Large-scale Measurement and Modeling of Backbone Internet Traffic

Matthew Roughan and Joel Gottlieb
AT&T Labs – Research
180 Park Av., Florham Park, NJ, 970960, USA
{roughan,joel}@research.att.com

ABSTRACT

There is a brewing controversy in the traffic modeling community concerning how to model backbone traffic. The fundamental work on self-similarity in data traffic appears to be contradicted by recent findings that suggest that backbone traffic is smooth.

The traffic analysis work to date has focused on high-quality but limited-scope packet trace measurements; this limits its applicability to high-speed backbone traffic. This paper uses more than one year's worth of SNMP traffic data covering an entire Tier 1 ISP backbone to address the question of how backbone network traffic should be modeled. Although the limitations of SNMP measurements do not permit us to comment on the fine timescale behavior of the traffic, careful analysis of the data suggests that irrespective of the variation at fine timescales, we can construct a simple traffic model that captures key features of the observed traffic. Furthermore, the model's parameters are measurable using existing network infrastructure, making this model practical in a present-day operational network.

In addition to its practicality, the model verifies basic statistical multiplexing results, and thus sheds deep insight into how smooth backbone traffic really is.

1. INTRODUCTION

A number of recent reports[1,2] have appeared claiming that traffic on backbone networks is "smooth". The sense used in these reports is that it can be generated by an independent Poisson process. These reports contradict some of the previous literature on traffic self-similarity (see for instance[3–6]), and yet have found considerable favour within the networking community because backbone traffic *does* appear to be smooth, at least to the eye. This paper marries the ideas that backbone traffic is smooth (in some sense), but still contains features consistent with LRD, and self-similarity.

The main problem is that there are varying definitions of smoothness. The definition applied in the reports above is much stronger than the implicit definition of smoothness used by the human eye. In fact traffic can be both Long-Range Dependent (LRD) and fractal and still apparently quite smooth.

The reasons that many people have failed to understand that LRD does not contradict smoothness are firstly that LRD is almost always discussed in the context of burstiness, and secondly that most authors concentrate on the parameter that describes the scaling law, the Hurst parameter H, but neglect the peakedness, or variance of the process. This is a critical parameter, along with the mean, when assessing the smoothness of the process. A process with strong LRD exhibited through a high Hurst parameter may still appear very smooth if the relative variance (the variance divided by the mean) is low. In fact, a LRD dependent process may appear more smooth than an uncorrelated process.

This paper sets out, in Section 2, some basic facts about the effect multiplexing has on backbone traffic, namely that multiplexing smooths the traffic, though it does not remove long-range correlations. Section 3 presents empirical results on a large volume of Tier 1 backbone traffic. The results, which agree with the theory, show that the variance of such traffic is proportional to the mean. Therefore the relative variance, our measure of smoothness, decreases as more sources are multiplexed together, and the traffic appears to be smoother. Hence the degree of variation in the traffic may be captured by a single *peakedness* parameter which takes values between [0.05, 0.14] with mean 0.09 (when the traffic is measured in Mbps). Section 4 presents some diagnostic results, and Section 5 the conclusion.

Apart from clearing up the controversy above, and gaining a better understanding of traffic modeling in general (allowing for more realistic simulations of backbone traffic), the results are directly applicable to capacity planning. The result is also useful in time-series analysis[7] of Internet data because the first step of such analysis is to obtain stationary data by removing the trend and periodic components – for instance, the daily variation in the traffic. The removal of these components depends on the way the traffic is modeled, in particular, the effect of the mean on the variance.

2. THEORY

The fundamental idea behind a packet switched network is that by multiplexing multiple sources, a resource (such as link bandwidth) can be shared with a resulting statistical multiplexing gain. While one source is quiescent, another can use the bandwidth.

One simple model for this multiplexing is the elegantly simple On/Off model.[8,9] In this model the sources (each of which is independent) are modeled as a single alternating renewal process in which the On and Off times are modelled as independent random variables with heavy-tails. In this model the number of sources which are On is a binomial random variable, which therefore has mean Np, and Variance $Np(1-p)$. For large N, and p not too small this can be approximated by a Gaussian random variable, and as N goes to infinity with Np a constant, it can be approximated by a Poisson random variable. The traffic rate will be the On rate r times the number of On sources, and the traffic variance will be $r^2Np(1-p)$.

There are therefore three modes by which we might vary:

1. variation in N,
2. variation in p,
3. variation in the On rate r.

Note that, in the core, where there is a large amount of aggregation, N is very large, and even for quite small p a Gaussian model seems appropriate. For p small, we can approximate the variance of the distribution by Np, and hence varying N or p has identical effects. That is, varying the number of connected sources is equivalent to varying the probability that a source is On. Clearly, in this case the variance is proportional to the mean. This is the basic law of multiplexing – as more sources are multiplexed together, the variance per source decreases. Obviously an increase in r will increase the traffic variance in proportion to r^2.

The now classical Norros model[10] also used in[11] for this type of traffic is

$$Y(t) = mt + \sqrt{am}Z(t), \tag{1}$$

where $Y(t)$ is the traffic to arrive up to time t. The parameter a is called the variance coefficient, or the peakedness (though this is sometimes used in other senses). We shall consider the increments of the $Y(t)$ which satisfy $Y(t) = \int_0^t X(s)ds$, and allow the mean to vary slowly over time, so that

$$X(t) = m(t) + \sqrt{am(t)}W(t), \tag{2}$$

where $X(t)$ is the instantaneous traffic rate at time t, $m(t)$ is the time varying mean, and $W(t)$ is the stochastic component which may be chosen to be Gaussian, and have properties such as LRD, but could also be chosen to be white noise. Note that under some types of stochastic process $W(t)$, including Gaussian white noise, the traffic rate may be negative. Typically we use parameters so that this is very rare, and set the rate in those cases to zero.

From Proposition 2.2 of[10] we can verify that this model satisfies the same superposition properties as described above when the stochastic component of the traffic is a (possibly fractional) Gaussian noise. One key aim of this paper is to verify the validity of the Norros model on real traffic data.

2.1. Self-similarity and Long-Range Dependence

Self-similarity and Long-Range Dependence[3] provide a natural, quantitative, and yet wonderfully elegant way of describing, and modeling data traffic. The simplest self-similar model, Fractional Brownian Motion (FBM), has only three parameters, and yet describes the characteristics of the traffic (including its burstiness) over scales ranging from milliseconds up to hours.

Self-similarity and long-range dependence are closely related phenomena. The former relates to the fact that under suitable scaling, the statistics of the traffic are the same at any time scale. The latter relates to correlations in the data which, though decreasing over wider ranges, never become insignificant.

In this paper we consider second order stationary stochastic processes $X(t)$, which consequently have constant mean $m = E[X(t)]$, and variance $\sigma^2 = E[(X(t)-m)^2]$, and an autocovariance (defined by $r(k) = E[(X(t+k)-m)(X(t)-$

$m)]$), which is a function of the *lag* $k = |t - s|$ only. The Fourier Transform of r is known as the *spectral density* and we denote it by f_X.

Self-similarity is the property that as you scale a process (zoom in on the details) you see the same structure repeated. In our case, we are interested in statistical self-similarity, in which the same statistical properties (for instance mean, variance, or marginal distribution) are the same (under a transformation) as you scale the process. The main parameter describing self-similarity is the *Hurst* parameter H.

The closely related phenomena Long-Range Dependence (LRD) is commonly defined by the slow, power-law decrease in the autocovariance function: $r(k) \sim c_\mathrm{r}|k|^{-(1-\alpha)}$, $k \to \infty$, $\alpha \in (0, 1)$, or equivalently as the power-law divergence at the origin of its spectrum: $f_X(\nu) \sim c_\mathrm{f}|\nu|^{-\alpha}$, $|\nu| \to 0$, ([13] p. 160). The power-law decay is such that the sum of all correlations (out from any lag) is always appreciable, even if individually the correlations are small. The past therefore exerts a long-range influence on the future.

The main parameter of LRD is the dimensionless scaling exponent α. It describes the qualitative nature of scaling – how behavior on different scales is related. The related parameters, c_r and c_f, are quantitative parameters which give a measure of the magnitude of LRD induced effects. The parameters may be estimated jointly using the Abry-Veitch wavelet based estimator,[14] or separately by a number of other techniques.[13] It is common practice to describe LRD through the *Hurst* parameter $H = (1 + \alpha)/2$.

One very powerful set of models that has been applied to traffic is that of the Gaussian processes. These are processes with a Gaussian marginal distribution. A simple example is White Gaussian Noise (WGN), an uncorrelated process. A more general example (which includes WGN) is that of FGN (the increments of FBM mentioned above) which is a LRD process, and therefore has correlations over long time scales. The advantage of FGN is that it is completely characterised by the three parameters: its mean, its variance* and H. WGN is simply the case with $H = 0.5$†.

Figure 1 shows two examples: WGN ($H = 0.5$), and FGN ($H = 0.9$) each added to a sine function. Although each has the same standard deviation (0.025) both look comparatively smooth (because their variance is small in comparison to the scale of the plot), and in fact the FGN plot looks in some repects smoother than the case of uncorrelated noise.

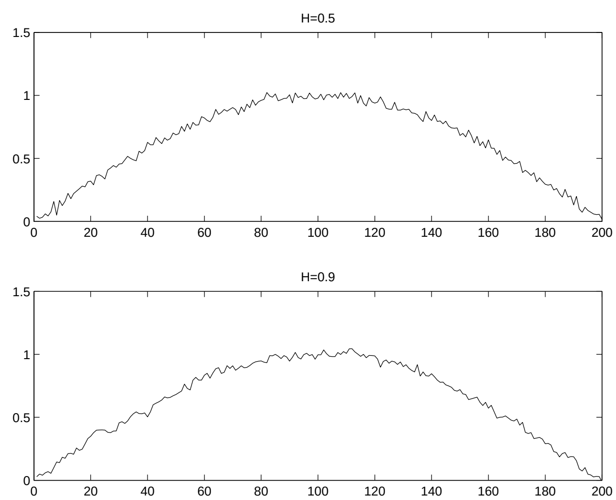

Figure 1. Plots of white Gaussian noise ($H = 0.5$), and fractional Gaussian noise ($H = 0.9$) added to a sine function. In each case the standard deviation of the noise is 0.025.

One quantitative measure of smoothness is the mean squared differences of the data. That is

$$S = \sum_{i=1}^{N-1} (y_{i+1} - yi)^2 , \tag{3}$$

*The variance of FGN is directly related to c_f and H.

†Note that in general $H = 0.5$ only implies a short-range dependent process, but in the case of FGN it also implies that the process is uncorrelated.

with a smaller value indicating a smoother process. By this measure the smoothness of the white GN process above is 0.0015, whilst the FGN process has a smoothness of 0.0007, less than half the value of the uncorrelated process. Hence we see that the LRD process is actually smoother than the uncorrelated process.

When LRD processes are multiplexed the resultant traffic has an asymptotic Hurst parameter equal to the highest Hurst parameter of the multiplexed traffic. In the case where the Hurst parameters are all equal, then the resultant traffic will have the same Hurst parameter (for instance see[10]).

2.2. Poisson limits

One of the arguments applied in[1, 2] is the traditional argument for why Public Switched Telephone Network (PSTN) calls appear as a Poisson process. The argument is that as we superpose large numbers of stochastic processes (the phone calls of each individual), and thin suitably (so that the mean remains finite) we approach a Poisson process in the limit. The authors of[1, 2] argue that this occurs as large numbers of independent connections are superposed in aggregate backbone traffic.

The argument lacks teeth in this context, however, as it relies on large numbers of *independent* traffic streams being superposed. In fact, when traffic streams are super-posed at reasonable utilization levels[‡] the network influences the traffic stream, firstly by not allowing packets to overlap, and secondly through TCP feedback mechanisms. Hence the Poisson process limit argument is not necessarily applicable here (although it might appear valid on some time-scales, and sets of measurements).

3. EMPIRICAL RESULTS

3.1. SNMP – the Simple Network Management Protocol

The models we build are critically dependent on the data on which they are built. In this paper, we analyze Simple Network Management Protocol (SNMP) traffic and fault data extracted from an archive that includes more than 1 year's worth of data collected from a large Tier-1 ISP's backbone network. SNMP is unique in that it is supported by essentially every device in an IP network. The SNMP data that is available on a device is defined in a abstract data structure known as a Management Information Base (MIB). An SNMP *poller* periodically requests the appropriate SNMP MIB data from a router (or other device). Since every router maintains a cyclic counter of the number of bytes transmitted and received on each of its interfaces, we can obtain basic traffic statistics for the entire network with little additional infrastructure.

Unfortunately, as a practical matter, SNMP data has many limitations – for instance missing data (it may be missing because SNMP uses UDP transport, or it may be lost while copying to our research archive), incorrect data (through poor router vendor implementations), and a coarse sampling interval. However, the most important limitation is that SNMP only provides aggregate link statistics. We cannot determine anything about the type of traffic using the link, nor its source or destination, and we can only sample at coarse intervals.

This work also uses topology information derived as in.[15] The archived data we have access to has similar limitations to SNMP data, being only updated daily, and with missing files.

These limitations do not, in general, restrict operational use of SNMP data, as most data is interpreted by eye for operations. However, the limitations make algorithmic analysis such as time series analysis difficult on this data. We have gone to considerable lengths to reduce the impact of these features of the data through careful post-processing: discarding ambiguous and incorrect data where possible, and using SNMP fault data to determine causes of some anomalies. Thanks to these efforts, and carefully choosing models and analysis that are not sensitive to the data quality we can use even such poor data in some quite sophisticated analysis.

One point to note is that many past analysis of such data have been done in the "busy hour", but such analyses suffers from one major feature. Given a strong daily and weekly cycle (which is the case here), and five minute traffic data (which is typical for SNMP, and what we have available here), one has only 12 sample points per week. To obtain enough data for a reasonably accurate analysis, one must average over many weeks, and over such time periods, the time at which the busy hour occurs may change. This motivates a more complete analysis.

We have more than 1 year's worth of SNMP data available for a set of OC48 (2.5 Gbps) inter-city backbone links from a Tier 1 ISP. This forms the basis of this study. The data is polled at five minute intervals. There are 52 links used in this study, and because of traffic asymmetry[§] there are 104 simplex links to consider. The links are typically not heavily

[‡]In[1, 2] the results are applied in the limit as links become heavily loaded.

[§]Traffic on the backbone is asymmetric not just because web traffic is asymmetric, but also because inter-peer routing is asymmetric.

loaded, but most run at a significant level of utilization.

3.2. Data Analysis

In order to compare the mean and variance of the process we must estimate these components. As the data shows strong evidence of non-stationarity – it has strong daily, and weekly cycles, and a long term exponential trend – we cannot get an overall estimate of these properties. We could attempt to remove the seasonal¶ and trend components, but methods for removing such components start with a model of the data, which is what we are trying to verify here.

The method we have chosen is to break the data into blocks, and estimate the mean, and variance on each block of data. There is a tradeoff in choice of blocksize: a smaller block will track changes in the mean and variance more accurately, but each block will have less data, and so the estimates will be less accurate; a larger blocksize will produce more accurate estimates, but may not be able to track the non-stationarity in the data with suitable accuracy. We have chosen a blocksize of one hour which seems to have the best tradeoff. Each block has 12 samples spaced at five minute intervals. We denote the mean of block i by x_i, and the variance by y_i.

We do not know a priori whether the data within a block are correlated, and hence we cannot provide an analytic measure of the variance of our estimates y_i‖. To provide such a measure, we have further grouped the data into 100 groups by the estimated mean over the hour. We take the average of the variance of traffic in each block within a group to be an estimate of the variance of the traffic with mean rate given by the rate in the group, and because in most cases we have multiple data points within a group, we can provide an empirical estimate of the variation of the variance estimate. More precisely, take groups $\mathcal{G}_j$ defined by

$$\mathcal{G}_j = \{i | hj \leq x_i < h(j+1)\}, \text{ for } j = 0, 1, 2, \ldots, \tag{4}$$

where h is small relative to the link loads (in our case chosen so that we end up with 100 groups, i.e. $h = x_{\max}/100$). The number of data points in each bin N_j, the mean traffic rate for the bin X_j, the variance estimate Y_j, and the empirical estimate of the variance of Y_j (which we denote σ_j^2) are given by

$$N_j = \sum_{i \in \mathcal{G}_j} 1, \tag{5}$$

$$X_j = h(j + 0.5), \tag{6}$$

$$Y_j = \frac{1}{N_j} \sum_{i \in \mathcal{G}_j} y_i, \tag{7}$$

$$\sigma_j^2 = \frac{1}{N_j(N_j - 1)} \sum_{i \in \mathcal{G}_j} (y_i - Y_j)^2. \tag{8}$$

We can then perform a weighted polynomial regression[16] on the (X_j, Y_j), using the σ_j^2 to form the weights. We only use groups which contain more than 10 blocks of data (that is $N_j > 10$), so that the estimate of σ_j^2 will be more reliable. We perform linear, and quadratic fits, and we further constrain these fits to pass near (0,0) by placing a highly weighted data point here. In Table 1 we show the results, giving Q the quality of the linear fit.

A good linear fit will have $Q > 0.1$, a possible fit $Q > 0.001$ and a poor fit $Q < 0.001$. Of the 102 sets of data, 18 have $Q > 0.001$. Figure 2 shows examples of good linear fits. The graphs show (X_j, Y_j) (dots) with the vertical bars indicating the 95th percentile confidence interval given by $\pm 1.96\sigma_j$. The blue bars show the data used in the weighted regression, and the magenta bars the data excluded due to too few data points being available for that group. The solid green line shows the linear fit, and the red line shows a quadratic fit to the data (used later).

However, a poor value of Q whilst indicating a poor linear fit in a strict mean squared error sense, does not necessarily mean that the data are not reasonably linear. Further, in Section 4.3 we present an example (in Figure 6) where almost exactly the same traffic is carried on two links (links 1 and 2) and one can see in Table 1 that the Q values for direction b are quite disparate, though in the case of link 1 the value is high. There are many other cases that have low Q values, but appear to have a generally good linear fit to the data.

¶We use seasonal to denote all of the period components.

‖x_i will also contain an error which we will not consider here as it should be smaller than the error in y_i

The mean squared error is quite sensitive to outliers, and so a small amount of outlying data can mean a low Q value. This data has many outliers due to either rapidly changing network behaviour when rerouting, or traffic surges as may be caused by Denial of Service (DoS) attacks and flash crowds. The fixed one hour block size is a liability in these cases, and can lead to variance estimates which are much greater than the true variance.

To provide another estimate of whether a linear fit to the data is appropriate we have compared the linear fit to a quadratic fit which would apply if the variance were proportional to the square of the mean. Table 1 also shows the mean squared errors for the linear and quadratic fits, denoted by χ_1^2 and χ_2^2 respectively. Clearly in many cases the two are quite close together. For comparison we use the relative difference between these,

$$D = \frac{\chi_1^2 - \chi_2^2}{\chi_2^2}. \tag{9}$$

Note that for the cases where $Q > 0.001$ the maximum value of D is 0.48, but the majority of values are below 0.18. Of the total 102 data sets, 51 have $D <= 0.18$ and 68 have $D <= 0.48$. Over half of the data sets appear not to have a significantly better quadratic fit. We could, of course, apply a more robust fit to the data to obtain the linear fit, but our aim here is not to obtain the most accurate possible estimate of the slope a, but to decide whether the linear fit is appropriate.

There are still some data sets for which there is a better quadratic fit, though in many of these cases the quadratic fit is not particularly good either. We present some examples of these in Figure 3 and shall discuss them further below.

Table 1 also shows the number of data points (x_i, y_i) used in each regression. The *goodfit* columns of the table is provided to make it easier to note which rows satisfy $Q > 0.001$ and or $D <= 0.18$. Note slope gives an estimate of the peakedness parameter a.

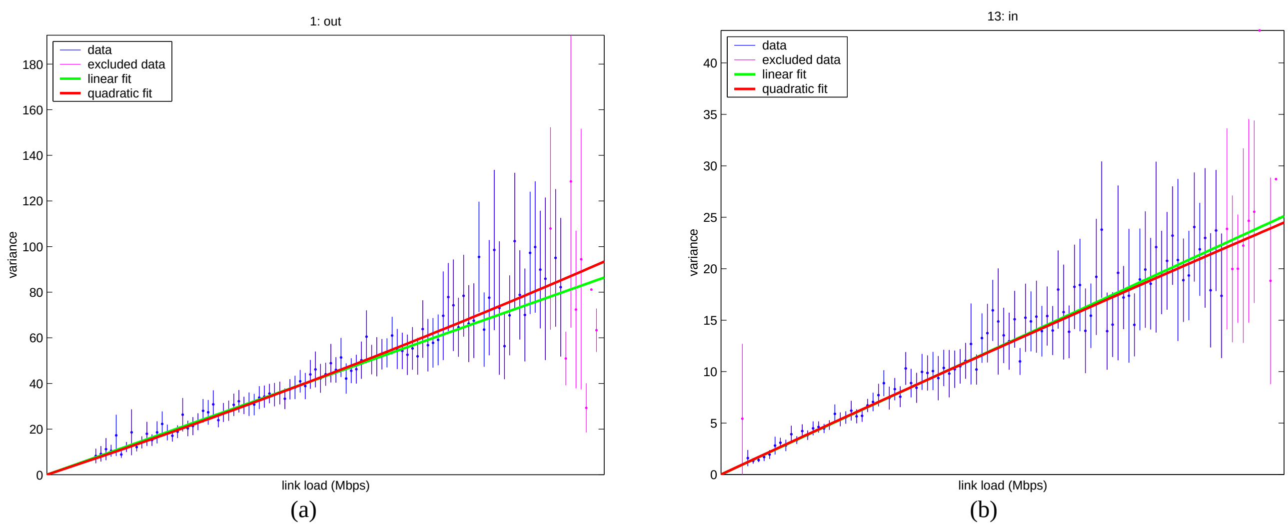

Figure 2. Examples: the figures show two examples for which Q is large. The graphs show (X_j, Y_j) (dots) with the vertical bars indicating the 95th percentile confidence interval given by $1.96\sigma_j$. The blue bars show the data used in the weighted regression, and the magenta bars the data excluded due to too few data points being available for that group. The solid green line shows the linear fit, and the red line shows a quadratic fit to the data. NB: the absolute values on the x-axis have been deleted because absolute link loads on the network are considered proprietary information, however, it is the relative value that matters here, not the absolute value. These links are significantly loaded, but not overloaded.

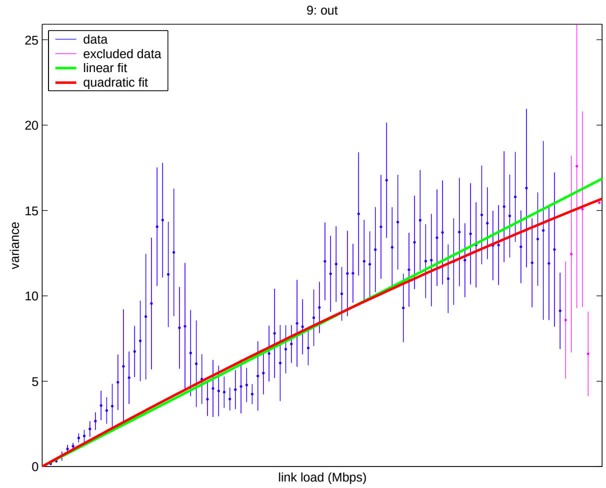

(a)

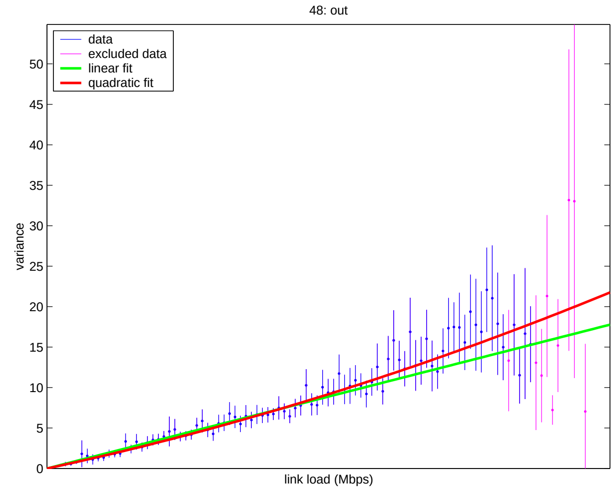

(b)

Figure 3. Examples: these figures show two cases with a poor linear fit.The graphs show (X_j, Y_j) (dots) with the vertical bars indicating the 95th percentile confidence interval given by $1.96\sigma_j$. The blue bars show the data used in the weighted regression, and the magenta bars the data excluded due to too few data points being available for that group. The solid green line shows the linear fit, and the red line shows a quadratic fit to the data. NB: the absolute values on the x-axis have been deleted because absolute link loads on the network are considered proprietary information, however, it is the relative value that matters here, not the absolute value. These links are significantly loaded, but not overloaded.

Table 1: Quality of a linear fit to the variance/mean relationship. i denotes the link (each direction is given separately). In cases with $Q > 0.1$ the fit is good, while where $Q > 0.001$ the fit is reasonable. The slope is the slope of the linear fit (in Mbps) and the χ^2 variables show the quality of the fit for the linear χ_1^2 and quadratic χ_2^2 fits. The final column indicates the quality of the fit based on the criteria $Q > 0.001$ and $D < 0.18$ where one (two) star(s) means one (two) of these is satified.

	direction a						direction b					
i	**Q**	**slope**	χ_1^2	χ_2^2	**N**	**goodfit**	**Q**	**slope**	χ_1^2	χ_2^2	**N**	**goodfit**
1	0.000000	0.078	237.3	149.8	7816		0.338982	0.125	95.0	88.1	7816	⋆⋆
2	0.000000	0.074	221.4	116.4	7737		0.000000	0.125	170.8	170.8	7737	⋆
3	0.000150	0.084	119.7	115.0	7384	⋆	0.000000	0.102	172.0	151.6	7384	⋆
4	0.000575	0.075	113.7	109.2	4494	⋆	0.000000	0.096	188.1	171.0	4494	⋆
5	0.000002	0.105	159.8	155.9	7785	⋆	0.337541	0.088	72.3	70.7	7785	⋆⋆
6	0.000001	0.104	146.4	131.1	7036	⋆	0.000000	0.085	147.9	99.2	7036	
7	0.000001	0.105	149.7	144.5	7730	⋆	0.001058	0.083	129.5	129.5	7730	⋆⋆
8	0.000000	0.101	391.7	309.2	6980		0.000000	0.133	193.4	186.4	6980	⋆
9	0.000000	0.123	511.2	507.6	8038	⋆	0.000000	0.140	525.7	512.1	8038	⋆
10	0.000000	0.080	170.5	115.2	6964		0.000000	0.118	235.3	199.9	6964	⋆
11	0.000000	0.124	223.2	221.0	7285	⋆	0.000000	0.084	310.8	140.3	7285	
12	0.155436	0.065	100.3	85.7	7652	⋆⋆	0.000000	0.101	204.8	108.4	7652	
13	0.001041	0.067	133.3	133.2	7711	⋆⋆	0.000000	0.106	175.0	104.3	7711	
14	0.000000	0.121	223.4	131.4	7520		0.000068	0.106	133.6	95.5	7520	
15	0.000022	0.106	145.2	133.1	7586	⋆	0.000000	0.114	352.8	114.5	7586	
16	0.000000	0.121	139.0	91.2	7467		0.000000	0.074	582.1	313.6	7467	
17	0.004553	0.076	118.0	99.4	7175	⋆⋆	0.000000	0.118	222.6	123.0	7175	
18	0.000000	0.125	209.5	157.5	7348		0.000000	0.112	199.5	154.7	7348	
19	0.000000	0.100	838.3	578.6	6389		0.000000	0.066	1276.5	993.1	6389	
20	0.344248	0.125	39.9	39.0	835	⋆⋆	0.037166	0.051	59.7	57.8	835	⋆⋆

Table 1: (continued)

	direction a						direction b					
i	**Q**	**slope**	χ_1^2	χ_2^2	**N**	**goodfit**	**Q**	**slope**	χ_1^2	χ_2^2	**N**	**goodfit**
22	0.000000	0.075	265.8	212.5	7160		0.000000	0.042	258.8	223.7	7160	⋆
23	0.000000	0.038	346.1	239.2	6097		0.000000	0.076	302.6	171.8	6097	
24	0.344248	0.125	39.9	39.0	835	⋆⋆	0.037166	0.051	59.7	57.8	835	⋆⋆
25	0.000000	0.080	276.6	273.0	7468	⋆	0.000000	0.112	248.3	155.8	7468	
26	0.000000	0.080	276.6	273.0	7468	⋆	0.000000	0.112	248.3	155.8	7468	
27	0.000000	0.121	194.5	153.5	5091		0.159576	0.069	67.6	64.4	5184	⋆⋆
28	0.314080	0.071	42.7	37.7	7371	⋆⋆	0.000000	0.096	107.6	95.4	7371	⋆
29	0.000000	0.080	156.2	118.3	6878		0.000001	0.062	130.5	114.7	6878	⋆
30	0.000000	0.087	281.5	281.4	7485	⋆	0.000000	0.102	357.3	207.8	7485	
31	0.000000	0.087	281.5	281.4	7485	⋆	0.000000	0.102	357.3	207.8	7485	
32	0.000000	0.071	294.1	286.3	7086	⋆	0.000000	0.055	285.8	122.0	7086	
34	0.000000	0.092	129.8	86.7	6901		0.000000	0.070	161.4	152.9	6901	⋆
35	0.000004	0.079	30.3	22.9	7710		0.000079	0.112	1.6	20.9	7710	⋆
36	0.025703	0.089	79.6	79.4	7376	⋆⋆	0.000433	0.084	95.5	73.9	7376	
37	0.000002	0.105	147.5	147.1	7381	⋆	0.000000	0.118	200.6	92.4	7381	
39	0.000000	0.080	182.4	178.6	6972	⋆	0.000443	0.093	121.2	113.5	6972	⋆
40	0.015293	0.092	102.7	100.6	6811	⋆⋆	0.000001	0.078	148.5	140.3	6811	⋆
41	0.000000	0.105	242.0	114.0	6536		0.000000	0.121	173.6	124.6	6647	
42	0.000000	0.136	170.9	117.1	7267		0.010092	0.132	115.8	114.0	7267	⋆⋆
43	0.066886	0.142	106.4	71.5	5989	⋆	0.000002	0.128	162.3	135.7	5989	⋆
44	0.000000	0.087	257.1	170.4	7518		0.000000	0.075	432.2	184.9	7518	
45	0.000000	0.087	481.3	381.9	7046		0.000000	0.080	384.5	210.1	7046	
46	0.000000	0.064	391.1	327.9	7138	⋆	0.000000	0.041	481.9	379.3	7138	
47	0.000000	0.123	225.6	211.1	7092	⋆	0.000000	0.057	211.2	103.3	7092	
48	0.000000	0.038	346.1	239.2	6097		0.000000	0.076	302.6	171.8	6097	
49	0.000000	0.061	191.0	119.8	6665		0.000000	0.041	347.2	219.1	6665	
50	0.000000	0.059	190.3	145.4	6526		0.000000	0.039	313.9	297.3	6526	⋆
51	0.590362	0.081	68.6	68.5	7188	⋆⋆	0.000000	0.079	163.0	153.6	7188	⋆
52	0.000100	0.090	116.2	108.2	6765	⋆	0.042494	0.082	93.9	91.5	6765	⋆⋆

One should not be overly surprised that many of the data sets do not show a good fit to a straight line. There are many reasons why the data would not show a good fit:

1. Non-stationarity in traffic parameters (other than the mean). For instance, the peakedness a may change over time, either because a trend in customer behaviour or mix, or because different sets of customers (with different activities) are active at different times of day. Note each chart is based on more than 1 year's worth of data!

2. We do not actually know the mean rate $m(t)$. The estimate provided by one hourly averages could be thrown out by a number of factors such as non-stationarity on finer time scales than 1 hour, or outliers in the data.

3. Transient events such as rerouting (for instance when a new link is put in the network), DoS attacks, and flash crowds may distort the traffic behaviour in unexpected ways.

4. Even on a backbone link it may be possible that a single, or small group of sources dominate the traffic.[17] This breaks the basic assumptions of large numbers of customers being multiplexed. This can occur particularly on lightly loaded links, of which there are quite a few.

5. Basic problems with the SNMP data can lead to artifacts in the data.

6. Under high link loads the link may act as a bottleneck some of the time, forcing TCP congestion control to kick in. This will distort the model, which assumes no feedback.

7. Rate increases are not always because of increased numbers of sources being multiplexed, but can also be because of increases in access rates of customers (for instance migration of dial-up customers to broadband access). SNMP

data is insufficient to check this because we cannot see into the individual flows in the data to see the reason for a traffic increase.

In fact, given these problems, it is amazing that such a good fit to a straight line is obtained in the number of cases seen over such a large set of data (more than one year's worth). We have deliberately included the negative cases in order to provide perspective on these results. We shall provide evidence for the above reasoning in the following section.

4. DIAGNOSTICS

In any data analysis we need to be careful that the results are not artifacts of assumptions in our data analysis. Furthermore, we have suggested some reasons above as to why the data does not always fit the model. This section provides some diagnostics to help assess whether what we are seeing is real, and to test the above assertions about why the model does not always hold.

4.1. Simple diagnostics

We start with a few simple tests to look for systematic effects in the data which should not be there. Figure 4 (a) shows a simple scatter plot of the estimated slopes where $Q > 0.001$. We can see that although these are not Gaussian, there do not appear to be any obvious artifacts in the data set. Figure 4 (b) shows the relationship between N and Q. There is a small negative correlation (-0.28) between Q and N. This is entirely due to the two data points with $N \sim 1000$, and without these points the correlation is small[**] (0.06). It is not surprising that the shorter data sets get a better fit because non-stationarity is less likely to play a role in shorter data sets (see below).

Figure 4 (c) shows the relationship between the average load on the links[††] and Q. The correlation coefficent here is 0.076, once again small. Figure 4 (c) shows the relationship between the slope and Q. The correlation coefficient is 0.064, once again small. The lack of correlation in these cases is a good sign – it means that the quality of fit was not dependent on the parameters we are attempting to fit to.

4.2. Non-stationarity of a

Earlier, we postulated that one of the reasons that the data don't fit a simple linear plot is non-stationarity, in particular in a. In this case the simple model suggested by (2) is not correct, and so there is no reason for our data to fit a straight line. The fact that our model breaks down over a period of a year should not be seen as a problem – the main reason we wish the model to remain stationary is to make broad ranging studies such as this possible. If the parameter remains reasonably constant even over periods as short as a day it could be useful in traffic analysis (many traffic studies have less than a day's traffic in any one segment of data).

To test the stationarity of a we shall break the data sets up into blocks each a month long, and compare the estimates of a on each month of data separately.

To illustrate the results we plot the monthly estimates in Figure 5. The figure shows directions a and b on link 1. The circles show the monthly estimates, with the vertical lines giving confidence intervals for the estimates, and the horizontal line being the overall mean. The cyan data points are those with $Q < 0.001$. We should ignore the cyan data points below, as they are being confounded by other factors.

From Table 1 we can see that direction b has good linear fit, while direction b has a poor linear fit. Figure 5 shows that the value of a is quite non-stationary in the direction-a data, but almost consistent with being stationary in direction b.

One should also note that in the shorter month long segments of data we see good values of Q more frequently. This is in part because over shorter times the stationary approximation is more accurate, but also because with a smaller data set, the variance of estimates is larger, and hence allow a "looser" fit.

The two cases above cannot explain all of the poor fits in the data. For instance in link (6b) the finer grained estimates appear to be stationary, but the value of Q is poor, and the χ^2 values have a large difference. Furthermore, link (9a) has very few good Q values even on the smaller monthly scale, although the difference of its χ^2 values is small. Therefore, non-stationarity is but one of a set of problems in this analysis.

[**] We use small here to mean that the correlation lies inside the region in which 95% of uncorrelated data would fall, given by $1.96/\sqrt{n}$ where n is the number of data points available ($1.96/\sqrt{n} = 0.198$ in this case).

[††] We have to erase the average link loads here as this data is considered proprietary.

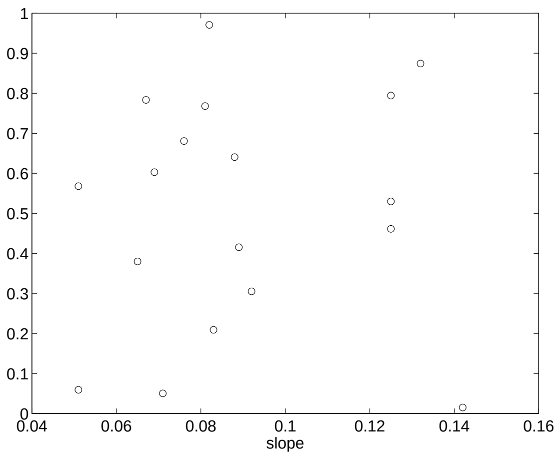

(a) slope scatter plot

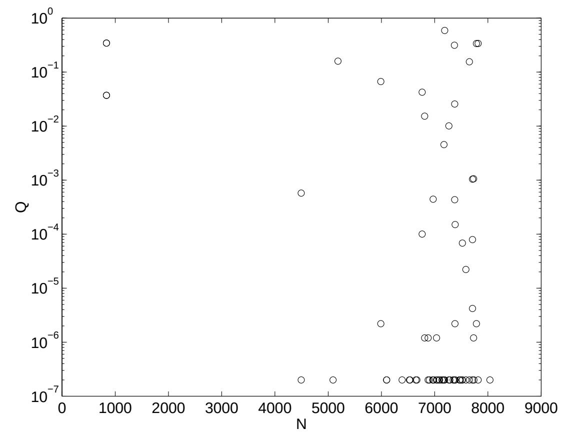

(b) N vs Q

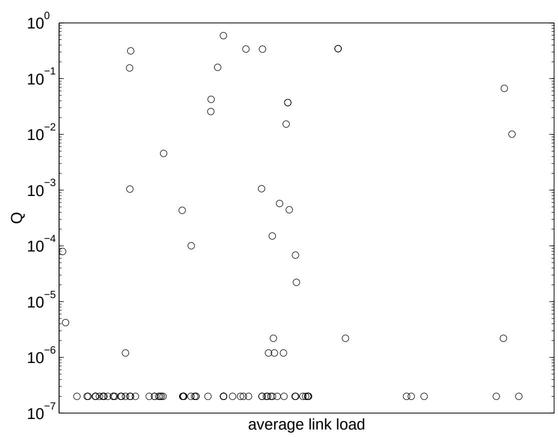

(c) Load vs Q

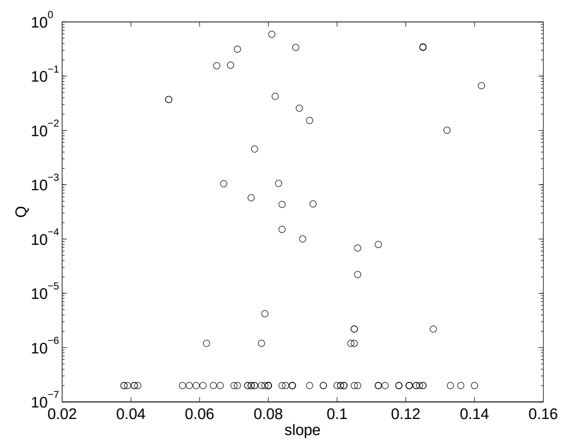

(d) Slope vs Q

Figure 4. Diagnostics: note zero values in log scale replaced by 2.0e-7.

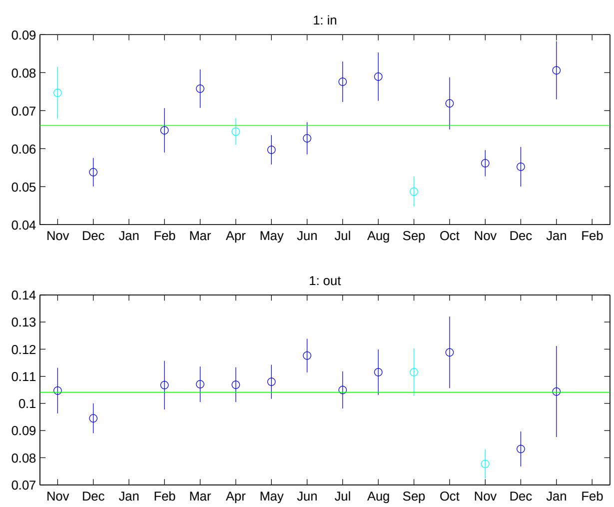

Figure 5. non-stationarity test

4.3. Topological information

There is more information in this data that we can exploit to assess the validity of a linear model. In particular some links occur in pairs, across which load is shared. An example of such a pair of links are 1 and 2 – this pair of links have almost the same traffic (particularly at a time scale of 5 minutes), except during rare rerouting events. This is consistent with the estimated slopes in each link, but not with the estimated value of Q. In one case of the four there is a very good fit. In the other three cases there is a poor fit.

Figure 6 shows the two links in the out direction. The two plots are very similar, but that the second has a number of outliers. As noted above the value of Q is quite sensitive to outliers. This adds weight to the idea that anomalies in the data can lead to outliers which then reduce the quality of the fit, as measured by Q. However, note that the difference between the χ^2 value for the linear and quadratic fit is small in both cases.

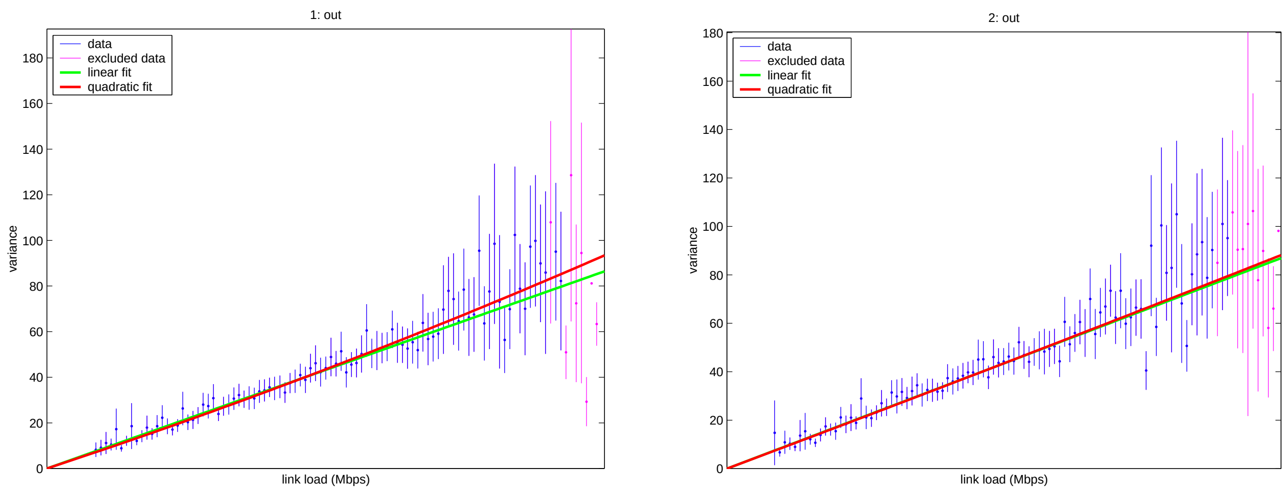

(a) link 1

(b) link 2

Figure 6. An example of data from a pair of links between the same pair of cities. Link 1 (direction b) is shown in (a), and link 2, direction b is shown in (b). Note that the plots appear very similar, but not the same, although the traffic on the two links is almost identical, except for rare events like rerouting. I particular link 1 has a much higher value of Q.

The results of the diagnostics above all support the model (2) by showing no confounding factors, and by justifying the reasons the model is not a good fit in some of the data sets.

5. CONCLUSION

To summarize the findings of this paper, we used a large amount of SNMP data which covered over 50 backbone, inter-city links of a Tier 1 ISP, over a period of more than one year. The data strongly supports the Norros model given in (2) in which the variance of the traffic is proportional to the mean. There are exceptions when data over a very long period is examined (in part due to non-stationarity over this time period), but even in more than a year's worth of data, we found more than half of the links examined fit the model well.

The main consequence of these results is an estimate of multiplexing gain. We can estimate this gain by examining the relative variance, as measured by the standard deviation divided by the mean, which is to say that as link speeds increase, and the amount of multiplexing increases (the number of users increase) the relative variance will decrease in proportion to the inverse square root of the mean.

Note though, that if the increase in traffic is not due to increased numbers of sources being multiplexed, but rather is due to an increase in source rates, such as might occur if subscribers are migrated to broadband access from dial-up, then the reduction in variance might not occur, and in fact the variance of traffic could potentially increase depending on changes in subscriber behavior.

The values of the peakedness parameter a take the range [0.05, 0.14] with mean 0.09 (when the traffic is measured in Mbps). More realistic values of a should aid in simulations of realistic traffic. If the packet arrivals were an independent, homogeneous Poisson process then the variance would be equal to the mean, and a would be 1. This is not to say that a Poisson model could not appear to be a reasonable approximation at some time-scale, but that the Poisson approximation is not valid at the time scale of these measurements. A key advantage of self-similar models is that they provide a parsimoneous model that is valid across a range of time scales.

Acknowledgement

The data used in this paper was collected by a group of people to who we are indebted. In particular we would like to mention Fred True and Carsten Lund who were particularly helpful in the collection and analysis of this data.

REFERENCES

1. J. Cao, W. S. Cleveland, D. Lin, , and D. X. Sun, “Internet traffic tends to Poisson and independent as the load increases,” tech. rep., Bell Labs, 2001. Available at `citeseer.nj.nec.com/426819.html`.
2. J. Cao, W. S. Cleveland, D. Lin, , and D. X. Sun, “The effect of statistical multiplexing on internet packet traffic: Theory and empirical study,” tech. rep., Bell Labs, 2001. Available at `http://cm.bell-labs.com/cm/ms/departments/sia/wsc/publish.html`.
3. W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, “On the self-similar nature of Ethernet traffic (extended version),” *IEEE/ACM Transactions on Networking* **2**, pp. 1–15, Feb 1994.
4. W. Willinger, V. Paxson, and M. S. Taqqu, “Self-similarity and heavy tails: Structural modeling of network traffic,” in *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*, R. Adler, R. Feldman, and M. S. Taqqu, eds., pp. 27–53, Birkhauser, Boston, 1998.
5. V. Paxson, “Empirically-derived analytic models of wide-area TCP connections,” *IEEE/ACM Transactions on Networking* **2**(4), pp. 316–336, 1994.
6. V. Paxson and S. Floyd, “Wide-area traffic: The failure of poisson modeling,” *IEEE/ACM Transactions on Networking* **3**(3), pp. 226–244, 1994. Available at `http://www.aciri.org/floyd/papers.html`.
7. P. Brockwell and R. Davis, *Time Series: Theory and Methods*, Springer-Verlag, 1987.
8. M. S. Taqqu, W. Willinger, and R. Sherman, “Proof of a fundamental result in self-similar traffic modeling,” *Computer Communications Review* **27**, pp. 5–23, 1997.
9. 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,” *Proceedings of the ACM/SIGCOMM’95* , 1995. Available at `http://www.acm.org/sigcomm/sigcomm95/sigcpapers.html`.
10. I. Norros, “A storage model with self-similar input,” *Queueing Systems* **16**, pp. 387–396, 1994.
11. W.-C. Lau, A. Erramilli, J.L.Wang, and W. Willinger, “Self-similar traffic parameter estimation: a semi-parametric periodogram-based algorithm,” in *Global Telecommunications Conference, GLOBECOM ’95*, 1995.
12. A. Erramilli. private communications.
13. J. Beran, *Statistics for Long-Memory Processes*, Chapman and Hall, New York, 1994.
14. P. Abry and D. Veitch, “Wavelet analysis of long-range dependent traffic,” *IEEE Trans. on Info. Theory* **44**(1), pp. 2–15, 1998.
15. A. Feldmann, A. Greenberg, C. Lund, N. Reingold, and J. Rexford, “Netscope: Traffic engineering for IP networks,” *IEEE Network Magazine, special issue on Internet traffic engineering* , pp. 11–19, March/April 2000.
16. W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C: The art of Scientific Computing*, Cambridge University Press, 2nd ed., 1988-92. Available at `http://www.nr.com/`.
17. S. Sarvotham, R. Riedi, and R. Baraniuk, “Connection-level analysis and modeling of network traffic,” in *ACM SIGCOMM Internet Measurement Workshop*, 2001. Available at `http://www.icir.org/vern/imw-2001/proceedings.html`.