [19] was introduced to capture the LRD in packet traffic (they use the phrase *self-similarity*). In fact the model suggested is a finite model which would not generate self-similarity but the authors hope it would approximate it. The model is further investigated in [12] and criticised as providing unrealistic estimates for queuing performance. The topology of the PSST model is shown in Figure 2.

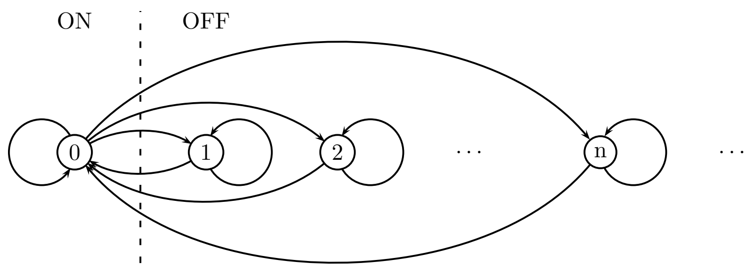

Figure 2: The topology of the PSST model.

Previous references have used a truncated finite version of this model, this provably cannot generate long-range dependence (hence the name pseudo-self-similar). However, there seems no particular reason to use this approximation and here the infinite model will be used. The model has two parameters, one of which sets the mean of the traffic produced and the other in some way controls the correlations. However, there is no direct correlation between the model parameters and the Hurst parameter. The long term behaviour of the model is interesting. For more details on see [6, Section 2.6].
It should be also noted that the choice of on and off in this model is controversial. A common assumption is that long-range dependence arises from correlations in the on periods (for example arising from the heavy-tailed distribution of file lengths in transfers). However in this model, the length of on periods decays exponentially and the length of the off periods has slower decaying correlations. In this paper, therefore, the PSST(b) model is also considered which is the same model with on and off reversed.

2.3 Arrowsmith–Barenco model

This model [2, 3] was introduced to capture the LRD seen in packet traffic and as a development of the Wang model. The topology is shown in Figure 3.
Because the model is a double-sided version of the topology of the Wang model then the decay of consecutive runs of zeros and ones can be individually controlled. An important result with this topology is [3, Theorem 4] which shows how if the decay of the f_i^L and f_i^R parameters are known then the Hurst parameter of the resulting traffic stream can be calculated.
It should be noted that this is really a whole family of models and the authors consider various methods for setting the f_i^L and f_i^R parameters. In this paper this topology is used to model the distribution of on and off times in the real data. This will be described in Section 3.3.

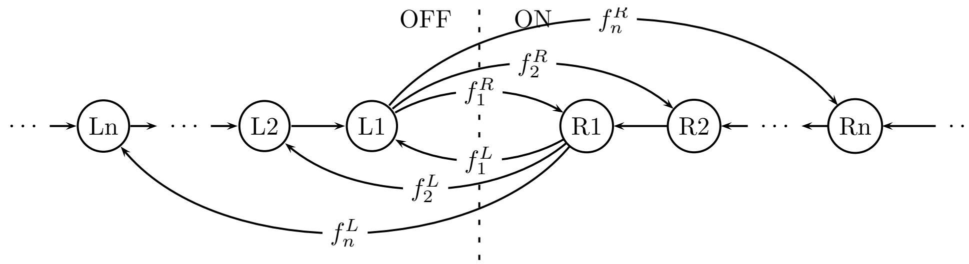

Figure 3: The topology of the Arrowsmith–Barenco model.

2.4 Clegg–Dodson model

The Clegg–Dodson model [8] uses the same topology as the Wang model but different transition probabilities. The model has two parameters π_0 which is related to the mean by $\pi_0 = 1 - \mu$ and $\alpha \in (0, 1)$ which determines the Hurst parameter.

2.5 Computer implementation

These MMP described are relatively simple and fast to implement. For this paper, they have been implemented in the computer language python. By necessity a robust computational implementation of such routines must deal with very small probabilities. The full details of how the necessity for high precision arthmetic can be avoided are given in [8].

3 Experimental setup

The experiments performed in this paper are extremely simple. The input data for one simulation is a set of arrival times and packet lengths. These data may originate from real measurements or from the models described. The packets are then simulated as arriving at a queue of known output bandwidth. The properties of the queue and the output traffic are then measured. The experiment may then be repeated with a smaller output bandwidth to see how this affects the queue. Obviously as bandwidth decreases it would be expected that the mean queue length and queuing delays would increase but the exact behaviour depends upon the statistical nature of the traffic.

It should be noted that LRD is a difficult subject to work with computationally. Hurst parameters near unity (α near zero) cause great problems. Theoretically, the nearer one the Hurst parameter, the slower the sample mean will converge to the mean. To give a concrete example, 10^6 iterations of the Clegg–Dodson model with $\mu = 0.5$ and $\alpha = 0.1$ ($H = 0.95$) gives for the first three samples, $\overline{Y} = 0.563$, $\overline{Y} = 0.401$ and $\overline{Y} = 0.426$. This does not indicate a problem with the computational model rather this slow convergence of the sample mean is inherent in the nature of LRD itself. Naturally, this has a critical importance to real experiments.

3.1 Data sets used

Two data sets are used for the simulation in this paper. In both cases, only the first 100,000 packets were investigated. The names and origin of the exact sources used are given here so that other researchers can make similar measurements.

CAIDA data: This data set is taken from a trace approximately an hour long. It is referred to on the CAIDA website as `20030424-000000-0-anon.pcap.gz` and was captured on the 24th April 2003. It was captured on an OC48 link with a rate of 2.45 Gb/s. The average packet length was 493 bits. The data is freely available to researchers who fill in a request form. More information about this data can be found at:
`www.caida.org/data/passive/`.

Bellcore data: This much-studied data set is described in [14]. The data here is taken from an August 1989 measurement referred to as `BC-pAug89.TL`. The data was collected on an Ethernet link which connected a LAN to the outside world. Note that in this case the data did not record the true length of packets, only the length less the Ethernet header (which is variable). The average packet length recorded was 464 bits. Hence, the experiment in this paper is only using an approximation of the real data. The data is freely available for researchers. More information about this data can be found at:
`ita.ee.lbl.gov/html/contrib/BC.html`.

3.2 Queuing model, pre-processing and digitisation of real data

The queuing model used in this paper is extremely simple. The system has a given bandwidth b (bits/sec). Items join the back of the queue. If a packet of length L bits arrives at the queue then it will take a time L/b to process. It is output from the queue at this time. Until this time the entire packet is considered to be part of the queue for purposes of calculating mean queue length (which can be calculated in terms of packets or bits).

A starting point for modelling is to establish a base case for comparisons. The real data simply consists of arrival times of packets and packet lengths. In order to attempt to match this data with real models, then a bandwidth b was selected. This was chosen to create an occupancy near ten percent as this was thought to be a reasonable occupancy for a congested network. The Bellcore data was reported as being taken from a network with an occupancy of twenty percent at peak times. The CAIDA data almost certainly had an occupancy much lower than this since it is from a modern high-speed link. The actual figure chosen is not really important since the data are then to be queued through lower and lower bandwidths.

For the Bellcore data the baseline bandwidth was chosen as 1.96Mb/sec and for the CAIDA data 128Mb/sec this gave occupancies of and 0.094 and 0.098 respectively. Traffic with the recorded arrival times and packet lengths was then passed through this queue and the output times from this queue were taken as the base case to simulate. The data referred to as "raw" for the rest of this paper is the output of this queue with either the Bellcore or CAIDA packet lengths as an input.

The traffic generation models are all *digitised* in a way that the real data was not. The models all simply produce a string of ones and zeros corresponding to a packet or a gap. To convert these into packets and departure times a timescale dt must be established and also a fixed packet length. The timescale dt is the length of time between packets in a packet train or the length of one inter-packet gap. The packet length l bits was chosen as the mean packet length of the real data (464 bits for the Bellcore data and 496 bits for the CAIDA). The timescale was then chosen related to the bandwidth as the time taken to transmit one packet of this length, that is $dt = l/b$.

Obviously this is a considerable simplification and it is therefore useful to investigate to what extent the real data would be altered if it were subject to this *digitisation*. Therefore a *digitised* version of the real data was produced where all the packets are of length l and broadcast at fixed multiples of dt. This data is referred to as the digitised data. It was created simply by simulating a queue where the packets arrival times and packet lengths were taken from the real data. At every time $ndt : n \in \mathbb{N}$, if the queue contained l or more bits then a packet of length l was sent from the queue at this time. The data referred to as *digitised* throughout the rest of this paper is the results of this process with the input as the raw data and queued using the same bandwidth b as the raw data.

Figure 4 shows the differences introduced by this digitisation process. These results are produced by queueing the Bellcore raw and digitised data in a queue with half the original bandwidth (b is reduced from 1.96Mb/s to 0.98Mb/s). The top figure shows the distribution of the queue size in bits and the bottom figure shows the distribution of the queue in packets. As can be seen on the data in bits, the real data has a much more complex graph, simply because packets can have a variety of different lengths. Interestingly the raw data tends to have a higher queue length in terms of bits but lower in terms of packets. The reason for this is not known. The digitised data certainly shows differences to the raw data but the queuing performance is not greatly dissimilar. Further comparisons will be shown in the next section.

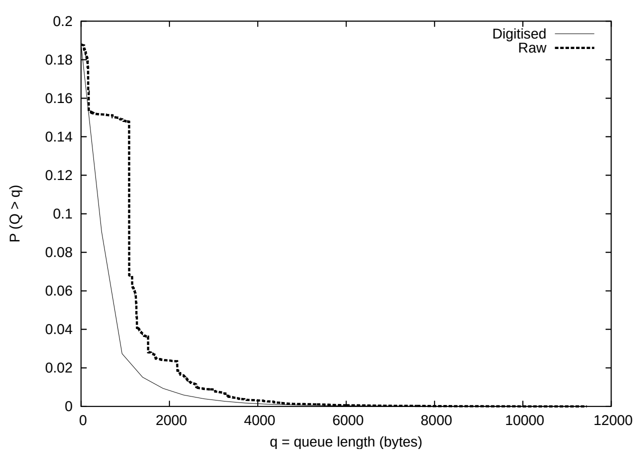

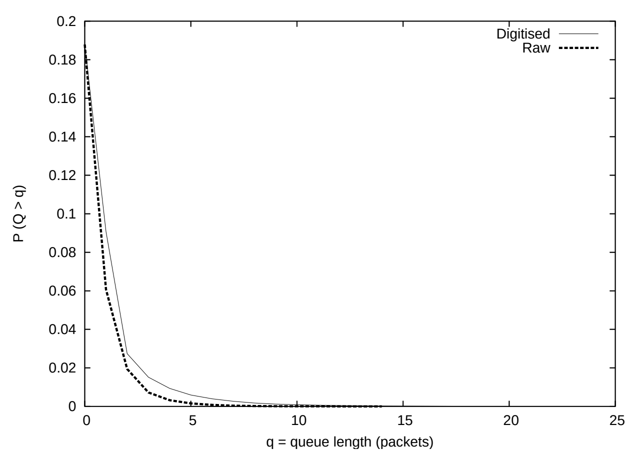

Figure 4: Plot of $\mathbb{P}[Q > q]$ with q as the queue length in bits(top) and packets(bottom).

3.3 Models used to simulate data

The models compared were the Poisson model and the Fractional Gaussian Noise (FGN) model as well as the MMP models, Wang, Clegg–Dodson, Arrowsmith–Barenco, PSST and PSST(b). This section describes how the models were tuned to be as close as possible to the real data.

The Poisson model has a single parameter and captures the mean of the real data.

The FGN model has two parameters and captures the mean and Hurst parameter of the real data. Note that a general FGN model also has a variance parameter but in an on/off model the variance is constrained by the mean and this is no longer a free parameter. The Wang and Clegg–Dodson models both have two parameters and also capture the mean and Hurst parameter of the real data.

As has been mentioned, the Arrowsmith–Barenco model is really a family of models with this topology. In this case, the specific model chosen captured exactly the distribution of on and off periods in the digitised version of the real data. This is simply done. Let $\mathbb{P}[\text{Off} = n]$ be the probability that a randomly chosen off period is of length exactly n then simply set $f_n^L = \mathbb{P}[\text{Off} = n]$. This means that the model will have the same distribution of lengths of off periods as the digitised version of the real data. Similarly, the f_n^R are set to replicate the distribution of the on times. It should be noted that this choice of parameters does not produce traffic which replicates the Hurst parameter of the real data.

The PSST model proved a particular problem. The model can be set to give data with a particular mean and the remaining parameter sets the amount of longer term correlation in the model. However, when the PSST model was set to give the required mean, for all choices of the second parameter the model produced a sample mean with an extremely high variance. That is, for several runs of the model, the number of packets produced varied greatly. The queuing performance of the model varied wildly between individual runs with the same parameters. For this reason no results using the PSST model are given. (The PSST model does produce results with a low variance on the sample mean for higher values of the mean, $\mu > 0.5$). The PSST(b) model produced more stable results for the mean required in these experiments. For both models there were problems calculating the Hurst parameter as described in the next section.

3.4 Calculating the Hurst parameter

Calculating the Hurst parameter of real data is not a simple matter. For a practical guide in the context of telecommunications and descriptions of the methods used in this paper see [7]. The methods used in this paper are the R/S estimator, Aggregated Variance, Periodogram, Wavelet analysis and the Local Whittle Estimator. Software using the statistics package R can be downloaded from:

`http://www.richardclegg.org/lrdsources/software/`. An excellent description of various methods and S-Plus code can be found at:

`http://math.bu.edu/INDIVIDUAL/murad/home.html`. The full results of the Hurst parameter tests done for this paper can be seen in [6, Section 3.4].

To calculate the Hurst parameter the data must first be converted into a time series. This can be done simply by counting the number of bytes processed if the data is split into sample times of a given period. The question is then which time period to choose. Too long a time period will give too small a sample size to work with (a ballpark figure is that several thousand time series points is a minimum). Too short a time period will give a time series which largely consists of zeros.

For the Bellcore date time periods of 0.1 seconds, 0.01 seconds and 0.001 seconds were tried. This resulted in 2521, 25208 and 252080 samples respectively. However, at the smallest sample size more than three quarters of the periods sampled have no packets at all. At the lowest sampling rate the estimators did not agree well but sampling at the two higher rates the estimators seem to agree on a Hurst parameter of around 0.8 (the exception being the Local Whittle estimator which gave 0.736 at the highest sampling rate for the raw data). The digitisation made very little difference to the estimated Hurst parameter of the data in almost all cases.

The same procedure was performed for the CAIDA data and the results broadly agreed with a value of $H = 0.6$. This is a low level of LRD (indeed it may be there is no LRD in this data) consistent with the rule of thumb that networks with lower utilisation often exhibit a lower level of long-range dependence.

Hurst parameter estimates were also made for simulated data sets of 100,000 packets as set by those models tested here which were supposed to produce traffic with LRD. The Clegg–Dodson and FGN models produced traffic where the measured Hurst parameter was close to their theoretical Hurst parameter for all the estimators used. The Wang model produced traffic which was consistent with the theoretical Hurst parameter for most estimators and when the sample size was increased to 1,000,000 packets the estimates improved.

Note that the method for choosing parameters for the Arrowsmith–Barenco model as described above does not attempt to reproduce the Hurst parameter and this was seen in these measurements. While the Bellcore data had a measured Hurst parameter around $H = 0.8$ the Arrowsmith–Barenco model tuned to have the same distribution of packet train lengths and inter-packet gaps has a measured Hurst parameter around $H = 0.6$.

The PSST(b) model performed in a spectacularly inconsistent manner. The H parameter is usually expected to be in the range $H = (1/2, 1)$ for LRD and $H = 1/2$ for no LRD. The estimates for the model varied across the entire available range and outside it. The estimates also change completely depending on the aggregation level considered. If the sample trace is longer (1,000,000 packets not 100,000) then more consistent estimates of the Hurst parameter are obtained for all models apart from the PSST and PSST(b) models.
Note that [12] reports that the PSST model was tuned to replicate the value of H for real data. It is possible that the authors only had a single H estimator available and thus did not notice these discrepancies. In this paper the remaining parameter for the PSST(b) model was chosen simply to be "large" for the high Hurst case and smaller for the low Hurst case. No more scientific fitting procedure was available for the model.

4 Results

The experiments performed are all of the same nature. The input to an experiment is data either from a real data source (raw or digitised) or from one of the models with its parameters tuned to match that of the digitised data. The input data is then sent through a queue with a given bandwidth b. The queuing performance of the model is then measured. While many performance measures could be considered, results are only given here in terms of the expected queue size $\mathrm{E}[q]$ and overflow probability $\mathbb{P}[q > n]$ for selected n. The experiment is then repeated with a smaller value of b until experiments have been performed with occupancies ranging from 0.1 to 0.6 (the latter representing a network with an extremely high degree of congestion).

4.1 Bellcore data

All models were run to produce traces 252 seconds long (the length of the original trace) with packets of length 464 bits.

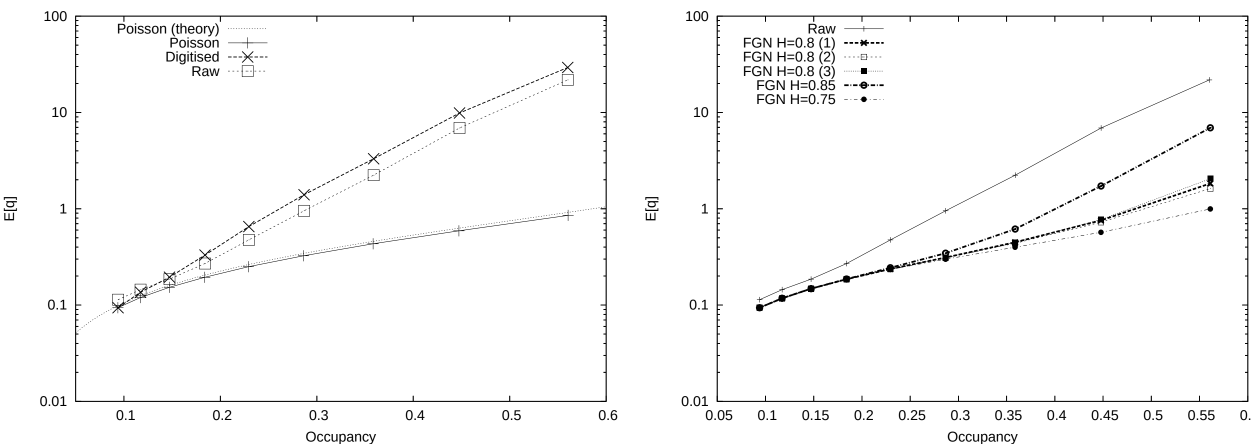

Figure 5: Comparison of Poisson (left) and FGN model (right) versus real traffic for Bellcore data.

Figure 5 (left) shows comparisons of the Poisson model with the real data (both raw and digitised). Note that the y axis is a logscale on this and all following figures. The theory line for the Poisson model is provided by the Pollaczek–Khinchin formula (with the discrepancy being accounted for by the fact that the model is, strictly speaking, Bernoulli not Poisson). As can be seen, and as would be expected, the Poisson model hugely underestimates the level of queuing in the network. This result is as would be expected from the literature.
Figure 5 (right) shows a comparison of the FGN model versus real traffic. Several realisations of FGN are tried. Three realisations have $H = 0.8$, one has $H = 0.85$ and one has $H = 0.75$. As can be seen the model produces more or less the same queuing performance with the same mean and Hurst parameter and a higher Hurst parameter produces larger queues. This is as would be expected from the literature. However, both are under predicting the queuing of the real data.
Figure 6 (left) shows a comparison with all of the models used against the real traffic trace. What is most striking about this is that none of the models are even close to replicating the real data. The raw and digitised data are relatively close together. The Clegg–Dodson and Wang models appear to be similar in performance (perhaps unsurprisingly since they have the same topology but different parameters). Both of these models overestimate queuing. The PSST model produces a higher queue level than these two models and is, obviously an overestimate of the queuing of the real data. The Poisson model, as has been mentioned, is an underestimate of the real queuing performance as is the FGN model. Interestingly, for low occupancies, the Poisson model is actually giving higher queues than the FGN model even though this model was motivated by addressing the

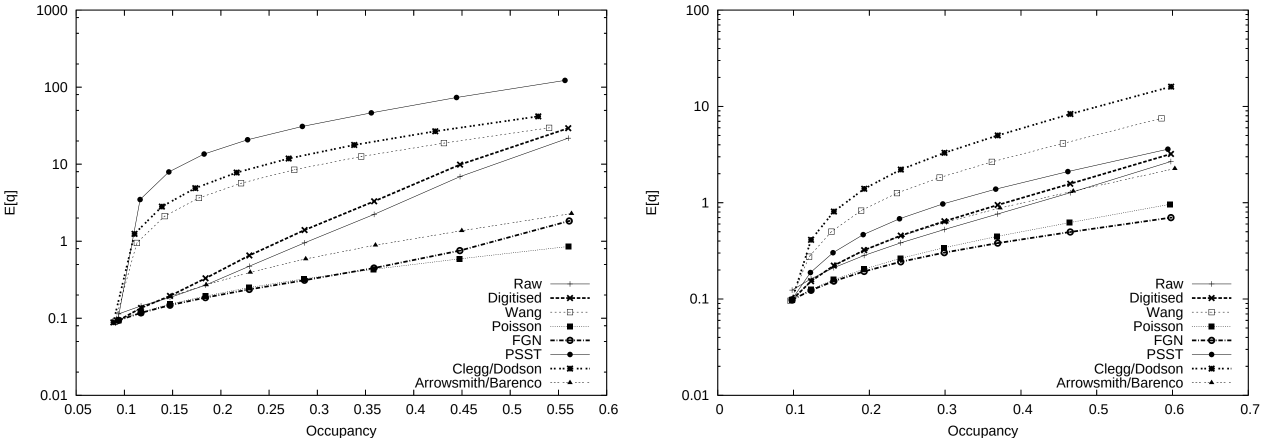

Figure 6: Comparison of queue lengths in all models on Bellcore (left) and CAIDA data (right).

underestimation of queuing in the Poisson model. The Arrowsmith–Barenco model is, perhaps, the closest model to the real data but this particular method must still be regarded as having failed to successfully model the queuing performance of the Bellcore data. Also the model used is a multi-parameter model as opposed to a one or two parameter model like the others and hence would be expected to be a much closer match.

A subtle but important difference in the figure is that in the regions with higher occupancy (the right hand side of the graph) the slope of the lines is very different. With the exception of the FGN model, in this region the models appear to have parallel lines on this figure but these lines have a very different gradient to the plots for the raw and digitised data. In other words, not only the level of congestion is different but the way the data responds to an increase in congestion is fundamentally different.

Another way of comparing the models is to look at the probability of given queue lengths. Here, a similar graph to those in Figure 6 is plotted with the y axis as the probability of the queue being equal to or greater than a given length. Figure 7 shows the probability of the queue equalling or exceeding five (left) or twenty (right). As can be seen, again none of the models are doing a good job of approximating these probabilities. The raw and digitised data remain similar to each other. At low occupancies the Poisson, FGN and Arrowsmith–Barenco models seem to be the best approximations and at high occupancies the Clegg–Dodson and Wang models seem to be closer. None of the models prove accurate over the whole range. The models prove poor approximations over the whole range of occupancies considered, no matter what the queue length chosen.

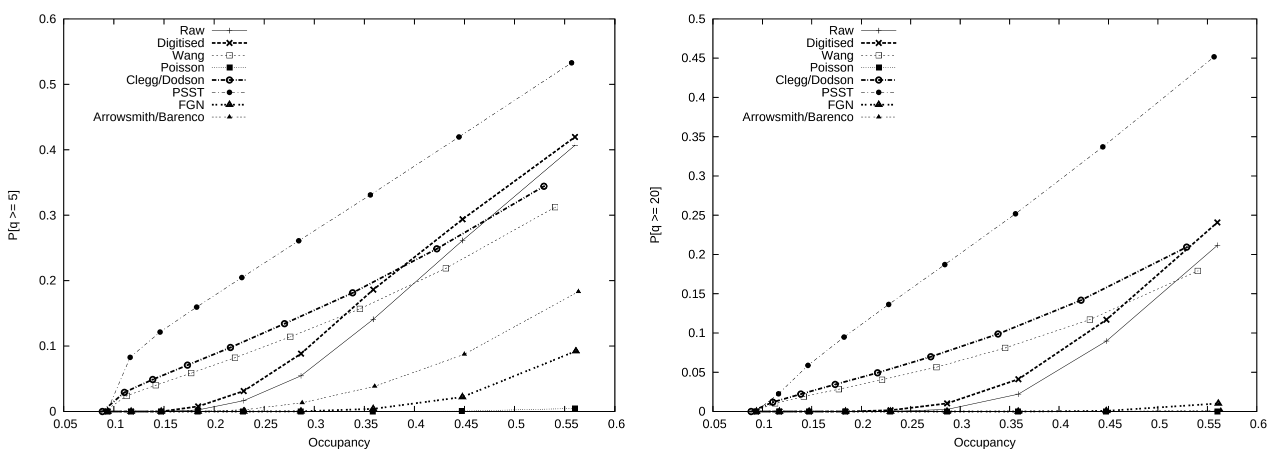

Figure 7: Comparison of all models on Bellcore data looking at the probability that the queue is five or greater (left) or twenty or greater (right).

4.2 CAIDA data

All models were run to produce traces 4.02 seconds long (the length of the original trace) with packets of length 496 bits.

Figure 6 (right) shows a comparison of all the models versus the real and digitised data. In most ways the results are similar to the results of modelling the Bellcore data. Again the Clegg–Dodson and Wang models are similar but provide an overestimate of the level of queuing. Again the FGN and Poisson models provide an underestimate of the queuing in this case with the FGN model giving a lower estimate of queuing than the Poisson model. In this case, however, the Arrowsmith–Barenco model has provided a very good estimate of the queuing lying somewhere between the raw and digitised data. The PSST model has provided a better approximation although it is still an over estimate. Again, however, the same feature can be seen as with the Bellcore data, in the high occupancy region (at the right hand side of the graph) the artificial models (with the exception of the FGN) seem to have run parallel (they appear to have approximately the same gradient). However, the real data appears to have a steeper gradient than any of the models in this region.

Again, the experiments can be repeated looking at the probabilities of the queue exceeding a given length. As for the Bellcore data, the probability of the queue length exceeding certain levels are plotted. Figure 8 shows the probability that the queue equals or exceeds five (left) or twenty (right). In the case of the queue exceeding five, the Arrowsmith–Barenco model gives an excellent approximation for most of the range of the experiment. Indeed it lies between the raw and digitised lines. The Wang and Clegg–Dodson (and possibly even PSST) models could also be seen to be acceptably close. For the more extreme event assessed by the probability that the queue exceeds twenty, the Wang and Clegg–Dodson models are clearly overstating the probability of these extreme queues. The PSST(b) model is a relatively good approximation of the real data. All other models seem to be underpredicting the likelihood of large queues at high occupancies.

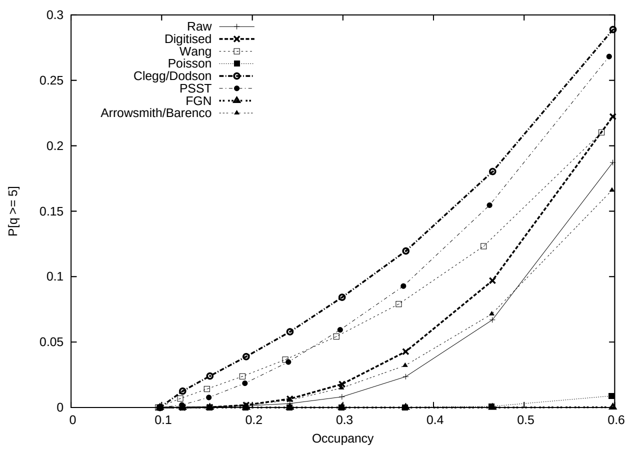

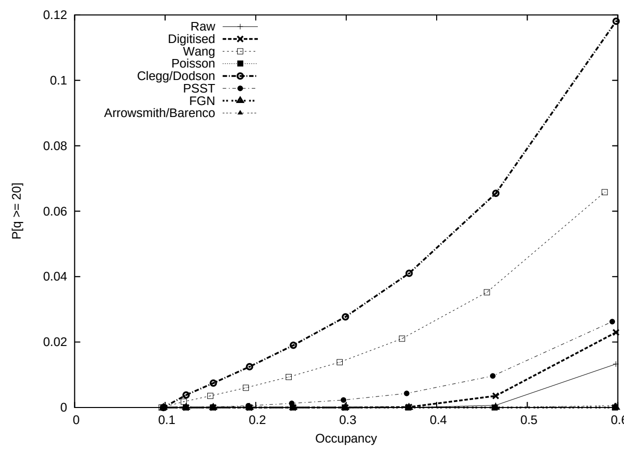

Figure 8: Comparison of all models on CAIDA data looking at the probability that the queue is five or greater (left) or twenty or greater (right).

4.3 Later sections of data

As has been seen, it is a difficult task to replicate the queuing performance of a sample of 100,000 packets of real data. It might then be asked, if this sample could, in principle be replicated with one hundred percent accuracy then would this modelling be appropriate for subsequent data samples from this data. Figure 9 (left) shows the second, third, fourth and fifth samples of 100,000 packets of the Bellcore data queued by the same process (raw data). As can be seen the queuing performance of the data varies greatly between these samples. Note that each point plotted corresponds to a different bandwidth for queuing but the data differs also in the mean packets transmitted and hence the occupancy differs between samples. However, the differences are far from being differences purely due to the time over which the packets transmitted. Consider, for example the first 100,000 packets compared to the third 100,000. The third 100,000 packets have a lower occupancy (that means they took longer to transmit and are sampled from a period of time where, on average, packets were being sent at a lower rate) but a higher queue.

Figure 9 (right) shows a similar plot for the CAIDA data. In this case, while the mean rate of data transmission still differs between samples, the queuing performance is broadly similar.

4.4 Discussion and criticism of results

The results presented here show an important weakness in a class of MMP models which have been used to emulate network traffic. Over all the data sets used, no models gave a good representation of the expected queue or queue overflow probabilities for the Bellcore data. It should be noted again that the model referred to as the Arrowsmith–Barenco model was just one

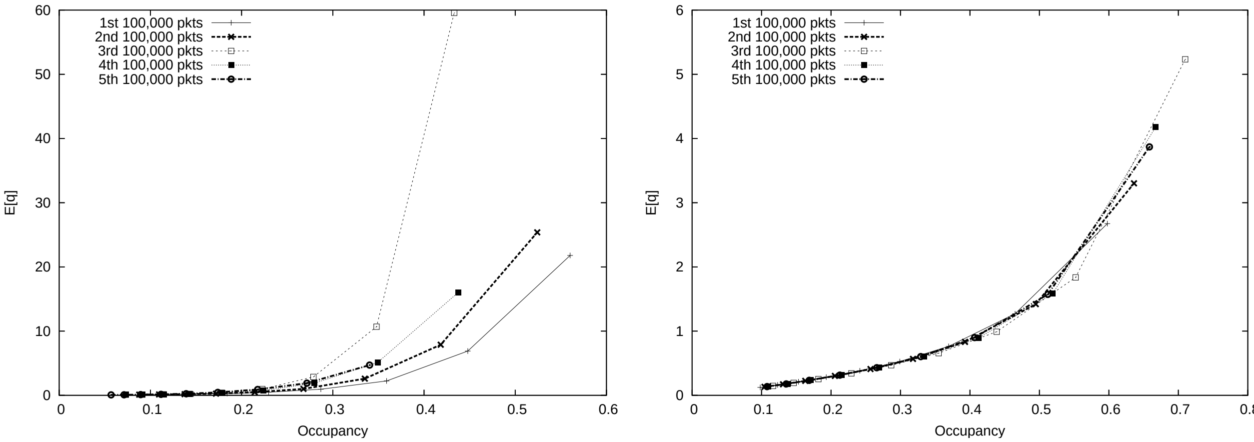

Figure 9: Comparison of subsequent sections of Bellcore data (left) and CAIDA data (right).

method for choosing parameters for this model and this is not a general criticism of using this topology for modelling data. For the expected queue length, the Arrowsmith–Barenco model gave a good representation of the CAIDA data and the PSST(b) model gave a fair representation. In the case of the PSST(b) model there is no good way the author knows of picking one of the model's parameters and it may be that this match is little more than coincidence. No models seemed able to predict the queue overflow probabilities of the CAIDA data for all queue lengths and all occupancies (although perhaps this would be demanding too much of such simple models).

The CAIDA data had a lower degree of long-range dependence. Also the CAIDA model had a more consistent performance between subsequent samples of 100,000 packets. It may be suspected that these two facts are related but it is hard to tell without further investigation.

One obvious criticism of the experiments performed here is that a real network would not behave like this under queuing. The TCP/IP protocol incorporates mechanisms which perform crude congestion control. In short, what is described here as the real data is not, in fact, how a real network would perform subject to the capacity constraints. This criticism is an important one and it is certainly true that a closed loop model incorporating this feedback would be a better representation of what actually happens when a real network becomes more congested. However, this said, good open loop models would greatly help the understanding of what factors in real network traffic impact on queuing performance. It would be much harder to understand how bandwidth impacted a closed loop system and it could be reasonably expected that changes to a network which were positive for an open loop system were also positive for a closed loop system (although this is by no means guaranteed).

Other models may be capable of producing a better model of the queuing performance of internet traffic. In [2] a method is described for tuning the parameters of the Arrowsmith–Barenco model to replicate the ACF of a traffic sample using genetic algorithms. Wavelets have been used not just for analysis of traffic traces (as in this paper) but also for simulating traffic [18] [1]. It is not clear, however, how an individual packet model could be generated from a time series produced using wavelets.

5 Conclusion

This paper presented a number of MMP models which produce *on/off* sequences. These models have all been suggested as potential models of telecommunications data (with the exception of the Wang model which arose in the field of Statistical Mechanics). In addition the FGN model was included since this is also a commonly suggested model for telecommunications data and the Poisson model was included as a baseline for comparisons. It is clear that in all cases the Poisson modelling was inappropriate and gave misleading queuing estimates.

While the Wang, Clegg–Dodson and FGN models captured the mean and Hurst parameter of the data, they did not accurately reflect the queuing performance of the system. Investigation of those models made it clear that, within those models, the Hurst parameter had a very important effect on queuing performance with a high Hurst parameter equating to worse queuing performance.

The PSST model proved hard to work with and the criticisms of this model by other authors [12] seem justified. However, this author is sceptical of the claim in [12] that the model can be fitted to provide a particular Hurst parameter. The PSST model does not seem to provide a traffic trace with a controllable Hurst parameter. In addition the original PSST model exhibits a remarkable degree of instability in its sample mean when the mean is set to produce a low level of traffic. This said, the PSST(b) model was the only two parameter model to have any degree of success in modelling the queue length either data set.

The Arrowsmith–Barenco model as used in this paper did not model the Bellcore data well but was a very good model for the CAIDA data. It should be recalled that in this paper, a particular method for fitting the model parameters was used and the Arrowsmith–Barenco model is more general than the particular model used here. The model used in this paper is a multi-parameter model and would be expected to provide a better fit than one or two parameter models.
It is important to note that in the case of the high Hurst parameter Bellcore data, even if a model were found which accurately represented the queuing performance of the first 100,000 packets subsequent samples of 100,000 packets behaved very differently. On the other hand, for the CAIDA data which had a lower Hurst parameter, subsequent samples of the data performed more consistently.
In short, the problem of replicating the statistics of real traffic traces is a difficult one. The models tried in this paper all have their attractions in terms of computational or mathematical simplicity but none of them proved adequate to model the queuing performance of a traffic trace taken from a real network. If researchers are to truly understand what causes and can relieve queuing in telecommunications networks then an open loop model of this type would be an important starting point.

References

[1] P. Abry, P. Flandrin, M. S. Taqqu, and D. Veitch. Self-similarity and long-range dependence through the wavelet lens. In P. Doukhan, G. Oppenheim, and M. S. Taqqu, editors, *Theory and Applications of Long-Range Dependence*, pages 526–556. Birkhäuser, 2003.

[2] M. Barenco. *Packet Traffic in Computer Networks*. PhD thesis, School of Mathematical Sciences, Queen Mary, University of London, 2002.

[3] M. Barenco and D.K. Arrowsmith. The autocorrelation of double intermittency maps and the simulation of computer packet traffic. *Dynamical Systems*, 19(1):61–74, 2004.

[4] J. Beran. *Statistics For Long-Memory Processes*. Chapman and Hall, 1994.

[5] R. G. Clegg. *Statistics of Dynamic Networks*. PhD thesis, Dept. of Math., Uni. of York., York., 2004. Available online at:
`http://www.richardclegg.org/pubs/thesis.pdf`.

[6] R. G. Clegg. Markov-modulated on/off processes for long-range dependent internet traffic, 2006. Available online at: `http://arxiv.org/abs/cs/0610135`.

[7] R. G. Clegg. A practical guide to measuring the Hurst parameter. *International Journal of Simulation: Systems, Science & Technology*, 7(2):3–14, 2006. Available online at:
`http://www.richardclegg.org/pubs/rgc_ijs.pdf`.

[8] R. G. Clegg and M. M. Dodson. A Markov based method for generating long-range dependence. *Phys. Rev. E*, 72:026118, 2005. Available online at:
`http://www.richardclegg.org/pubs/rgcpre2004.pdf`.

[9] A. Erramilli, R. P. Singh, and P. Pruthi. Chaotic maps as models of packet traffic. In *Proc. 14th Int. Teletraffic Cong.*, volume 1, pages 329–338, 1994.

[10] D. Heath, S. Resnick, and G. Samorodnitsky. Heavy tails and long-range dependence in on/off processes and associated fluid models. *Math. of Oper. Res.*, 1:145–165, 1998.

[11] F. Huebner, D. Liu, and J. M. Fernandez. Queueing performance comparison of traffic models for internet traffic. *Global Telecommunications Conference, GLOBECOM 98*, 1:471–476, 1998.

[12] R. El Abdouni Khayari, R. Sadre, B. R. Haverkort, and A. Ost. The pseudo-self-similar traffic model:application and validation. *Perform. Eval.*, 56:3–22, 2004.

[13] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic. In D. P. Sidhu, editor, *Proc. ACM SIGCOMM*, pages 183–193, San Francisco, California, 1993.

[14] W. E. Leland and D. V. Wilson. High time-resolution measurement and analysis of LAN traffic: Implications for LAN interconnection. *Proc. IEEE INFOCOM*, pages 1360–1366, 1991.

[15] S. Li and C. Hwang. Queue response to input correlation functions: Discrete spectral analysis. *IEEE/ACM Trans. on Networking*, 1(5):522–533, October 1993.

[16] A. L. Neidhardt and J. L. Wang. The concept of relevant time scales and its application to queuing analysis of self-similar traffic (or is Hurst naughty or nice?). In *Proceedings of the 1998 ACM SIGMETRICS joint international conference on Measurement and Modeling of Computer Systems*, pages 222–232, 1998.

[17] Y. Pomeau and P. Manneville. Intermittent transition to turbulence in dissipative dynamical systems. *Commun. Math. Phys.*, 74(2), 1980.

[18] R. H. Riedi, M. S. Crouse, V. J. Ribeiro, and R. G. Baraniuk. A multifractal wavelet model with application to network traffic. *IEEE Special Issue On Information Theory*, 45(April):992–1018, 1999.

[19] S. Robert and J.-Y. Le Boudec. New models for pseudo-self-similar traffic. *Perform. Eval.*, 30:57–68, 1997.

[20] Z. Sahinoglu and S. Tekinay. On multimedia networks: Self similar traffic and network performance. *IEEE Communications Magazine*, pages 48–52, January 1999.

[21] X. J. Wang. Statistical physics of temporal intermittency. *Phys. Rev. A*, 40(11):6647–6661, 1989.

[22] W. Willinger, V. Paxson, R. H. Riedi, and M. S. Taqqu. Long-range dependence and data network traffic. In P. Doukhan, G. Oppenheim, and M. S. Taqqu, editors, *Theory And Applications Of Long-Range Dependence*, pages 373–407. Birkhäuser, 2003.

[23] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson. Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Trans. on Networking*, 5(1):71–86, 1997.