

On the Distribution of Traffic Volumes in the Internet and its Implications

Abstract—Getting good statistical models of traffic on network links is a well-known, often-studied problem. A lot of attention has been given to correlation patterns and flow duration. The distribution of the amount of traffic per unit time is an equally important but less studied problem. We study a large number of traffic traces from many different networks including academic, commercial and residential networks using state-of-the-art statistical techniques. We show that the log-normal distribution is a better fit than the Gaussian distribution commonly claimed in the literature. We also investigate a second heavy-tailed distribution (the Weibull) and show that its performance is better than Gaussian but worse than log-normal. We examine anomalous traces which are a poor fit for all distributions tried and show that this is often due to traffic outages or links that hit maximum capacity.

We demonstrate the utility of the log-normal distribution in two contexts: predicting the proportion of time traffic will exceed a given level (for service level agreement or link capacity estimation) and predicting 95th percentile pricing. We also show the log-normal distribution is a better predictor than Gaussian or Weibull distributions.

Index Terms—Traffic modelling, network planning, bandwidth provisioning, traffic billing

I. INTRODUCTION

Internet traffic characterisation is an important problem for network researchers and vendors. The subject has a long history. Early works [1], [2] discovered that the correlation structure of traffic exhibits self-similarity and that the durations of individual flows of packets exhibit heavy-tails [3]. These works were later challenged and refined (see Section VI for a summary). By comparison the distribution of the amount of traffic seen on a link in a given time period has seen comparatively less research interest. This is surprising as this quantity can be extremely useful in network planning.

In this paper we use a rigorous statistical approach to fitting a statistical distribution to the amount of traffic within a given time period. Formally, we choose some timescale T and let X_i be the amount of traffic seen in the time period $[iT, (i+1)T)$. We investigate the distribution of the random variable X over a wide range of values of T . We show that the distribution of the variable has considerable implications for network planning; for assessing how often a link is over capacity and in particular for service level agreements (SLAs), and for traffic pricing, particularly using the 95th percentile scheme [4].

Previous authors have claimed that X has a normal (or Gaussian) distribution [5]–[7]. Others claim X is Gaussian plus a tail associated with bursts [8], [9]. A variable X has a log-normal distribution if its logarithm is normally distributed $\ln(X) \sim N(\mu, \sigma^2)$ where $\mu \in \mathbb{R}$ is the mean and $\sigma > 0$ is the standard deviation of the distribution. We use a well-

established statistical methodology [10] to show that a log-normal distribution is a better fit than Gaussian or Weibull¹ for the vast majority of traces. This holds over a wide range of timescales T (from 5 msec to 5 sec). This paper is the most comprehensive investigation of this phenomenon the authors know about. We study a large number of publicly available traces from a diverse set of locations (including commercial, academic and residential networks) with different link speeds and spanning the last 15 years.

The structure of the paper is as follows. In Section II we describe the datasets used. In Section III we describe our best-practice procedure for fitting traffic and demonstrate that log-normal is the best fit distribution for our traces under a variety of circumstances. We examine those few traces that do not follow this distribution and find it occurs when a link spends considerable time either having an outage or completely at maximum capacity. In Section IV we demonstrate that the log-normal distribution is the most useful for estimating how often a link is over capacity. In Section V we show that the log-normal distribution provides good estimates when looking at 95th percentile pricing. In Section VI we give related work. Finally, Section VII gives our conclusions.

II. NETWORK TRAFFIC TRACES

A key contribution of our work stems from the spatial and temporal diversity of the studied traces. The dataset spans a period of 15 years and comprises 229 traces.

CAIDA traces. We have used 27 CAIDA traces captured at an Internet data collection monitor which is located at an Equinix data centre in Chicago [11]. The data centre is connected to a backbone link of a Tier 1 ISP. The monitor records an hour-long traces four times a year, usually from 13:00 to 14:00 UTC. Each trace contains billions of IPv4 packets, the headers of which are anonymised. The average captured data rate is 2.5 Gbps. At the time of capturing, the monitored link had a capacity of 10 Gbps. Traces were captured between 2013 and 2016.

MAWI traces. The MAWI archive [12] consists of a collection of Internet traffic traces, captured within the WIDE backbone network that connects Japanese universities and research institutions to the Internet. Each trace consists of IP level traffic observed daily from 14:00 to 14:15 at a vantage point within WIDE. Traces include anonymised IP and MAC headers, along with an *ntpd* timestamp [12]. We have looked at 107 traces (each one being 15 minutes long). Traces were

¹A variable X has a Weibull distribution with parameters $k > 0$ (known as shape) and $\lambda > 0$ (known as scale) if its probability density function follows $f(x) = \frac{k}{\lambda} (\frac{x}{\lambda})^{k-1} \exp(-(x/\lambda)^k)$ when $x \geq 0$ and is 0 otherwise.

captured between 2014 and 2018. On average, each trace consists of 70 million packets; the average captured data rate is 422 Mbps. The monitored link had a capacity of 1 Gbps.

Twente University traces. We used 40 traffic traces captured at five different locations (8 traces from each location). Traces are diverse in terms of the link rates, types of users and captured time [13]. Each trace is 15 minutes long. The first location is a residential network with a 300 Mbps link, which connects 2000 students (each one having a 100 Mbps access link); traces were captured in July 2002. The second location is a research institute network with a 1 Gbps link which connects 200 researchers (each one having a 100 Mbps access link); traces were captured between May and August 2003. The third location is at a large college with a 1 Gbps link which connects 1000 employees (each one having a 100 Mbps access link); traces were captured between February and July 2004. The fourth location is an ADSL access network with a 1 Gbps ADSL link used by hundreds of users (each one having a 256 Kbps to 8 Mbps access link); traces were captured between February and July 2004. The fifth location is an educational organisation with a 100 Mbps link connecting 135 students and employees (each one having a 100 Mbps access link); traces were captured between May and June 2007.

Waikato University VIII traces. The Waikato dataset consists of traffic traces captured by the WAND group at the University of Waikato, New Zealand [14]. The capture point is at the link interconnecting the University with the Internet. All of the traces were captured using software that was specifically developed for the Waikato capture point and a DAG 3 series hardware capture card. All IP addresses within the traces are anonymised. In our study, we have used 30 traces captured between April 2011 and November 2011.

Auckland University IX traces. The Auckland dataset consists of traffic traces captured by the WAND group at the University of Waikato [15]. The traces were collected at the University of Auckland, New Zealand. The capture point is at the link interconnecting the University with the Internet. All IP addresses within the traces are anonymised. In our study, we have used 25 traces captured in 2009.

III. FITTING A STATISTICAL DISTRIBUTION TO INTERNET TRAFFIC DATA

In this section we present an extensive statistical analysis applied to the datasets described in the previous section. The aim is to discover which statistical distribution best fits the traces. In contrast to the existing research (see Section VI), we are basing our analysis on the framework proposed by Clauset et al. [10], a comprehensive statistical framework developed specifically for testing power-law behaviour in empirical data². The framework combines maximum-likelihood fitting methods with goodness-of-fit tests based on the Kolmogorov–Smirnov statistic and likelihood ratios. The method reliably tests whether the power-law distribution is the best model for

a specific dataset, or, if not, whether an alternative statistical distribution (e.g., log-normal, exponential, Weibull) is. The framework performs the tests described above as follows: (1) the parameters of the power-law model are estimated for a given dataset; (2) the goodness-of-fit between the data and the power-law is calculated, under the hypothesis that the power-law is the best fit to the provided traffic samples. If the resulting p -value is greater than 0.1 the hypothesis is accepted (i.e. the power law is a plausible fit to the given data), otherwise the hypothesis is rejected; (3) alternative distributions are tested against the power-law as a fit to the data by employing a likelihood ratio test.

For the vast majority of the traces examined, the hypothesis was rejected; i.e. the power-law distribution was not a good fit. Consequently, we investigate alternative distributions by performing the likelihood ratio (LLR) test (following Clauset’s methodology), as follows:

$$\mathfrak{R}, p = \text{fit.distributionCompare}(\text{powerlaw}, \text{alternative})$$

where \mathfrak{R} is the normalised LLR³ between the power-law and alternative distributions and p is the significance value for this test. \mathfrak{R} is positive if the power-law distribution is a better fit for the data, and negative if the alternative distribution is a better fit for the data. A p -value less than 0.1 means that the value of \mathfrak{R} can be trusted to make a conclusion that one candidate distribution (power-law or alternative, depending on the sign of \mathfrak{R}) is a good fit for the data. In contrast, a p -value greater than 0.1 means that there is nothing to be concluded from the likelihood ratio test.

A. Fitting the log-normal distribution to Internet traffic data

Figure 1 shows the results of the LLR test for all 229 traces with log-normal, exponential and Weibull distribution as the alternative to power-law. For this test we have aggregated traffic at a timescale $T = 100$ msec. The points marked with a circle are the ones with $p > 0.1$. It is clear that the log-normal distribution (black line in Figure 1) is the best fit for the studied traces; i.e. $\mathfrak{R} < 0$ and $p < 0.1$ for most traces when the alternative distribution (to the power-law which is almost always rejected) is the log-normal one⁴. The log-normal distribution is not the best fit for 1 out of 27 CAIDA traces, 2 out of 30 Waikato traces, 1 out of 25 Auckland traces, 5 out of 40 Twente traces and 9 out of 107 MAWI traces. We examined these traces in more detail and discuss them in Section III-B.

For the vast majority of traces the power-law distribution is favoured over the exponential one (i.e. $\mathfrak{R} > 0$), as shown in Figure 1. Thus, the exponential distribution cannot be considered as a good model for our traffic traces. On the other hand, the Weibull appears to be a better fit over the power-law distribution; however, when compared to the log-normal distribution, it still performs poorly (i.e. $\mathfrak{R} > 0$ or $\mathfrak{R} < 0$ but $p > 0.1$) for a substantial amount of traces.

³ \mathfrak{R} is calculated as $\mathcal{R}/(\sigma\sqrt{n})$, where \mathcal{R} is the log likelihood ratio [10].

⁴For clarity, in Figures 1(e) and 2(e) we only plot traces 60 – 107. For traces 1 – 59, \mathfrak{R} is less than 0 and the respective p -value is less than 0.1; i.e. the alternative distribution is the best fit for the respective trace

²We have used the source code discussed in [16].

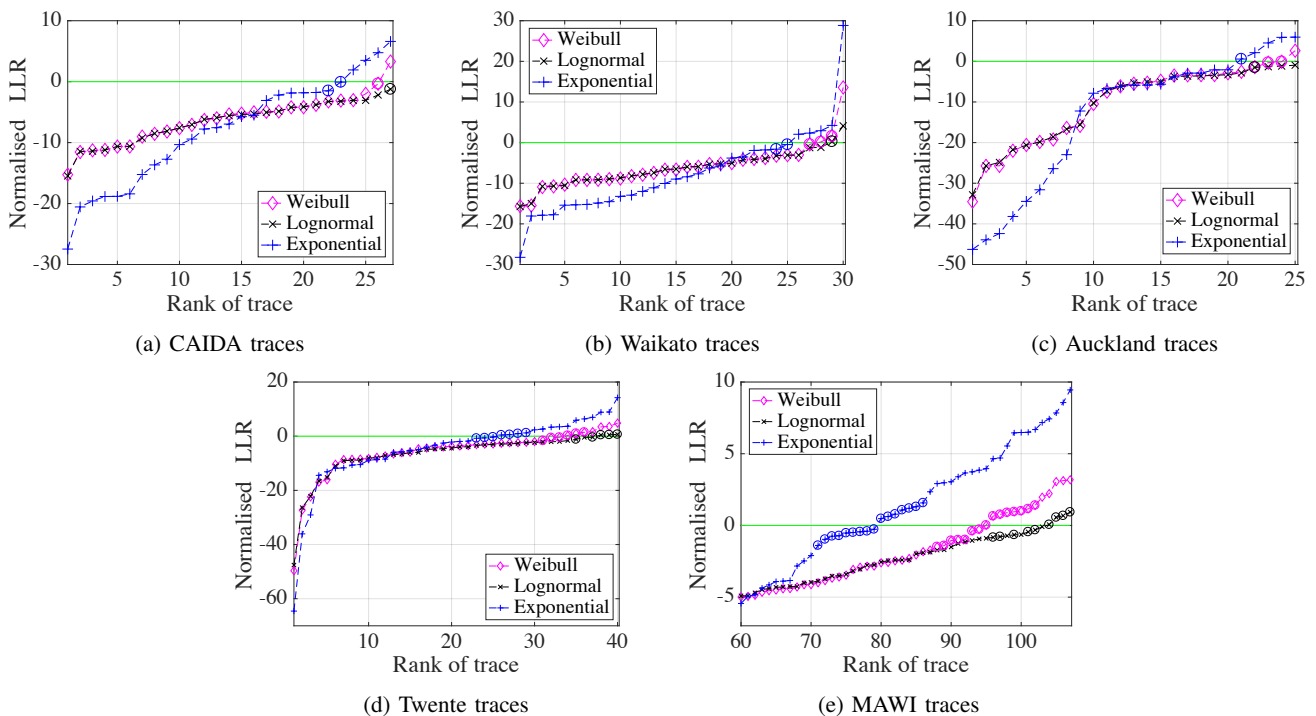


Fig. 1: Normalised Log-Likelihood Ratio (LLR) test results for all studied traces and candidate distributions. Aggregation timescale T is 100 msec. Circled points in the plot are the ones with p -value greater than 0.1; i.e. likelihood test is inconclusive with respect to fitting any of the candidate distributions to the traffic data.

Identifying the log-normal distribution as the best fit for the vast majority of traffic traces at $T = 100$ msec is very encouraging. This specific traffic aggregation timescale has been commonly studied in the literature [17], [18]. Next we investigate what the best model is for a range of aggregation timescales. The results are shown in Figure 2. As reflected by the \mathfrak{R} and p -values, the log-normal distribution is the best fit for the vast majority of captured traces at all examined timescales (5 msec to 5 sec)⁵. This is a strong result suggesting the generality of our observations. The good log-normal fit at time scales as small as 5 msec is important for practical applications of the log-normal model.

We also examined Q-Q plots for a large number of traces⁶. The log-normal distribution appeared to be a better fit than other tested distributions and no deviations from the expected pattern were observed in the body or tail of the distribution.

B. Anomalous traces

As mentioned in Section III-A, there are a small number of traces for which the log-normal distribution is not a good fit (none of the other examined distributions is, either). Figure 3(a) shows the PDF plot for one of the 8 anomalous

⁵Note that it is possible that the network traffic may not follow a log-normal distribution at very fine or coarse aggregation granularities.

⁶Due to lack of space, Q-Q plots are not included as we would have to present plots for each trace, separately.

MAWI traces. Figure 3(b) shows the PDF for another MAWI trace for which the log-normal distribution is a good fit. It is obvious from Figure 3(a) that the link was either severely underutilised (see large spike on the left part of the plot area) or fully utilised (see smaller spike at the right part of the plot area) for higher data rates. All traces for which the log-normal distribution was not a good fit exhibited similar behaviour and (aggregated) traffic patterns. On the contrary, we did not observe any such behaviour for the majority of traces for which the log-normal distribution was the best fit. A likely explanation for the anomalous traces is that those traces contain either periods of over-capacity (traffic is at 100% of link capacity) or periods where the link is broken (no traffic).

C. Fitting the log-normal and Gaussian distributions using the correlation coefficient test

The linear correlation coefficient test has been widely used to assess the fit of a distribution to empirical data. To reinforce the results of Section III-A, we employ the linear correlation coefficient assuming that the log-normal distribution is the best fit (as we showed in Section III). We compare the results of this test for both the log-normal and Gaussian distributions. We use the linear correlation coefficient as defined in [19]:

$$\gamma = \frac{\sum_{i=1}^n (S_{(i)} - \hat{\mu})(x_i - \hat{x})}{\sqrt{\sum_{i=1}^n (S_{(i)} - \hat{\mu})^2 \cdot \sum_{i=1}^n (x_i - \hat{x})^2}} \quad (1)$$

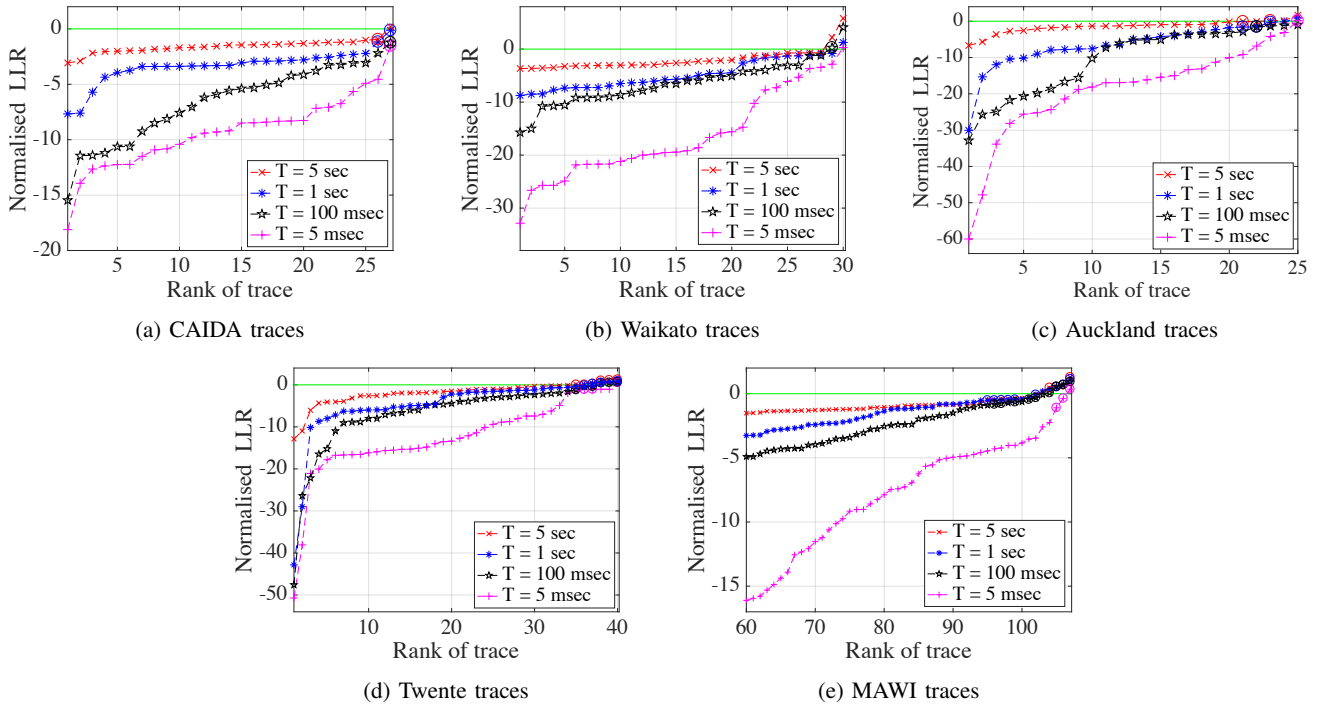


Fig. 2: Normalised Log-Likelihood Ratio (LLR) test results for all studied traces and log-normal distribution. Aggregation timescales are 5 sec, 1 sec, 100 msec and 5 msec. Circled points in the plot are the ones with p -value greater than 0.1, i.e. likelihood test is inconclusive with respect to fitting the log-normal distribution to the traffic data.

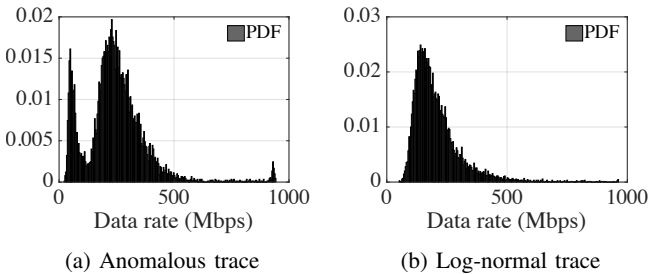


Fig. 3: PDF of an anomalous and non-anomalous trace.

where $S_{(i)}$ is the observed sample i , and $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n S_{(i)}$ is the samples' mean value. x_i is sample i from the reference distribution (log-normal in our case), which can be calculated from the inverse CDF of the reference random variable $x_i = F^{-1}\left(\frac{i}{n+1}\right)$ and $\hat{x} = \frac{1}{n} \sum_{i=1}^n x_i$ is the respective mean value. The value of the correlation coefficient can vary between $-1 \leq \gamma \leq 1$, with a 1, 0 and -1 indicating perfect correlation, no correlation and perfect anti-correlation, respectively. Strong goodness-of-fit (GOF) is assumed to exist when the value of γ is greater than 0.95 [17].

We measure the linear correlation coefficient for all datasets at four different aggregation timescales (ranging from 5 msec to 5 sec) and plot the results in Figures 4(a) to 4(e) for the log-

normal distribution and Figures 4(f) to 4(j) for the Gaussian distribution. Traces are ordered by the value of γ for the given timescale. It can be clearly seen that $\gamma > 0.95$ for most traces when employing the test for the log-normal distribution, but this is not the case for the Gaussian distribution. γ is larger for smaller aggregation timescales indicating that the log-normal distribution is an even better fit as the aggregation gets finer. For very small values of T , i.e. lower than 1 msec, data samples exhibit binary behaviour, where either a packet is transmitted or not during each examined time frame [18]. We have examined γ for very short (and large) aggregation timescales, and can confirm the absence of a model describing the data (for brevity, we have omitted the relevant figures).

Next, we calculate v_γ (the variation of γ) for each dataset. v_γ gives an indication of the stability of γ for each dataset, for all timescales tested. This metric is defined as:

$$v_\gamma = \sqrt{\text{var}(\gamma_{T_1}, \gamma_{T_2}, \gamma_{T_3}, \gamma_{T_4})} \quad (2)$$

where $T_1 = 5$ sec, $T_2 = 1$ sec, $T_3 = 100$ msec and $T_4 = 5$ msec. Figures 4(k) to 4(o) show the results for each dataset with the traces ranked by v_γ . For log-normal model, v_γ is very small (below 0.045) for all traces, therefore we can conclude that γ is almost constant for all studied aggregation timescales. While v_γ is higher for Gaussian model. Furthermore, the error bars in Figures 4(p) to 4(t) represent the standard deviation of the correlation coefficient at different timescales (see x-axis).

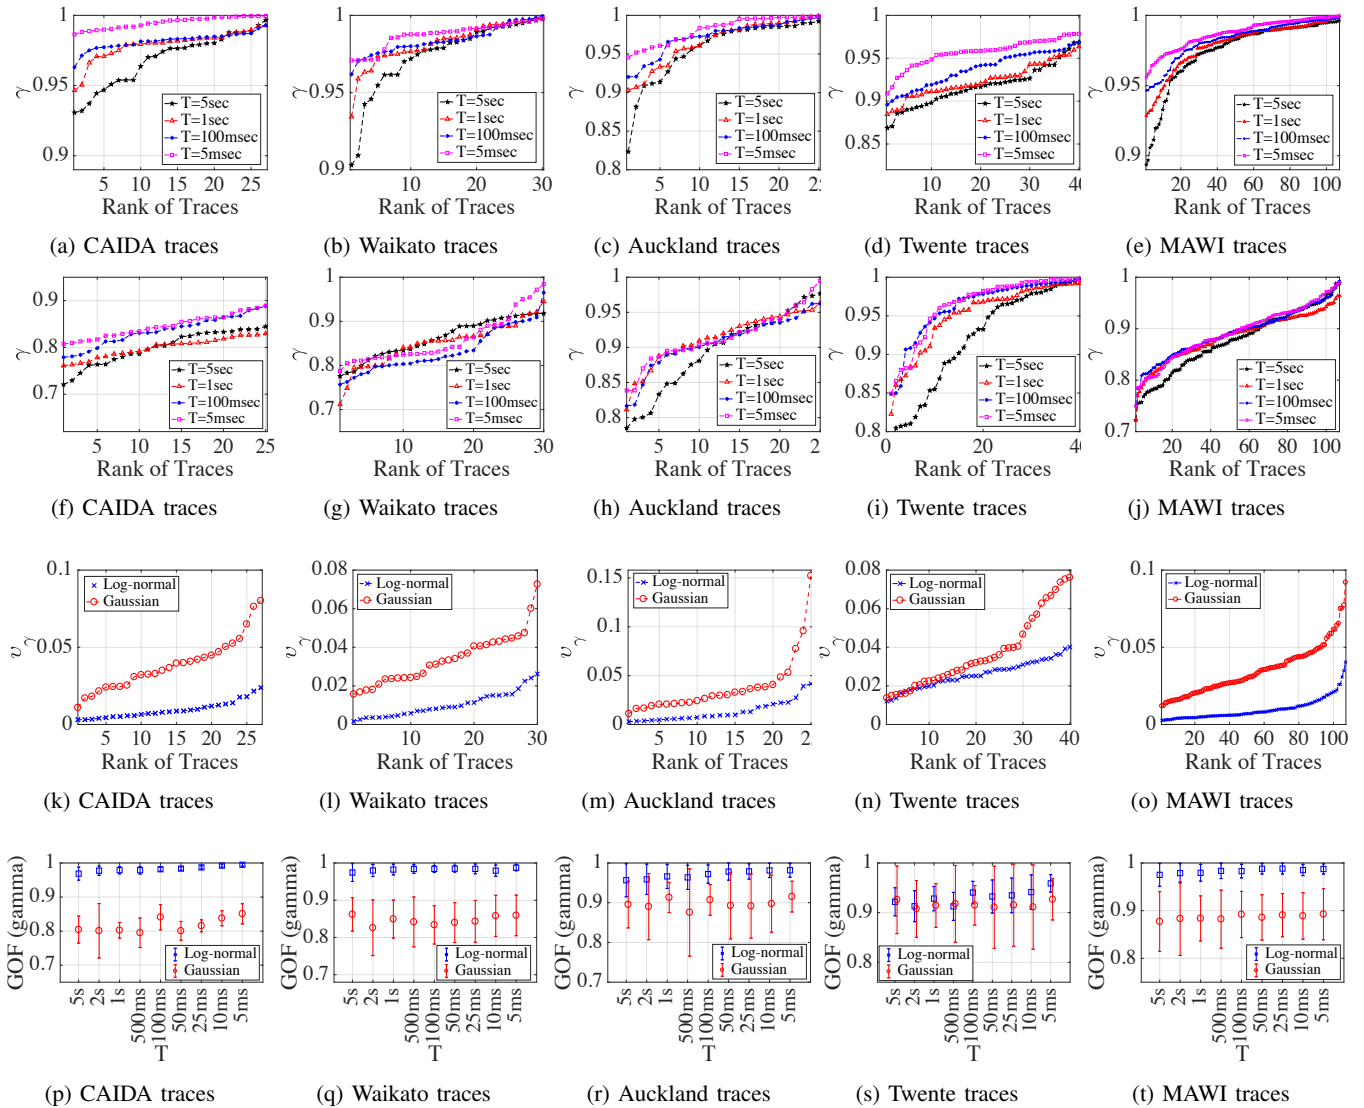


Fig. 4: Correlation coefficient test results for all studied traces and different timescales.

This again shows that for log-normal model γ is larger than 0.95 (at different T values) for most CAIDA and MAWI traces, while it is larger than 0.9 for all other datasets. This is not the case with the Gaussian model, where most γ values are less than 0.9.

Overall, the correlation coefficient test reinforces the results extracted in Section III-A, providing strong evidence that the log-normal distribution is the best fit for all studied traces. Superior performance of our model can also be seen from comparison of our results for correlation coefficient with those in [20] where the Gaussian model was used.

IV. BANDWIDTH PROVISIONING

It has been previously suggested that network link provisioning could be based on fitted traffic models instead of relying on straightforward empirical rules [20]. In this way, over-

or under-provisioning can be mitigated or eliminated even in the presence of strong traffic fluctuations. Such approaches rely on having a statistical model that accurately describes the network traffic. This is therefore an excellent area for applying our findings on fitting the log-normal distribution to Internet traffic data. In the literature, the following inequality (the authors call it the “link transparency formula”) has been used for bandwidth provisioning [18]:

$$P(A(T) \geq CT) \leq \varepsilon. \quad (3)$$

In words, this inequality states that the probability that the captured traffic $A(T)$ over a specific aggregation timescale T is larger than the link capacity has to be smaller than the value of a performance criterion ε . The value of ε is chosen carefully by the network provider in order to meet a specific SLA [20]. Likewise, the value of the aggregation

time T should be sufficiently small so that the fluctuations in the traffic can be modelled as well, taking into account the buffering capabilities of network switching devices⁷.

We compare bandwidth provisioning using Meent's approximation formula [20] (assuming Gaussian) and using a log-normal traffic model.

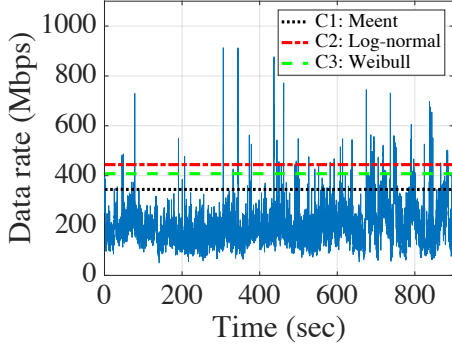


Fig. 5: Data rate of a MAWI trace ($T = 100$ msec and $\varepsilon = 0.01$). The horizontal lines represent the calculated link capacity based on different models.

A. Bandwidth provisioning using Meent's formula

To find the minimum required link capacity, Meent et al. [20] proposed a bandwidth provisioning approach that is based on the assumption that the traffic follows a Gaussian distribution. Meent's dimensioning formula is defined as follows [20]:

$$C1 = \mu + \frac{1}{T} \sqrt{-2\log(\varepsilon) \cdot v(T)} \quad (4)$$

where μ is the average value of the traffic, $v(T)$ is the variance at timescale T and ε is the performance criterion. The link capacity is obtained by adding a safety margin value

$$\text{Safety margin} = \sqrt{-2\log(\varepsilon)} \cdot \sqrt{\frac{v(T)}{T^2}}$$

to the average of the captured traffic (see Equation 4). This safety margin value depends on ε and the ratio $\sqrt{v(T)/T^2}$. As the value of ε decreases the safety margin increases. For example, when the value of ε decreases from 10^{-2} to 10^{-4} , then value of the safety margin increases by 40%. This is different from conventional link dimensioning methods, where the safety margin is fixed to be 30% above the average of the presented traffic [20], [21]. Traffic tails are represented using the Chernoff bound, as follows:

$$P(A(T) \geq CT) \leq e^{-SCT} E[e^{SA(T)}]. \quad (5)$$

Here $E[e^{SA(T)}]$ is the moment generation function (MGF) of the captured traffic $A(T)$.

⁷Large traffic fluctuations at very short aggregation timescales are smoothed by the presence of buffers at network routers and switches.

B. Bandwidth provisioning based on the log-normal model

Here we investigate whether we could achieve more reliable bandwidth provisioning by adopting the log-normal traffic model. We calculate the mean and variance from the captured trace and generate the respective log-normal model. Then, we use the CDF function (F) to solve the link transparency formula shown in Equation 3. Hence, F is defined as $F(C) = P(A(T)/T < C)$, which can be solved to find C , as follows:

$$C2 = F^{-1}(1 - \varepsilon). \quad (6)$$

C. Comparison of bandwidth provisioning approaches

In this section, we compare the bandwidth provisioning approaches described above. The performance indicator is the empirical value of the performance criterion, which is denoted by $\hat{\varepsilon}$ and defined as follows:

$$\hat{\varepsilon} = \frac{\#\{A_i | A_i \geq CT\}}{n}, i \in 1 \dots n. \quad (7)$$

In words, this empirical value is the percentage of all the data samples of the captured traffic which are measured larger than the estimated link capacity. Ideally, $\hat{\varepsilon}$ would be equal to the target value of the performance criterion ε . The difference between $\hat{\varepsilon}$ and ε is due to the fact that the chosen traffic model is not accurately describing the real network traffic. A simple example of the described comparison approach is illustrated in Figure 5, in which we plot the captured data rate for a MAWI trace ($T = 100$ msec)⁸. The calculated capacity values from each approach when the target ε is 0.01 are $C1 = 344.8$ Mbps and $C2 = 444.3$ Mbps (represented by the horizontal lines in Figure 5). The empirical value can be calculated by using Equation 7, which gives $\hat{\varepsilon}_1 = 0.042$ and $\hat{\varepsilon}_2 = 0.012$. Obviously, with the first approach the network operator would not be able to meet the target $\varepsilon = 0.01$, while with the second approach the empirical value is close to the target.

We next compare results of bandwidth provisioning calculations based on the (a) Meent's formula, (b) Weibull model and (c) proposed log-normal model. Figure 6(a)-(d) shows the average of the empirical value ($\text{avg}(\hat{\varepsilon})$) for all traces in each dataset at $T = 0.1$ sec, $T = 0.5$ sec and $T = 1$ sec. The value of T is chosen to be sufficiently small so that the fluctuations in the traffic can be modelled as well. Each model is tested for four different values of the performance criterion: $\varepsilon = 0.5$, $\varepsilon = 0.1$, $\varepsilon = 0.05$ and $\varepsilon = 0.01$. In Figure 6(a)-(d) we clearly see that the log-normal model is able to satisfy the required performance criterion ε at different aggregation time-scales for all datasets. In contrast, Meent's formula failed to allocate sufficient bandwidth, which results in missing the target performance criterion ε for all datasets and target performance values, as depicted in Figure 6(i)-(l) (see horizontal red line). The Weibull distribution performs better comparing to Meent's formula, but bandwidth provisioning using the log-normal model is far superior, as can be seen from Figures 6(a)-(d) and 6(e)-(h).

⁸Note that in all subsequent figures we have also included results for a Weibull model to get insights about bandwidth provisioning using a heavy-tailed distribution.

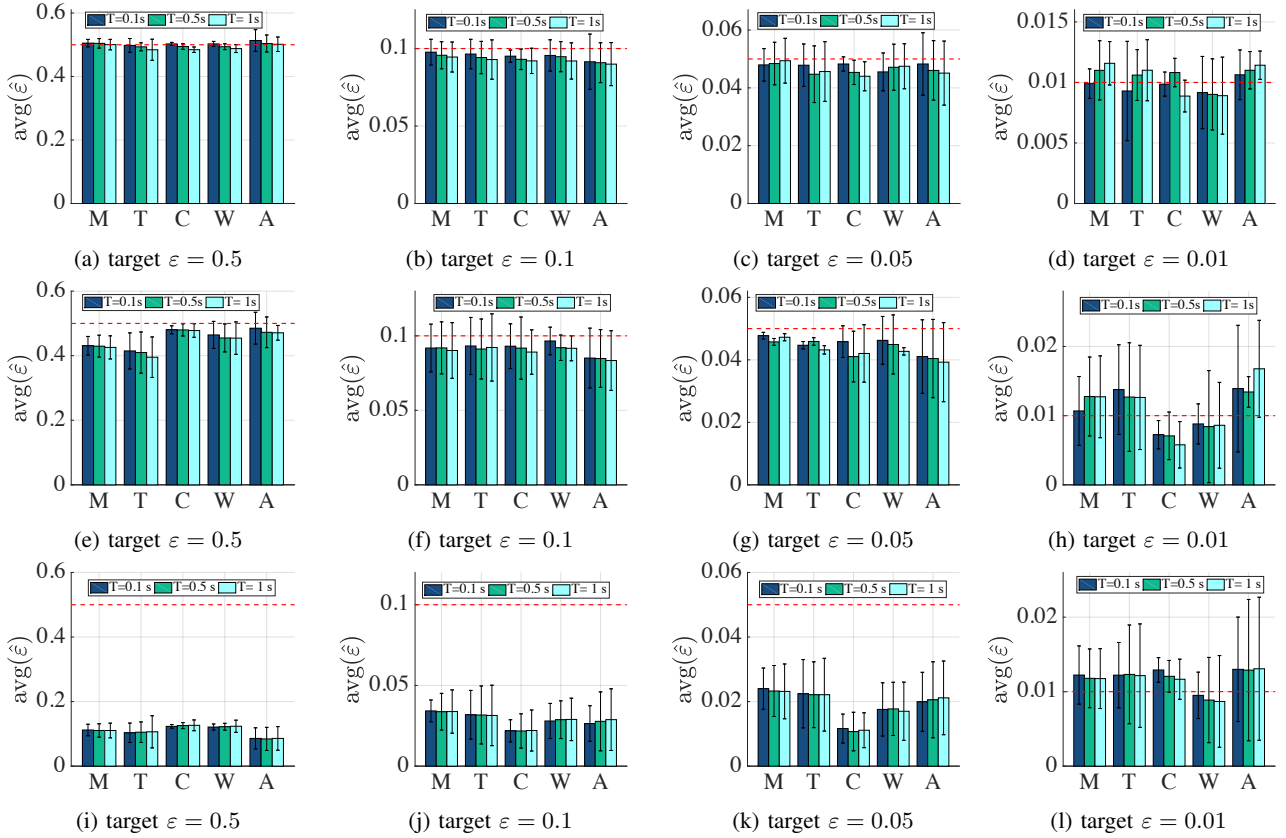


Fig. 6: Link dimensioning based on (a-d) log-normal model, (e-h) Weibull model and (i-l) Meent's formula: $\text{avg}(\hat{\varepsilon})$ for different datasets (M: MAWI, T: Twente, C: CAIDA, W: Waikato, A: Auckland), aggregation timescales (100 msec, 500 msec and 1 s), and target values of ε (0.5, 0.1, 0.05 and 0.01). Error bars represent $\text{stderr} |\varepsilon - \hat{\varepsilon}|$.

V. 95TH PERCENTILE PRICING SCHEME BASED ON LOG-NORMAL MODEL

Traffic billing is typically based on the 95th percentile method [22]. Traffic volume is measured at border network devices (typically aggregated at time intervals of 5 minutes) and bills are calculated according to the 95-percentile of the distribution of measured volumes; i.e. network operators calculate bills by disregarding occasional traffic spikes. Forecasting future bills, which is important for ISPs and clients, can be done using a model of the traffic calculated through previously sampled traffic. In this section, we apply our findings on Internet traffic modelling in predicting the cost of traffic according to the 95th percentile method. For each network trace we calculate the actual 95th percentile of the traffic volume. The majority of the studied traffic traces were 15-minute long but operators typically use measurements traffic volumes for much longer periods, therefore we scale down the calculation of the 95th percentile by dividing each trace (900 seconds) into 90 groups (10 seconds length each). The authors appreciate that by using 15-minute rather than day long traces we omit any study of diurnal effects in the distribution. We note that the sum of several log-normal distributions is itself very accurately represented by a log-normal distribution [23].

Hypothetically, therefore, if 96 consecutive 15-minute traces fit a log-normal distribution (with different parameters for each) then the resulting 24 hour trace is also likely to be a good fit to a log-normal. We also note that the distributions tested were on a level playing field in that they would all be affected equally by the shorter duration of the data sets.

We calculate the 95th percentile for the observed traffic. We then fit a Gaussian, Weibull and log-normal distribution to each trace (for $T = 100$ msec) and calculate the 95th percentile of the fitted distribution. We plot the actual 95th percentile against the three predictions in Figure 7 with a red reference line to show where perfect predictions would be located. It is clear that the log-normal model provides much more accurate predictions of the 95th percentile than the Gaussian model. As with the bandwidth dimensioning case discussed in Section IV, the Weibull is better than the Gaussian model but worse than the proposed log-normal model.

We employ the normalised root mean squared error (NRMSE) as a goodness of fit to the results in Figure 7. NRMSE measures the differences between values predicted by a hypothetical model and the actual values. In other words, it measures the quality of the fit between the actual data and the predicted model. Table I shows the NRMSE for all datasets

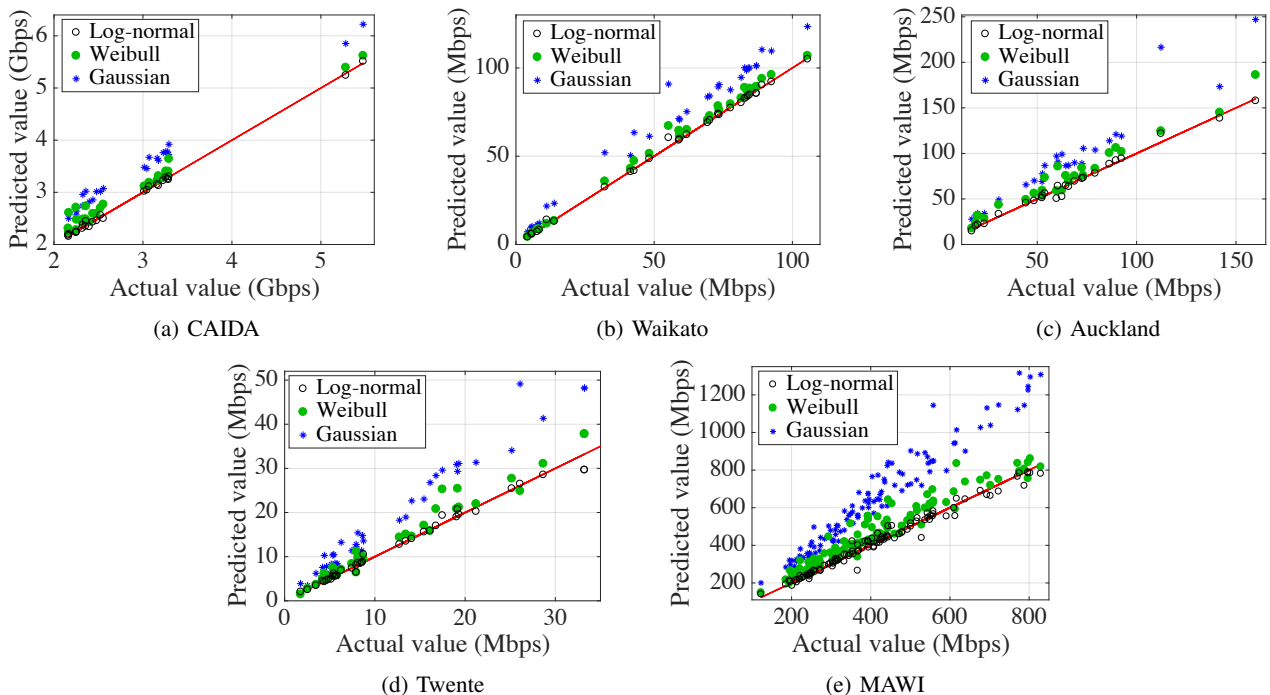


Fig. 7: 95th percentile values (actual vs predicted rates) based on log-normal, Weibull and Gaussian models. An ideal model would result in points in the plot area that fall exactly on the red line.

TABLE I: Goodness of fit (GOF) using normalised root mean squared error (NRMSE)

Model/Dataset	CAIDA	Waikato	Auckland	Twente	MAWI
Log-normal	0.0399	0.0401	0.1058	0.0979	0.1528
Weibull	0.2410	0.1148	0.2984	0.2123	0.4145
Gaussian	0.5544	0.4193	0.6866	0.5741	0.9828

and the three considered models. It is clear that the lowest NRMSE value is for the log-normal model, which is the best model compared to the Gaussian and Weibull ones.

VI. RELATED WORK

Reliable traffic modelling is important for network planning, deployment and management; e.g. for traffic billing and network dimensioning. Historically, network traffic has been widely assumed to follow a Gaussian distribution. In [5], [7], the authors studied network traces and verified that the Gaussianity assumption was valid (according to simple goodness-of-fit tests they used) at two different timescales. In [24], the authors studied traffic traces during busy hours over a relatively long period of time and also found that the Gaussian distribution is a good fit for the captured traffic. Schmidt et al. [8] found that the degree of Gaussianity is affected by short and intensive activities of single network hosts that create sudden traffic bursts. All the above mentioned works agreed on the Gaussian or ‘fairly Gaussian’ traffic at different levels of aggregations in terms of timescale and

number of users. The authors in [19], [25] examined the levels of aggregation required to observe Gaussianity in the modelled traffic, and concluded that this can be disturbed by traffic bursts. The work in [9], [26] reinforces the argument above, by showing existence of large traffic spikes at short timescales which result in high values in the tail. Compared to existing literature, our findings are based on a modern, principled statistical methodology, and traffic traces that are spatially and temporally diverse. We have tested several hypothesised distributions and not just Gaussianity.

An early work drawing attention to the presence of heavy tails in Internet file sizes (not traffic) is that of Crovella and Bestavros [2]. Deciding whether Internet flows could be heavy-tailed became important as this implies significant departures from Gaussianity. The authors in [27] provided robust evidence for the presence of various kinds of scaling, and in particular, heavy-tailed sources and long range dependence in a large dataset of traffic spanning a duration of 14 years.

Understanding the traffic characteristics and how these evolve is crucial for ISPs for network planning and link dimensioning. Operators typically over-provision their networks. A common approach to do so is to calculate the average bandwidth utilisation [6] and add a safety margin. As a rule of thumb, this margin is defined as a percentage of the calculated bandwidth utilisation [21]. Meent et al. [20] proposed a new bandwidth provisioning formula, which calculates the minimum bandwidth that guarantees the required performance,

according to an underlying SLA. This approach relies on the statistical parameters of the captured traffic and a performance parameter. The underlying fundamental assumption for this to work is that the traffic the network operator sees follows a Gaussian distribution. This approach has been used again in [18].

The 95th percentile method is used widely for network traffic billing. Dimitropoulos et al. [22] have found that the computed 95th percentile is significantly affected by traffic aggregation parameters. However, in their approach they do not assume any underlying model of the traffic; instead, they base their study on specific captured traces. Stanojevic et al. [4] proposed the use of Shapley value for computing the contribution of each flow to the 95th percentile price of interconnect links. Works [28]–[31] propose calculating the 95th percentile using experimental approaches. Xu et al. [32] assume that network traffic follows a Gaussian distribution “through reasonable aggregation” and propose a cost efficient data centre selection approach based on the 95th percentile.

VII. CONCLUSION

The distribution of traffic on Internet links is an important problem that has received relatively little attention. We use a well-known, state-of-the-art statistical framework to investigate the problem using a large corpus of traces. The traces cover several network settings including home user access links, tier 1 backbone links and campus to Internet links. The traces are from times from 2002 to 2018 and are from a number of different countries. We investigated the distribution of the amount of traffic observed on a link in a given (small) aggregation period which we varied from 5 msec to 5 sec. The hypotheses compared were that the traffic volume was heavy-tailed, that the traffic was log-normal and that the traffic was normal (Gaussian). The vast majority of traces fitted the log-normal assumption best and this remained true all timescales tried. Where no distribution tested was a good fit this could be attributed either to the link being saturated (at capacity) for a large part of the observation or exhibiting signs of link-failure (no or very low traffic for part of the observation).

We investigate the impact of the distribution on two sample traffic engineering problems. Firstly, we looked at predicting the proportion of time a link will exceed a given capacity. This could be useful for provisioning links or for predicting when SLA violation is likely to occur. Secondly, we looked at predicting the 95th percentile transit bill that ISP might be given. For both of these problems the log-normal distribution gave a more accurate result than a heavy-tailed distribution or a Gaussian distribution. We conclude that the log-normal distribution is a good (best) fit for traffic volume on a normally functioning internet links in a variety of settings and over a variety of timescales, and further argue that this assumption can make a large difference to statistically predicted outcomes for applied network engineering problems.

REFERENCES

- [1] P. Pruthi and A. Erramilli, “Heavy-tailed on/off source behavior and self-similar traffic,” in *Proc. of ICC*, 1995.
- [2] M. E. Crovella and A. Bestavros, “Self-similarity in world wide web traffic: evidence and possible causes,” *IEEE/ACM ToN*, 1997.
- [3] P. Loiseau, P. Goncalves, G. Dewaele, P. Borgnat, P. Abry, and P. V. B. Primet, “Investigating self-similarity and heavy-tailed distributions on a large-scale experimental facility,” *IEEE/ACM ToN*, 2010.
- [4] R. Stanojevic and et. al., “On economic heavy hitters: Shapley value analysis of 95th-percentile pricing,” in *Proc. of ACM IMC*, 2010.
- [5] R. V. D. Meent, M. Mandjes, and A. Pras, “Gaussian traffic everywhere?” in *Proc. of IEEE ICC*, 2006.
- [6] R. d. O. Schmidt, H. van den Berg, and A. Pras, “Measurement-based network link dimensioning,” in *Proc. of IFIP/IEEE*, 2015.
- [7] R. d. O. Schmidt, R. Sadre, and A. Pras, “Gaussian traffic revisited,” in *Proc. of IFIP Networking*, 2013.
- [8] R. d. O. Schmidt, R. Sadre, N. Melnikov, J. Schnwlder, and A. Pras, “Linking network usage patterns to traffic gaussianity fit,” in *Proc. of IFIP Networking*, 2014.
- [9] X. Yang, “Designing traffic profiles for bursty Internet traffic,” in *Proc. of IEEE GLOBECOM*, 2002.
- [10] A. Clauset, C. S. Rohilla, and M. Newman, “Power-law distributions in empirical data,” *arXiv:0706.1062v2*, 2009.
- [11] “The caida uscd anonymized internet traces,” 2016. [Online]. Available: http://www.caida.org/data/passive/passive_dataset.xml
- [12] “Mawi archive,” 2018. [Online]. Available: <http://mawi.wide.ad.jp/>
- [13] R. R. R. Barbosa, R. Sadre, A. Pras, and R. van de Meent, “Simpleweb/university of twente traffic traces data repository,” <http://eprints.eemcs.utwente.nl/17829/>, Tech. Rep., 2010.
- [14] “Wits: Waikato internet traffic storage,” 2013. [Online]. Available: <https://wand.net.nz/wits/waikato/8/>
- [15] “Wits: Auckland x,” 2009. [Online]. Available: <https://wand.net.nz/wits/auck/10/>
- [16] J. Alstott, E. Bullmore, and D. Plenz, “powerlaw: a python package for analysis of heavy-tailed distributions,” *arXiv:1305.0215*, 2014.
- [17] M. Mandjes and R. van de Meent, “Resource dimensioning through buffer sampling,” *IEEE/ACM Transactions on Networking*, 2009.
- [18] R. d. O. Schmidt, R. Sadre, A. Sperotto, H. van den Berg, and A. Pras, “Impact of packet sampling on link dimensioning,” *IEEE Transactions on Network and Service Management*, 2015.
- [19] J. Kilpi and I. Norros, “Testing the gaussian approximation of aggregate traffic,” in *Proc. of SIGCOMM*, 2002.
- [20] A. Pras, L. Nieuwenhuis, R. van de Meent, and M. Mandjes, “Dimensioning network links: a new look at equivalent bandwidth,” *IEEE Network*, 2009.
- [21] “Best practices in core network capacity planning,” online, accessed July 2018. [Online]. Available: https://www.cisco.com/c/en/us/products/collateral/routers/wan-automation-engine/white_paper_c11-728551.pdf
- [22] X. Dimitropoulos, P. Hurley, A. Kind, and M. P. Stoeklin, “On the 95-Percentile Billing Method,” in *Proc. of PAM*, 2009.
- [23] R. Mitchell, “Permanence of the log-normal distribution,” *J. Optical Society of America*, 1968.
- [24] J. L. García-Dorado, J. A. Hernández, J. Aracil, J. E. López de Vergara, and S. Lopez-Buedo, “Characterization of the busy-hour traffic of IP networks based on their intrinsic features,” *Computer Networks*, 2011.
- [25] A. B. Downey, “Evidence for Long-tailed Distributions in the Internet,” in *Proc. of ACM SIGCOMM Workshop on Internet Measurement*, 2001.
- [26] H. Abrahamsson, B. Ahlgren, P. Lindvall, J. Nieminen, and P. Tholin, “Traffic characteristics on 1gbit/s access aggregation links,” in *Proc. of IEEE ICC*, 2017.
- [27] R. Fontugne and et. al., “Scaling in internet traffic: A 14 year and 3 day longitudinal study, with multiscale analyses and random projections,” *IEEE/ACM Transactions on Networking*, 2017.
- [28] L. Golubchik and et. al., “To send or not to send: Reducing the cost of data transmission,” in *Proc. of IEEE INFOCOM*, 2013.
- [29] N. Laoutaris, M. Sirivianos, X. Yang, and P. Rodriguez, “Inter-datacenter bulk transfers with netstitcher,” in *Proc. of ACM SIGCOMM*, 2011.
- [30] I. Castro, R. Stanojevic, and S. Gorinsky, “Using Tuangou to Reduce IP Transit Costs,” *IEEE/ACM Transactions on Networking*, 2014.
- [31] H. Xu and B. Li, “Joint request mapping and response routing for geo-distributed cloud services,” in *Proc. of IEEE INFOCOM*, 2013.
- [32] —, “Cost efficient datacenter selection for cloud services,” in *Proc. of IEEE ICC*, 2012.