Experience in Measuring Internet Backbone Traffic Variability: Models, Metrics, Measurements and Meaning

Matthew Roughan[a], Albert Greenberg[a], Charles Kalmanek[a], Michael Rumsewicz[b], Jennifer Yates[a] and Yin Zhang[a]

[a]AT&T Labs – Research, 180 Park Av., Florham Pk, NJ, 07932, USA
Email: {roughan,albert,crk,jyates,yzhang}@research.att.com

[b]Teletraffic Research Centre, University of Adelaide, Adelaide 5005
Email: mrumsewi@trc.adelaide.edu.au

Understanding the variability of Internet traffic in backbone networks is essential to better plan and manage existing networks, as well as to design next generation networks. However, most traffic analyses that might be used to approach this problem are based on detailed packet or flow level measurements, which are usually not available throughout a large network. As a result there is a poor understanding of backbone traffic variability, and its impact on network operations (e.g. on capacity planning or traffic engineering). This paper introduces a metric for measuring backbone traffic variability that is grounded on simple but powerful traffic theory. What sets this metric apart, however, is that we present a method for practical implementation of the metric using widely available SNMP traffic measurements. In addition to simulations, we use a large set of SNMP data from an operational IP network on the order of 1000 nodes to test our methods. We also delve into the degree and sources of variability in real backbone traffic, providing insight into the true nature of traffic variability.

1. Introduction

Despite a significant amount of research addressing Internet traffic models (for instance see [1–4]), there is not yet wide-spread agreement about the characteristics of backbone Internet traffic. This problem is exacerbated by exaggerated reports on Internet traffic growth and variability [5,6], by the challenges associated with Internet traffic measurements [7], and a lack of understanding of the applicability of results such as the discovery of self-similarity in traffic [1–3]. For instance, in [5], dire claims are made on the basis of the notion that large volumes of traffic *slosh* around the Internet in a highly irregular way.

Obtaining the data necessary to develop an accurate and current view of backbone traffic requires significant investment in measurement infrastructure. Although detailed packet traces are collected on a limited scale in many networks, these traces are not just difficult to collect on high speed links (OC48 and greater) but also represent huge volumes of data, and are often aggregated into simple statistics before being collected for analysis, where they are available at all. The result is that this information is almost never available in the detailed form needed for most applicable traffic models. Nearly all network management tasks are therefore carried out on coarse aggregate statistics of the traffic.

Nonetheless, understanding Internet backbone traffic is crucial for evolving the Internet architecture, doing capacity planning, traffic engineering, and meeting service level agreements. In particular, our investigation was specifically motivated by the question: to what extent does traffic variability justify the need for a re-configurable optical network below the IP layer to provide bandwidth management. Such an optical network would allow IP routers equipped with the appropriate interfaces to request additional point-to-point capacity when needed, and to reconfigure existing capacity between routers [8–10]. Routers might need additional capacity due to congestion resulting from any of a number of causes: major events (September 11th), re-routing events triggered by failures, transient overloads due to Denial of Service (DoS) attacks or flash crowds, or externally induced traffic shifts from peer networks. Alternatively, we can view this problem through the lens of over-provisioning, namely to what extent does the IP layer need to be over-provisioned to meet its service level agreements with high reliability.

We address the problem of backbone traffic variability by looking at aggregate link statistics collected via the Simple Network Management Protocol (SNMP). From these statistics it is clear that the traffic has both daily and weekly periodic components, as well as a longer-term trend. Superimposed on top of these components are shorter time scale stochastic variations. Given these characteristics, we develop a simple, but powerful stochastic model for backbone traffic (based on the Norros model [11]), and then use that model to derive an empirical metric referred to here as the peakedness parameter (though this term is often used in different ways in traffic modeling), that provides a measure of the traffic variability. However, note that the model does not require any specific form of stochastic component, and either a Long-Range Dependent (LRD) or Short-Range Dependent (SRD) model could be used with equal facility. We believe that this metric will be useful to network operators in both architecture evolution and traffic management, e.g., allowing network operators to determine whether (or when) it makes sense to layer IP over a re-configurable optical network, assisting in provisioning backbone capacity, tuning OSPF links weights, etc. An important feature of this model is parsimony – only one parameter is required to describe the most important features of the stochastic variation in the traffic, and this parameter can be estimated from standard SNMP traffic measurements.

We test this approach on a set of real SNMP data from one of the largest operational Internet backbones in North America (AT&T). A major insight of this paper is that backbone traffic is predominantly composed of a regular and predictable component, though it does have a significant stochastic component. The majority of exceptions to this rule, the relatively rare large fluctuations are generally short lived anomalous events.

2. Models

2.1. Data

The models we build are critically dependent on the data on which they are built. In this paper we analyze 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, and so we can collect data from 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. Also SNMP only provides aggregate link statistics, not the type of traffic using the link, nor its source or destination.

These limitations make 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 detailed analysis.

One point to note is that many past analyses of such data have been done in the "busy hour", but such analyses suffer from one major feature. Given a strong weekly cycle (which is the case here), and five minute traffic data (which is typical for SNMP measurements), one has only 12 sample points per week. To obtain enough data for a reasonably accurate analysis, one must average over many weeks. Over such time periods in the Internet non-stationarity effects may effect results [12]. This motivates using a model to describe and analyze the seasonal and trend components. Such a model is also useful in detecting anomalies that can occur at times outside the busy hour.

2.2. Traffic modeling

In this section, we describe the basic traffic model that we will use throughout the paper, based on standard techniques from time series analysis [13]. The most obvious characteristics of IP backbone traffic are the strong diurnal (daily) and weekly cycles, as well as long-term trends (for example see [14]).

Figure 1 shows the total traffic entering the network at a Point of Presence (PoP) over two consecutive weeks in May 2001, and illustrates these daily and weekly variations in the traffic. Also striking is the similarity between consecutive weeks of data. The obvious model for such traffic is a simple non-stationary model in which the traffic statistics (for instance, the mean and variance) vary over time in a regular and predictable way.

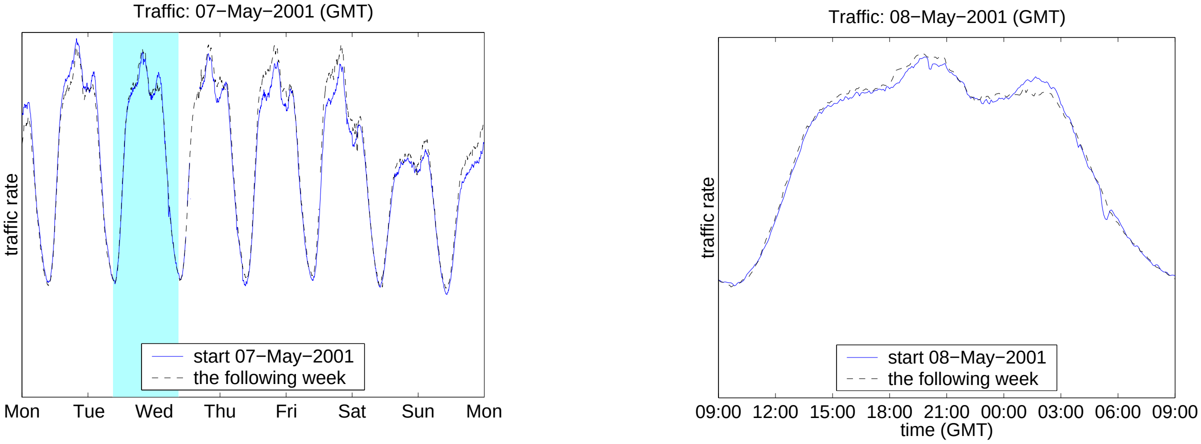

Figure 1. Total traffic into a region over two consecutive weeks. The solid line is the first weeks data (starting on May 7th) the dashed line shows the second weeks data. The second figure zooms in on the shaded region of the first.

We can quantify this intuitive view, by segmenting the traffic into a regular, predictable component, and a stochastic component. The most natural way to segment the two is using a generalization of the Norros model [11]. This model is ideal for backbone traffic modeling because it captures the effect of statistical multiplexing of many sources. The two components, a mean traffic rate m_t which can be some general function of time, and the stochastic component W_t are combined to give the total traffic rate by

$$x_t = m_t + \sqrt{a m_t} W_t, \tag{1}$$

where W_t is a stationary stochastic process with zero mean, and unit variance, and a is a parameter sometimes referred to as the *peakedness*. This type of model has been widely used – see for example [15–17] – and considerable evidence for the model's applicability to ISP backbone traffic is given in [12].

The mean traffic rate m_t is intended to capture the predictable components, namely the trend, and the weekly and daily cycles. The variation in the periodic components increases in proportion to the total volume of the traffic, and so it is natural [13] to take the mean to be,

$$m_t = T_t S_t, \tag{2}$$

where T_t denotes the trend, and S_t denotes the periodic *seasonal* component. We assume that we know the period T of the seasonal component, and that for the length of the data examined here the seasonal component can be well approximated as being periodic, i.e. $S_{t+nT} = S_t$, for all $n = 0, 1, 2, \ldots$

We discuss the parameter a in detail in Section 3, where we introduce it as a metric of traffic variability. In the original Norros model [11] the stochastic component W_t was Fractional Gaussian Noise (FGN), a simple Long-Range Dependent (LRD) process (see [1,2]). Here, our SNMP measurements do not allow a detailed characterization of W_t, so we allow any finite variance process.

This model satisfies a number of desirable properties. One important characteristic of the model is that when sets of traffic that obey the model are multiplexed they continue to obey the model. For instance, take N traffic streams x_i with constant mean m_i, peakedness a_i, and stochastic components which are independent realizations of the same Gaussian process. The mean of the new process is $m = \sum_{i=1}^N m_i$, and the peakedness (derived from the variance) is $a = \frac{1}{m} \sum_{i=1}^N a_i m_i$, which is a mean weighted average of the component peakednesses.

The units of a are unit-seconds [18], because the integral of the stochastic component W should be dimensionless. For instance, if the measurements are in kbps, then a is measured in kilobit-seconds. If the measurements are in packets per second, then a is measured in packet-seconds.

3. Metrics

One major problem with many of the previous commentaries on traffic stability is a lack of quantifiable definitions. In this section, we propose such a quantitative metric. A useful metric should capture the important characteristics of traffic in a parsimonious way, e.g., we might want a metric that captures the effect of the traffic variability on capacity planning perhaps by specifying how much over-provisioning is required to carry the specified traffic at a given Quality of Service (QoS). However, detailed capacity estimates are difficult, particular to the technology and network, and require measurements of traffic properties at finer scales than our SNMP measurements provide. Hence we limit ourselves to simple metrics that can be quickly calculated on large data sets. The properties of the peakedness enumerated above make it an ideal metric for our purposes, but in order to measure a we must first extract the non-stationary components from the data.

3.1. Estimating the Mean

In this section we introduce and use basic time series analysis to compute an estimate of the (non-stationary) mean in equation (2). The starting point is a Moving Average (MA) which is simply a convolution of the time series with a low-pass filter. We use only centered rectangular windows in this work. Thus the MA of width $2n+1$ applied to time series x_t is

$$\hat{T}_t = \frac{1}{2n+1} \sum_{i=-n}^{n} x_{t+i}. \tag{3}$$

If the filter has length greater than the period T of the seasonal component (1 week) then the MA acts to remove the periodic variations. Therefore $n = T/2$ yields an estimate of the trend, i.e., we estimate the trend using a MA with width ~1 week. Once estimated we can form a detrended data set by $y_t = x_t/\hat{T}_t$.

The standard method in time series analysis used to estimate the seasonal component exploits the periodicity by using a Seasonal Moving Average (SMA) where we take a MA over a series of data points separated by the period, so as to estimate the periodic component at those time points [13]. In our analysis, we perform our average over the whole data set, i.e.

$$\hat{S}_t = \frac{1}{N_t} \sum_{i=0}^{N_t-1} y_{t+iT}, \tag{4}$$

where $t \in [0, T)$, and N_t is the largest integer such that $t + (N_t - 1)T \leq N$, where N is the length of the data. We refer to this as the Seasonal Average (SA). We may estimate the mean by $\hat{m}_t = \hat{S}_t \hat{T}_t$.

3.2. Estimating peakedness

Once we understand the periodic and trend components of the traffic, the next thing to capture is the random variation around the mean. Most metrics of variation used in capacity planning do not account for the time-varying component, and so are limited to busy-hour analyses. In comparison, we have estimated $\hat{m}_t$ and so can use (1) to estimate the stochastic component, by $z_t = (x_t - \hat{m}_t)/\sqrt{\hat{m}_t}$. We can now measure the variability of the random component of the traffic using the variance of z_t. If we knew m_t exactly the variance would be the peakedness. Thus the metric defined here is an estimator for a, which we will denote $\hat{a}$, and refer to as the *empirical peakedness*.

We must also include in the estimation a correction for bias in the estimate. The correction arises for the same reason that the prefactor in the unbiased estimator of the variance of a set of data x_i with unknown mean is $\frac{1}{N-1}$ rather than $\frac{1}{N}$ as one might naively expect [13]. In our case, note that the computation of $\hat{a}$ can be rewritten as

$$\hat{a} = \frac{1}{T} \sum_{t=0}^{T-1} \mathrm{Var}_n \left(z_{t+nT} \right), \tag{5}$$

where T is the period of the seasonal component. When computing $\text{Var}_n\left(z_{t+nT}\right)$ we must use the prefactor for the unbiased sample variance, taking into account that each is based on N_t data points:

$$\text{Var}_n\left(z_{t+nT}\right) = \frac{1}{N_t - 1} \sum_{n=0}^{N_t-1} \left(z_{t+nT} - \bar{z}_t\right)^2, \tag{6}$$

where $\bar{z}_t = \frac{1}{N_t-1} \sum_{n=0}^{N_t} z_{t+nT}$. If $N_t = M$, a constant over the data set (the data length is an exact multiple M of the period of the data), then the above reverts to estimating the variance of z_t with a correction factor $\frac{M}{M-1}$. For example, with four weeks of data, and a weekly period, $M = 4$ so the correction factor is a not insignificant $4/3$.

We should note that not only is a ideal for our purposes, but even when the model (1) breaks down, for instance when there are outliers in the data, then parameter $\hat{a}$ is a useful and meaningful measurement (see [16,17] and Section 3.4). Furthermore, the metric may be easily adapted to deal with missing data, a feature we use below.

3.3. Simulation results

In the previous section we present an estimator called the empirical peakedness, but this estimator is not an unbiased estimator of peakedness, because the model (1) is more complicated than in typical time series cases. In this section, we present some simulation results to confirm that the metric above has only small bias on a sample data set similar to actual measurement data. We simulate a time series corresponding to one month's worth of 5 minute SNMP measurements, according to (1) and (2), where the stochastic component is FGN with $H = 0.5$, 0.75 and 0.95 (generated using the technique in [19]); the trend is exponential corresponding to a doubling every three months[1] (with a base rate of 26.6); and with the seasonal, or periodic component determined by a sinusoid plus a constant, i.e. $S_t = c + k \sin\left(2\pi \frac{t}{T}\right)$, where c and k are constants. The results in Figure 2 (a) are based on the product of two sinusoids, one with a period of 1 day, the other with a period of 1 week to simulate both weekly and daily cycles in the data. The parameters used in the presented simulation are $c_{\text{daily}} = 2$, $k_{\text{daily}} = 1$, $c_{\text{weekly}} = 1$, $k_{\text{weekly}} = 0.2$ (though we have tested results for a much wider range of parameters). In each of the simulations we also remove two blocks of 12 hours of data to simulate missing data.

Figure 2 (a) shows the results of estimates over a range of values of a based on 10 simulations each (with 95% confidence intervals shown). We can see that while there is a statistically significant bias in some results (those with $H = 0.5$ and 0.75) it is very small. For larger values of H, there is no significant bias. The missing data has little effect on the results (as long as not too much data is missing) and nor do the parameters of the sinusoid used to generate the seasonal component.

We have compared this metric with alternatives, such as the coefficient of variation, and peak to mean ratio, and found the empirical peakedness to be less biased, and less sensitive to form of the function m_t, making $\hat{a}$ a better measurement.

3.4. Effect of anomalies

Anomalous events include transient peaks/dips in the data caused by DoS attacks, flash crowds, rerouting of traffic, etc, that do not fit model (1). In Figure 2 (b) we show the effect of an outlier in the data – that is, a set of points not corresponding to the model in (1). The model used in the simulation has the same parameters as in the previous example, and the x-axis of the figure shows the size of the outlier (the outlier was chosen to effect three consecutive data points) with respect to those parameters. One can see that the outlier causes $\hat{a}$ to deviate from a (as one would expect). The deviation is approximately quadratic in the size of the event, and is linear in the length of the event, and the number of events in the data.

Figure 2 (b) demonstrates that the metric has the very desirable property that when the modeling assumptions that lie behind it are violated, the estimate responds in a smooth, predictable way. If, in contrast, the measurement varied wildly, then it would loose any meaning, because in almost all data

[1] This rate of increase is likely to be far more extreme than the rate of increase in real traffic [20,21], but otherwise the trend would be barely noticeable over a one month period.

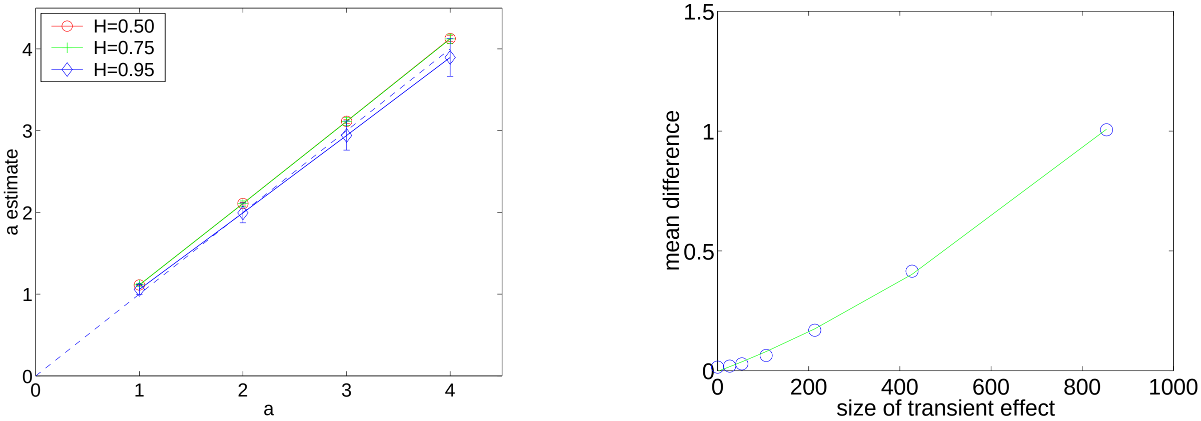

(a) Results of simulated measurements of $\hat{a}$.

(b) Measurements of the mean difference $\hat{a} - a$ when there is a transient outlier in the data. The circles show the empirical data. The line is a quadratic fit to the data.

Figure 2. Simulation results.

sets there are some outliers. The predictable nature of the variation allows us to draw meaning from the results even when outliers are present. In fact these outliers may be events of interest in themselves, and so it is not unreasonable for them to effect the measurement, so long as they do so in a controlled manner.

4. Measurements

This section presents measurements of $\hat{a}$. We shall look at approximately one month sequences of data, for which a purely periodic seasonal component was found to work well. At longer time scales a SMA is required as the seasonal component itself may change significantly, as the network topology, or the customer mix changes. For longer time scales, one should also ensure stationarity of a using tests such as in [12]. We present results from two periods: May 1st to 31st 2001 (an ordinary month), and September 5th to October 11th 2001. The dates in the Sept-Oct data set were chosen to cover several large events that we wish to investigate.

We measure $\hat{a}$ on the 5 minute traffic data from the OC48 links that form the majority of the backbone network at the time of study. Figure 3 (a-top) shows an example of 1 month of traffic data from a link, while Figure 3 (a-middle) shows the estimated mean m_t. Figure 3 (a-bottom) shows the value of z_t. Although this series appears to have some residual non-stationarity, the appearance could be the result of long-range correlations in the data [23]. In Figure 3 (b-top) we zoom in on the first week of the data shown in Figure 3. In Figure 3 (b-bottom) we show a very simple simulated version of the traffic. The measured value of the metric is $\hat{a} = 0.78$ Mbs for this data.

We have applied the measurements to of the order of 50 OC48 inter-city backbone links. Figure 4 shows the Cumulative Distribution Function (CDF) of the measured values of $\hat{a}$. Most values are in 0.5-3 Mbs, though 15%-25% of values are larger. The May data set also has many more values > 5 Mbs.

The dashed lines in Figure 4 show the measurements when anomalies in the data are removed from the data. There is a dramatic reduction in $\hat{a}$, in particular the larger values all but vanish even though only a small number of data points are removed from the data (on average 148 measurements from 8064). The reasons for this are discussed in Section 3.4: anomalous values can have a significant effect on the value of $\hat{a}$. A very small proportion of atypical data is contributing a large proportion of the variability. In detail, these points appear as spikes, or drops, and they can almost all be attributed to rerouting of traffic.

5. Meaning

The main value of metrics such as the average is that we feel we have some intuition about their meaning. At this point, the metric $\hat{a}$ is somewhat abstract. Here we shall give some meaning to the metric, through simulating a simple WDM model to see the effect a has on the decision to use dynamic

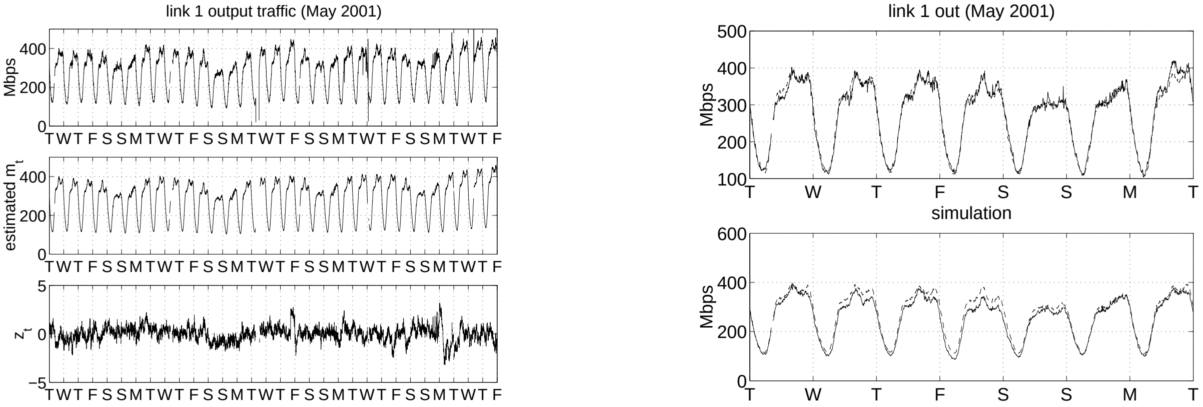

(a) Real measurements of $\hat{a}$. The data shown is that from an inter-city OC48 backbone link. The first plot shows the traffic, the second the estimate of m_t, and the third shows the transformed traffic z_t. For this data $\hat{a} = 0.78$.

(b) The first plot shows the first week of data from Figure 3. The second plot shows a simple synthetic data stream, based on (1). Note that in the plots, the dashed line shows the same $\hat{m}_t$.

Figure 3. Real Measurements of $\hat{a}$.

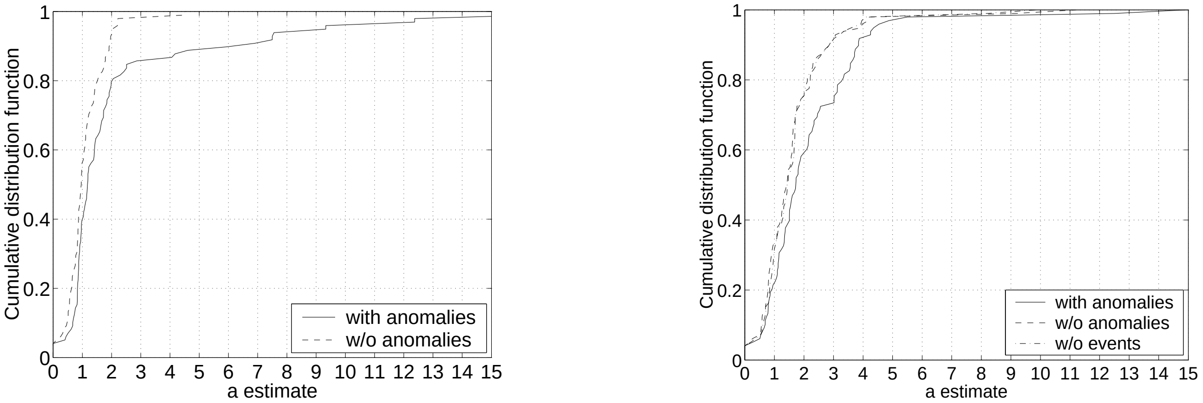

(a) May data.

(b) Sept-Oct data.

Figure 4. The CDF of $\hat{a}$ (solid), and of $\hat{a}$ once anomalous events have been isolated from the data.

reconfiguration in the network. In Section 5.1 the relationship between the peakedness and the need for network reconfigurability, and in Section 5.2 we look at data to understand the peakedness seen in AT&T's network over time.

5.1. Peakedness and the need for reconfigurability

We first approach the problem of giving meaning to a by placing it within a context: the network topology, link bandwidths, routing, etc. We simulate a simple Wavelength Division Multiplexing (WDM) access network based on the Metropolitan Area Network described in [24] where the topology, shown in Figure 5 (a), is feeder ring with multiple access nodes. Each access node is an IP router and an optical add-drop multiplexer (O-ADM) giving it individually allocated wavelengths that provide transport to the gateway node. The gateway node allocates wavelengths to the different access nodes, and so can be modeled as a simple multiplexer/demultiplexer as shown in Figure 5 (b).

Each of the N access nodes has dedicated bandwidth to the uplink allocated in wavelengths of capacity C. We assume negligible internal traffic on the access network, and model the external traffic from access node i as x_t^i using (1), with three additional assumptions

- Stationarity: We assume that the mean of the traffic is constant in time.
- Homogeneous: We assume that $m_i = m$ and $a_i = a$ for all i.
- Gaussianity: We assume the stochastic component is Gaussian.

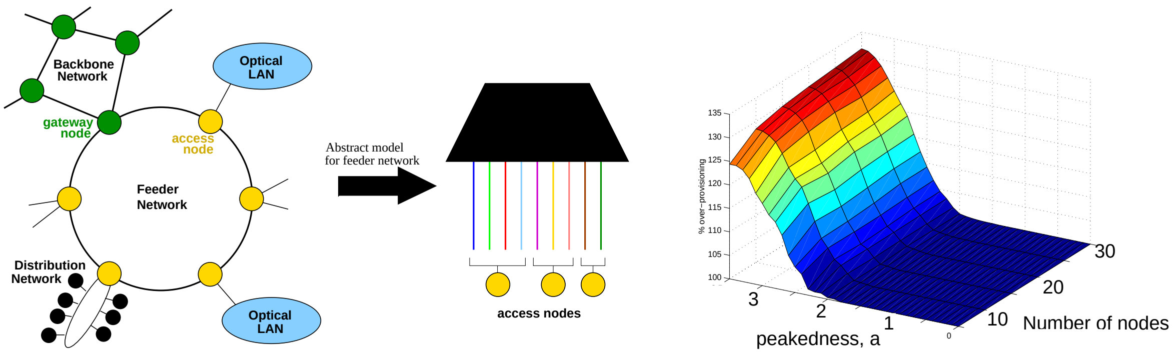

Figure 5. The simple Metropolitan Area Network model (a,b) and simulation results (c).

In the simplest scenario, the wavelengths are statically allocated according to traffic forecasts and provisioning rules. However, improved performance is expected if the wavelengths are dynamically assigned to different access nodes to follow fluctuations in the offered loads. The aim of this model is to assess the relative capacity gain of dynamic reconfiguration. As a simple first analysis, we allocate enough wavelengths to each access node to carry all of the offered traffic. Considering the above model, with traffic x_t^i to node i at time t, we find that the number of wavelengths required $y_t^i = \lceil x_t^i / C \rceil$, where $\lceil x \rceil$ denotes the smallest integer larger than x. The total number of wavelengths required in each case is

- **Static configuration:** $W = \sum_i \max_t y_t^i$.
- **Dynamic reconfiguration:** $W = \max_t \sum_i y_t^i$.

We simulate using 20 sample realizations of 4 weeks of 5 minute measurements, and show one illustrative result with traffic characteristics close to real data (bandwidth per wavelength $C = 155$ Mbps, average traffic per node is 350 Mbps, and a varies from 0 to 3.5 Mbs).

Figure 5 (c) shows a plot of the over-provisioning required by the static case to carry the same traffic as the dynamic case. The figure shows that there is a threshold in a below which there is no need for reconfiguration. Above the threshold, a statically configured system requires increasing amounts of over-provisioning. The threshold in a is largely insensitive to number of nodes N, but does depend on the mean rate m. The results show that the value of a has a clear impact on and network design. Given that very few of measured values of a lie above 3.0 (excluding anomalies) the case for reconfigurability is weak, though when we include network anomalies a may take large enough values to justify reconfigurability.

5.2. Causes of variability

Another way in which we can gain some understanding of the measurement $\hat{a}$ is to determine what factors effect it's value. Is it simply variation in the traffic over time, or do particular events influence its value? We know, for instance, that in the link data the value of $\hat{a}$ is increased by atypical events (see Figure 4). Such events could be flash crowds; DoS attacks; and self-propagating worms or viruses; rerouting due to link outages, or externally induced changes via BGP; or important natural and man-made events (for instance earthquakes, September 11th, holidays). In the May data, the majority of anomalous events occur as the result of rerouting of traffic, however there are events such as the May 28th Memorial day holiday may result in changes in traffic (a noticeable dip on this holiday). The Sept-Oct data was chosen because it covers some more dramatic events:

1. The attack on the WTC on September 11th: Figure 6 (a) shows three consecutive weeks from Sept-Oct for the New York region, and the time of the first crash. There is a clear drop in the traffic from New York at this time. The drop appeared in most other PoPs as well, suggesting that the drop was not caused by link outages (a fact supported by the fault data).

2. The Nimda Worm [25,26]: once infected, a host would send probes to infect new machines. There was a large enough volume of these probes to cause a DoS effect in some networks. The worm was first noticeable at 13:00 GMT on Sept. the 18th and its peak activity continued until the 19th.

These times are shaded on Figure 6 (a), but note that there is no discernible increase in traffic in New York (and elsewhere) during these times. The reason that this has not effected backbone traffic totals is that the probes are generally fairly small, and there are not really that many compared with the total backbone traffic. It takes a *really big* event to affect the backbone.

3. The final event of interest is a BGP problem associated with a malformed AS-path that propagated through the Internet from 19:40 on the 7th to 1:40 on the 8th of Oct. (see NANOG mailing list for details). The event is clearly visible in BGP traffic, which surged by about a factor of 10. The malformed AS-path may have caused some customer routers to crash, and thereby caused the increased BGP traffic, and we might therefore expect to see a decrease in data traffic. Figure 6 (b) shows the weeks in question. There is a small decrease in traffic on the day in question, but it not clearly caused by the BGP event. Without more detailed traffic data, it is hard to confirm the true cause, but it does seem to be evidence that BGP routing instability can affect traffic.

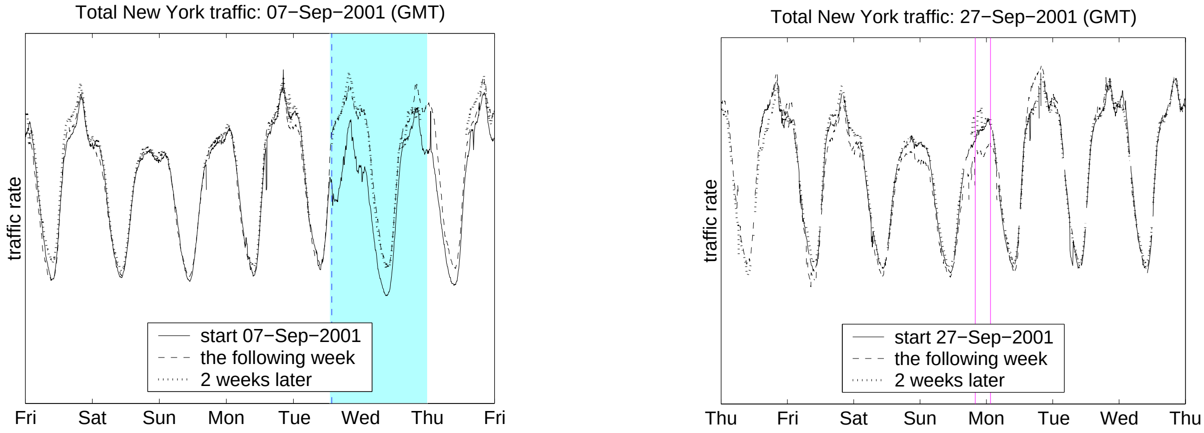

(a) The solid line shows the drop on the 11th of Sept. The shaded region shows the period of the Nimda attack.

(b) The vertical lines show the period of increased BGP activity.

Figure 6. Total traffic into New York over three consecutive weeks (for two different months).

To understand what sort of effect these unusual events have on our measure of variability $\hat{a}$, we also compute the empirical peakedness of the traffic, excluding these events. Figure 4 (b) shows the CDF of the measured $\hat{a}$, with all data, without simple anomalies, and without all anomalies, including those above. The latter two are almost indistinguishable, because the peakedness's response to anomalies is quadratic in their 'size', but only linear in the duration. Hence long lived events such as the WTC 20% traffic drop don't impact peakedness as much as short rerouting events which may double the traffic on a link. Together with the results of 5.1, we see that optical reconfigurability helps most to handle traffic changes arising from sudden events such as network failures, which may be external to the system under administrative control.

6. Conclusion

Understanding the characteristics of backbone traffic is crucial for both the engineering and design of large networks. This paper presents a novel technique for measuring the variability of backbone Internet traffic, and investigates this techniques using SNMP measurements from a Tier-1 ISP backbone and through simulations. A key insight of this paper is that large deviations from traffic predictions (not due to routing changes) are rare. Most normal variation has a in the range 0.5-3.0 Mbs for 5 minute SNMP measurements. This value of a appears to represent relatively stable traffic. However, we note that a can be larger even when we exclude obvious transient events. At the very least this provides a realistic set of parameter values for simulations of backbone traffic.

Our original motivation for this work was to look at the benefits of building IP backbones on top of a re-configurable optical network. Though the diurnal variations in traffic are significant, these are tightly coupled across North American, and so do not allow temporal sharing of capacity. Furthermore, the stability of the stochastic component suggests that the case for a re-configurable optical network layer based solely on traffic variations is weak – certainly some claims of variability have been dramatically exaggerated. However, it may still make sense to use a such a network to deal with traffic load changes resulting from IP layer re-routing due to failures [22], or due to global differences in the diurnal cycle.

Acknowledgments: We would like to thank Fred True, Joel Gottlieb, and Tim Griffin for their work in collecting, and managing the data used here, and Bob Doverspike, and Panita Pongpaibool for stimulating discussion of optical reconfiguration.

REFERENCES

1. 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*, vol. 2, pp. 1–15, Feb 1994.
2. 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.
3. V. Paxson, "Empirically-derived analytic models of wide-area TCP connections," *IEEE/ACM Transactions on Networking*, vol. 2, no. 4, pp. 316–336, 1994.
4. J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun, "On the nonstationarity of Internet traffic," in *ACM SIGMETRICS 2001*, (Cambridge), 2001.
5. B. S. Arnaud, "Current optical network designs may be flawed," *Optical Networks Magazine*, vol. 2, March/April 2001.
6. J. Mooney, "Gaining the edge in flexible metro service provisioning," *Lightwave*, February 2001.
7. A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving traffic demands for operational IP networks," *IEEE/ACM Trans. on Networking*, pp. 265–279, 2001.
8. J. Y. Wei, "IP over WDM network traffic engineering approaches," in *OFC*, 2002.
9. D. Awduche and Y. Rekhter, "Multiprotocol lambda switching: combining MPLS traffic engineering control with optical crossconnects," *IEEE Communications Magazine*, vol. 39, pp. 111–116, March 2001.
10. D. Banerjee, "Wavelength routed optical networks: linear formulation, resource budgeting tradeoffs and a reconfiguration study," in *IEEE INFOCOM'97*, pp. 269–276, 1997.
11. I. Norros, "A storage model with self-similar input," *Queueing Systems*, vol. 16, pp. 387–396, 1994.
12. M. Roughan and J. Gottlieb, "Large-scale measurement and modeling of backbone internet traffic," in *ITCOM, Boston, MA, USA*, 2002.
13. P. Brockwell and R. Davis, *Time Series: Theory and Methods*. Springer-Verlag, 1987.
14. J. W. Roberts, "Traffic theory and the Internet," *IEEE Communications Magazine*, Jan 2001.
15. W.-C. Lau, A. Erramilli, J.L.Wang, and W. Willinger, "Self-similar traffic parameter estimation: a semi-parametric periodogram-based algorithm," in *GLOBECOM'95*, 1995.
16. D. L. Jagerman, B. Melamed, and W. Willinger, "Stochastic modeling of traffic processes," in *Frontiers in Queueing: Models, Methods and Problems* (J. Dshalalow, ed.), CRC Press, Boca Raton, 1995.
17. D. L. Jagerman, "Burstiness descriptors of traffic streams: Indices of dispersion and peakedness," in *Proceedings of the Conference on Information Sciences and Systems*, (Princeton, NJ), pp. 24–28, 1994.
18. A. Erramilli. private communications.
19. V. Paxson, "Fast, approximate synthesis of fractional Gaussian noise for generating self-similar network traffic," *Computer Communications Review*, vol. 27, pp. 5–18, Oct 1997.
20. A. M. Odlyzko, "Internet growth: Myth and reality, use and abuse," *J. Computer Resource Management*, vol. 102, pp. 23–27, 2001.
21. K. G. Coffman and A. M. Odlyzko, "Internet growth: Is there a "Moore's law" for data traffic?," in *Handbook of Massive Data Sets* (J. Abello, P. M. Pardalos, and M. G. C. Resende, eds.), Kluwer, 2001.
22. P. Pongpaibool, R. Doverspike, M. Roughan, and J. Gottlieb, "Handling IP traffic surges via optical layer reconfiguration," in *OFC*, 2002.
23. J. Beran, *Statistics for Long-Memory Processes*. Chapman and Hall, New York, 1994.
24. J.Yates, G. Hjálmtýsson, and A.Greenberg, "Reconfiguration in IP over WDM access networks," *OFC*,2000.
25. R. China, "NIMDA worm/virus report - final." Incidents.org, October 2001.
26. CERT Advisory CA-2001-26. Available at `http://www.cert.org/`.
27. J. Crowie, A. Ogielski, B. Premore, and Y. Yuan, "Global routing instabilities during Code Red II and Nimda worm propagation." NANOG 23, 2001.