

Characterising A Grid Site's Traffic

Tiejun Ma
Department of Computing
Imperial College London
London, UK
tma@doc.ic.ac.uk

Yehia El-khatib
Computing Department
Lancaster University
Lancaster, UK
yehia@comp.lancs.ac.uk

Michael Mackay
Computing Department
Lancaster University
Lancaster, UK
m.mackay@comp.lancs.ac.uk

Christopher Edwards
Computing Department
Lancaster University
Lancaster, UK
ce@comp.lancs.ac.uk

ABSTRACT

Grid computing has been widely adopted for intensive high performance computing. Since grid resources are distributed over complex large-scale infrastructures, understanding grid site data traffic behaviour is important for efficient resource utilisation, performance optimisation, and the design of future grid sites as well as traffic-aware grid applications. In this paper, we study and analyse the traffic generated at a grid site in the Large Hadron Collider (LHC) Computing Grid (LCG). We find that most of the generated traffic is TCP-based and that a small set of grid applications generate significant amounts of the data. Upon analysing the different traffic metrics, we also find that the traffic exhibits long-range dependence and self-similarity. We also investigate packet-level metrics such as throughput, packet rate, round trip time (RTT) and packet loss. Our study establishes that these metrics can be well represented by Gaussian mixture models. The findings we present in this paper will enable accurate grid site traffic monitoring and potentially on-the-fly traffic modelling and prediction. It will also lead to a better understanding of grid site's traffic behaviour and contribute to more efficient grid site planning, traffic management, data transmission protocol optimisation, and data-aware grid application design.

Categories and Subject Descriptors

C.4 [Performance, Network Measurement]: Traffic Modelling, Performance Measurement

General Terms

Network Performance, Grid Computing

Keywords

Traffic Modelling, Network Performance, Grid Computing

1. INTRODUCTION

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DIDC10 Chicago, Illinois - June 22, 2010

Copyright 2010 ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

Over the last few years, grid computing has been widely adopted within the research community for scientific computing. This has made it possible to deploy a number of very large-scale experiments that were deemed infeasible until only recently. The grid computing paradigm offers seamless access to a dynamic set of heterogeneous computational resources across different organisations, providing substantial computing and storage capabilities. Within grid environments, computational jobs are submitted to be run on any suitable resources, and the data is submitted to be transparently stored across geographically distributed nodes. The performance of the network infrastructure therefore plays an important role in the effectiveness of grid systems and it becomes important to characterise how data transfers at such a scale affects the underlying network infrastructure.

The structure and design aims of grid sites are often quite different compared with those of Web clusters and, more recently, data centres. Consequently, the protocol and application stacks deployed in each also differ. Protocols such as Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), etc. are widely adopted in Web-based clusters and data centres, while GRAM [2] for resource managing, GridFTP [5] for data transferring and Berkeley Database Information Index (BDII) [1] for resource discovery are often deployed within grid environments to enable resource sharing. Currently, most grid infrastructures focus on how to share computational and storage resources, how these resources are used, and how to maximise the performance of these resources. Less attention is given to other aspects such as the performance and behaviour of the network that interconnects such large-scale distributed resources, which could present a bottleneck in the performance of the whole system. This leaves many questions unanswered, such as:

- What, if any, are the distinguishable characteristics of the traffic generation behaviour of grid sites?
- Is the network connecting the grid sites used efficiently?
- Does the data traffic generated by grid sites display long-range dependence and/or self-similarity like Web traffic [15, 19]?
- What are the underlying traffic volumes, packet inter-arrival times (packet rate) and packet loss that define grid site traffic?
- Which applications will have a significant impact or generate the most packets on a grid site?

Answering these questions should give us a better understanding of grid sites network behaviour and consequently allow network

infrastructure designers and managers to optimise their design, implementation and runtime system management to better suit grid applications. To the best of our knowledge, very little is known about the traffic generation characteristics of grid sites.

In general, data traffic analysis work focuses on analysing examples of real data traffic to extract its characteristics. Specifically, this includes mapping aspects of the traffic flow to known mathematical models and identifying certain characteristics (such as self-similarity, etc.). Such studies give us a better understanding of application network behaviour and operational requirements, which in turn allows more informed improvements to the fields of network design, protocol design, security, traffic engineering, and pricing.

In this paper, we present a statistical evaluation of the characteristics of the traffic generation behaviour of a real-world grid site, with production-level services deployed at the University of Innsbruck (UIBK). This grid site is a part of the Large Hadron Collider (LHC) Computing Grid (LCG) infrastructure, which is used for performing LHC scientific computing tasks. We captured all data traffic generated at this site over the period of seven days. The grid site is operating as normal days during the data collection period, thus we regard the data collection duration is a “typical” seven days of this grid site. For this grid site, we first evaluate the data and connection behaviour of some widely used grid applications, such as GridFTP and BDII. Then, we analyse the properties of the grid traffic and use our results to determine if self-similar behaviour exists for grid data traffic. In addition, we evaluate the packet-level network characteristics, such as throughput, packet size and packet rate. We then develop analytical traffic models using prevalent statistical distributions.

The rest of this paper is organised as follows. Section 2 discusses related work. Section 3 describes node and network setup of the site we monitored to abstract operational details. Section 4 introduces the findings of the statistical analysis we carried out on the gathered traffic traces. Implications of our work and future work are discussed in section 5, while section 6 offers some concluding comments.

2. RELATED WORK

The statistical evaluation and analysis of data traffic is an important aspect of modern networking. As network throughput and load gradually increases along with the prevalence of applications requiring some level of service from the network, as does the requirement to explore the way that these applications utilise the network. It is therefore becoming critical to understand the traffic sent over the network and the statistical analysis and modelling of applications offers an ideal solution for this [16].

To date, traffic analysis has tended to focus on the fields of traditional network protocols, such as TCP [17] and underlying technologies such as Ethernet [15]. In the following, we have categorised the related work into three category groups.

The first group of research discussed here relates to various workload related data traffic, such as Web traffic [11, 24, 13], Variable-Bit-Rate video traffic [9], ethernet traffic [32], enterprise traffic [26]. These works mainly focus on modelling and evaluating performance related issues of the network systems in terms of the infrastructure, network transmission protocol and user behaviour. Empirical probability distribution fitting for relevant metrics is often adopted to derive empirical models of the system, *c.f.* [8].

For example, Ersoz *et al.* [13] analyses a clustered three-tier data centre based on a testbed. This paper adopts empirical distribution fitting based evaluation for the arrival rate and inter-arrival time distribution of the requests to individual server nodes, the data traffic between tiers, and the average size of messages exchanged

between tiers. The authors find that the request inter-arrival rates follow log-normal distribution, and self-similarity exists when the data centre is heavily loaded. The message sizes can be presented by the log-normal distribution and service times follows Pareto distribution. Compared with [13], our work is focusing on the traffic generation behaviour of a grid site. We, however, collect a large amount of traffic data from a real production level grid site, which is much more elaborate. As already discovered in [28], general empirical distribution fitting tests are biased against large or “messy” data sets [18]. Also a single distribution would be difficult to give good fitting results in the real-world. Thus, in our paper we adopt the discrepancy statistical analysis method introduced in [30] to reduce the level of intricacy of the collected data. Furthermore, we employ Gaussian Mixture Models for more accurate traffic pattern approximation.

The second group of analysis investigates self-similarity of traffic, such as [29, 11, 32, 14]. The main focus of this group of work is whether the observed traffic exhibits self-similar characteristics and how to build mathematical models to represent this. For example, [11] examines the dependence structure of Web traffic. The paper shows that Web traffic exhibits behaviour consistent with self-similar traffic models. The paper also demonstrates that the self-similarity in network traffic can be explained in terms of file system characteristics and user behaviour. In our paper, we also evaluate the self-similarity property of the traffic generated by this grid site.

The third group is focussed on analysing TCP traffic, congestion avoidance behaviour and the impact of a random loss process on a TCP connection, such as [25, 22, 6]. These works focus on how to build stochastic models of the TCP throughput, which is represented using packet loss, round trip delay, etc. Among these works, Kumar [22] proposes a stochastic model to study the throughput performance of various versions of Transport Control Protocol (TCP) in a local network. TCP throughput is computed as the reward rate in a certain Markov renewal reward process. The paper shows that for large packet-loss probabilities, TCP-Reno performs no better, or worse, than TCP-Tahoe but TCP-NewReno is a considerable improvement over TCP-Tahoe. In addition, Padhye *et al.* [25] develop an analytic characterisation of the steady state throughput of TCP-Reno, as a function of loss rate and round trip time for the TCP flow. It captures the essence of TCP’s congestion avoidance behaviour. The model is evaluated using a number of TCP connections between the United States and Europe.

Despite this plethora of works on various types of traffic and workload analysis, very little effort has been made to deliberately study grid sites and the network. The reasoning behind this lack of understanding of grid traffic properties is twofold. First, grid application developers tend to assume that all grid traffic is made up of *elephants and mice*, *i.e.* huge bulks of data and (relatively) tiny control signals, respectively. This assumption has been evident in grid literature for a long time, but is invalid as illustrated in [12] and in the results of our study. Second, there is seemingly a common lack of interest within the community of Grid users about network issues, with the premise that over-provisioning the network will deal with any shortcomings, if any. This is problematic as many Grid applications do not run on dedicated or over-provisioned infrastructures. System designers and researchers who wish to work on new technologies for the Grid have to rely on Grid simulation packages to evaluate their prototypes. Unfortunately, current Grid simulators do not offer any realistic models of Grid traffic, which follows from the lack of Grid traffic studies we outlined. Most Grid simulators assume that Grid traffic is merely Internet traffic which is not an accurate assumption, as will be discussed later in this pa-

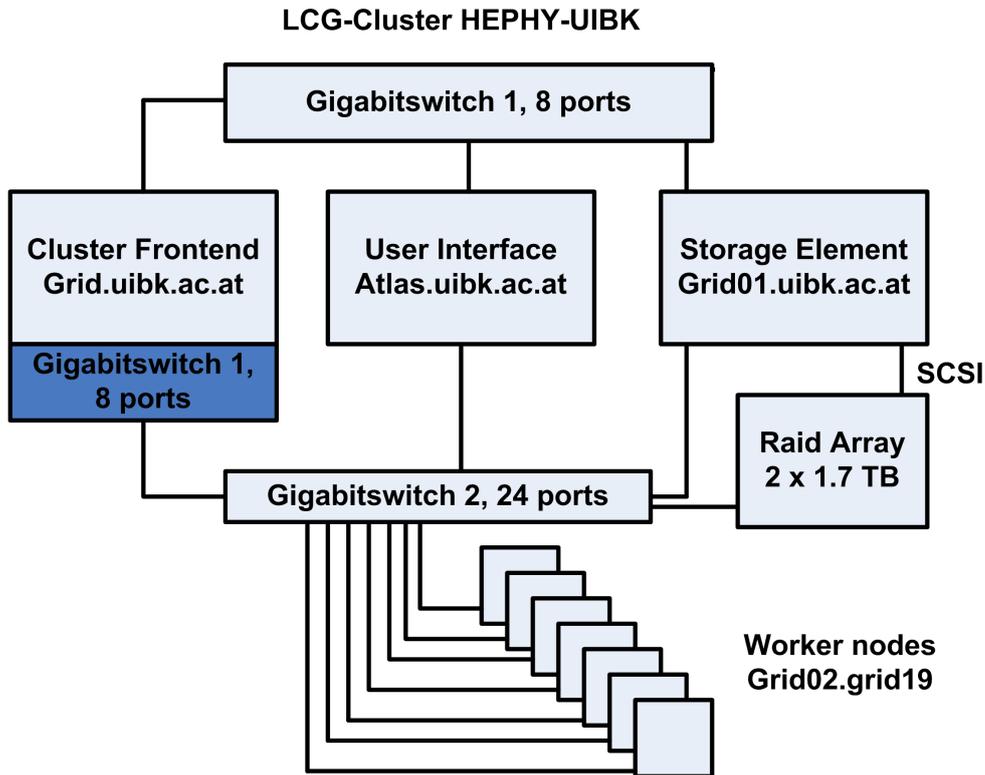


Figure 1: The Architecture of the UIBK LCG Grid Site

per.

3. GRID SITE INFRASTRUCTURE

The UIBK grid site consists of one cluster frontend (dual-core Intel Xeon 2.66GHZ, 2GB RAM, 80GB IDE), which is connected with other grid sites, via the EGEE service grid infrastructure, by a 300Mbps link. The UIBK cluster comprises of 19 Worker Nodes (WNs) which provide the computational resources (dual-core Intel Xeon 3.0 GHz, 2GB RAM, 80GB SATA) (see Figure 1). In addition, there is a storage server (2 Intel Xeon 2.66 GHz, 2GB RAM, 32GB SCSI) which manages two 1.7TB RAID 5 arrays. Several standard grid applications are deployed on the cluster, including Globus GRAM, GridFTP, BDII, gLite software stacks and the Network File System (NFS). All of these distributed solutions generate data over the infrastructure network.

In order to evaluate the traffic behaviour of this grid site, we captured all the traffic generated from the site. This was achieved by running tcpdump on the headnode of the site through which all traffic from the cluster passes. Due to the huge amount of processing power and storage space consumed by the monitoring process, we limited the monitoring period to the duration of seven days. This is because we perform our analysis at the "second" level for seven days of traffic data collection, which already contained more than half a million records. This is sufficiently large for us to perform a reasonably fine-grained analysis and capture detailed traffic behaviours at the packet-level. If characteristics like self-similarity, correlations or long-range dependencies exist, such a duration is long enough to capture these traffic behaviours. In addition, our ultimate traffic modelling aim is to build analytical traffic models "on-the-fly". The model should be able to quickly adjust itself without needing to observe a very large set of traffic traces. This should

achieve a desired level of accuracy while at the same time avoiding non-stationary behaviour [20]. Thus the observation duration should be sufficiently long for our purposes.

The experimental data provides details of all packets sent out from this grid site. We derived network traffic performance metrics of interest, each of which can be regarded as a stochastic process, and then analysed their characteristics. We perform packet-level analysis of throughput, packet size, packet rate, delay and loss for measuring the packet generation behaviour. With these performance metrics defined, we then developed a set of Perl- and AWK-based scripts to process the collected data traffic information. In the following section, we will introduce more detail about the traffic characteristic analysis.

4. TRAFFIC CHARACTERISTICS AND EVALUATION

In order to analyse which applications generate significant amounts of data and TCP connections from the grid site, we first examined how much data is sent over each TCP port and how many connections are initiated by the site nodes over the monitored period. This provides us with useful information about grid site traffic behaviour to learn about the application level traffic characteristics. Then we study the self-similarity properties and analyse suitable statistical models for each traffic metric, which provides more insight into the packet-level grid traffic characteristics.

4.1 Application-Based Traffic

Table 1 shows a summary from a data-oriented and connection-oriented perspective for the grid applications considered¹. As out-

¹Due to space limit, we only show individual port traffic which gen-

Application	Port(s)	Data Sent (%)	Connections (%)
SSH	22	0.38%	0.05%
NFS	2049	29.38%	0.87%
GRAM	2119	0.01%	1.7%
BDII	2170	28.39%	62.25%
GridFTP	2811	0.003%	0.35%
Ganglia	8649	0.97%	6.69%
Globus TCP Port Range	40000–40500	0.42%	1.04%
Other Ports	above 30000	40.108%	26.4049%

Table 1: Application-Based Grid Traffic Analysis

lined in Table 1, BDII and NFS both generate a large proportion of data. BDII generated 28.39% of the total data sent and initiated 62.25% of the 350,000 connections. This indicates that the GLUE [7] dynamic resource information publishing service consumes quite a large proportion of network resources and this traffic is mostly short-live TCP connections. However, this may not be desirable for traffic optimisation and so a more economical scheme for publishing grid site resource information could be considered to strike a trade-off between network resource consumption and the update speed of grid resource information. Further, we see that the Ganglia [3] cluster monitoring system generates a considerably substantial amount of TCP connections. Specifically, we see that Ganglia generates 1% of the data and 6.69% of the total connections. We also know that Globus jobs for this grid site are submitted through port range 40000–40500. Thus, we calculated the data communicated through these Globus job ports and found that they represent around 1.04% of the total connections and just 0.42% of the total data. These results indicate that the data traffic of this grid site is not heavily dominated by Globus job behaviour. Finally we also summarise the total data and connections above TCP port 30000 (including Globus) due to the fact that we observe most other traffics are through these high ports. They represent around 26.4049% of the total connections and 40.108% of the total data.

Overall, from Table 1, we can see that the traffic generated by BDII for publishing resource information and NFS for transferring files have significant impacts on the traffic behaviour of this grid site. It is also observed that the traffic of different applications on this grid site are quite different from each other. For example, applications such as NFS generate more data and less connections, indicating that a relatively large amount of data is transferred for each established TCP connection. In contrast, applications such as BDII and Ganglia generate less data for each established connection but they create these TCP connections quite frequently. This is due to the dynamic nature of grid monitoring and resource publishing applications. As such, when modelling individual grid applications, these characteristics must be taken into consideration. We also see that there are very little HTTP or Web traffic (port 80 or 8080) generated, which is already a distinctive difference from many Web clusters and data centres. For other traffics, due to limited knowledge of the site’s port configuration, we do not know exactly which applications generate them. However, in later sections, we involve all of these traffics for the packet level synthetic traffic analysis.

4.2 Self-Similarity Evaluation

We then analysed the self-similar and long-range dependence of significant amounts of data / number of connections (larger than 1%).

Metric	Hurst Exponent Estimate	Correlation Coefficient
Packet Rate	0.501–0.800	98.35%–99.85%
Throughput	0.516–0.790	98.26%–99.91%
Packet Loss Rate	0.592–0.835	98.40%–99.96%
Packet RTT	0.648–0.698	97.93%–98.67%

Table 2: Self-Similarity Analysis

properties [19, 31] for each evaluated packet-level traffic metric (see Table 2) for the overall traffic behaviour. A self-similar process looks qualitatively the same irrespective of the time scale. Long-range dependence is different in that it measures the memory of a process by measuring its statistical decay rate in comparison to exponential decay. This involves calculating the *autocorrelation decay*, a statistical measure of the relationship between a random variable and itself at different time lags. The Hurst parameter H is a key measure for the persistence of self-similarity and long-range dependence [31]. Technically, a stochastic process $X(t)$ is statistically self-similar with parameter H ($0.5 < H < 1$) if, for any real value $a > 0$, the process $a^{-1}x(at)$ has the same statistical properties as $X(t)$. Here, the closer H is to 1, the stronger the dependence of the process is. Whereas if $H = 0.5$, it follows that the model has short range dependence. It has been shown in [29] that a Poisson model can be used to describe short-range dependent traffic, which does not take into account the autocorrelation concept due to the independence hypothesis in the data. Thus, we measure the autocorrelation (see Figure 2) of each traffic metric as well as an assistant measurement.

With these considerations, we plotted the autocorrelation and estimated the correlation coefficient for each metric. We estimated H on a daily basis, where we adopted rescaled range (R/S) estimator [15] and computed H . Table 2 presents each Hurst parameter estimate and the correlation coefficient results. These experimental results show that packet rate, throughput, packet loss and packet delay are time-dependent and self-similar. This is because the Hurst parameter estimate for both packet rate and throughput trend are larger than 0.5. Together with the autocorrelation results, all of these parameters potentially exhibit long-range dependencies.

4.3 Statistical Analysis

We derive approximate traffic models using Gaussian mixture distributions for the traffic generation behaviour of this grid site. This involves applying the λ^2 discrepancy measure proposed in [30] and the expectation maximisation method to estimate the parameters for Gaussian mixture distributions being tested. Given an initial parameter estimate vector, these methods try to improve it on each subsequent iteration and the algorithm terminates when the stopping criteria is satisfied. With such underlying techniques, we can then find the best approximated models with suitable parameters.

4.3.1 Mathematical Methodology

It has been found that large “messy” data sets [18] might cause approximated models to become biased against the real traffic behaviour. Also discovered in [18], the same occurs when it exhibits significant autocorrelation, which can be seen from our autocorrelation results in Figure 2. Thus, we apply the λ^2 discrepancy measure proposed in [30]. This method has been shown to be effective for examining large messy traffic data [10]. Instead of testing how well a particular data set can be approximated by a probability distribution, the method shown in [30] discovers the discrepancy between

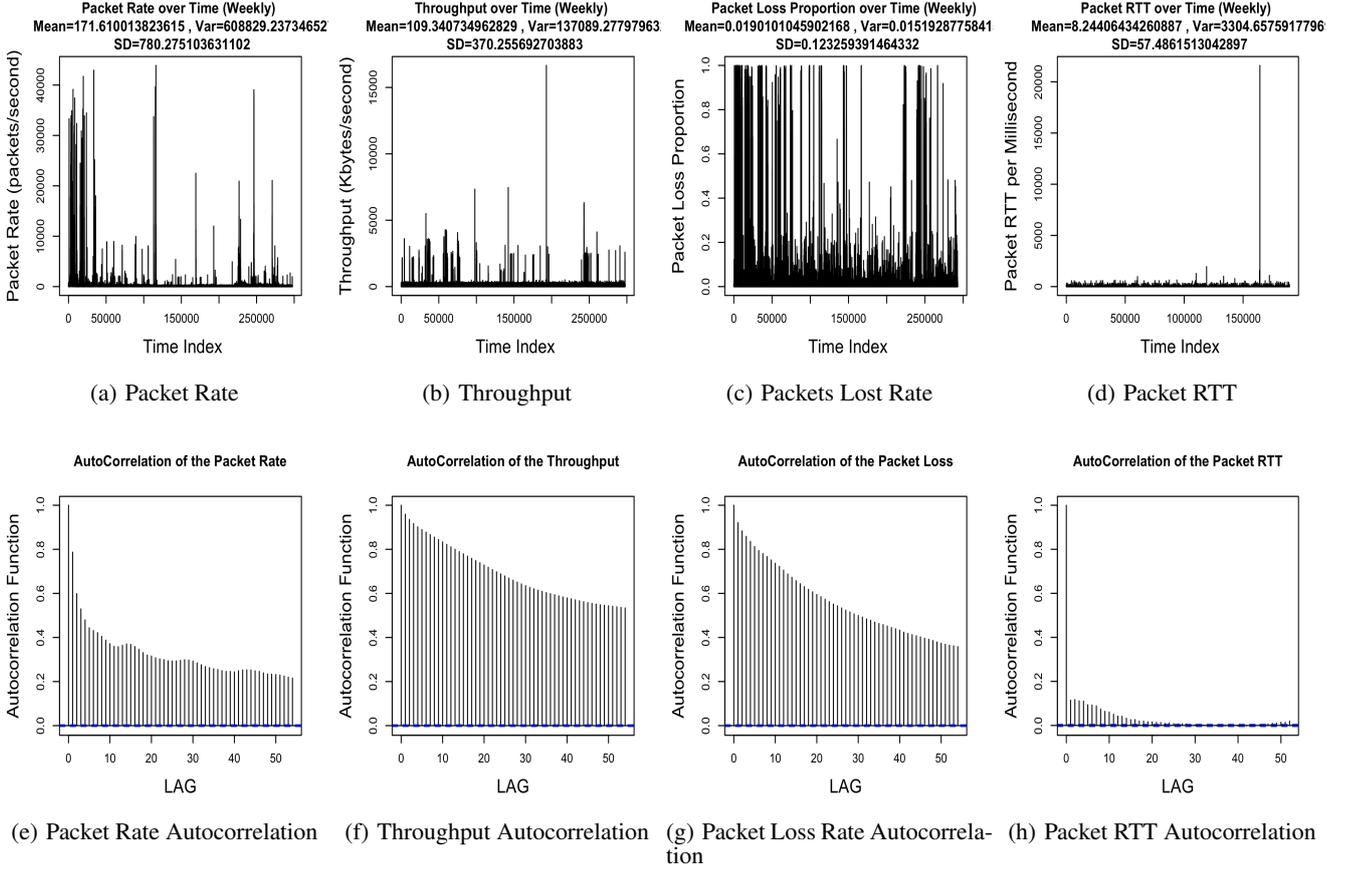


Figure 2: Network Performance for Consecutive 7 Days

the actual data and the statistical model. The larger the discrepancy the worse the model, and vice versa. In practice, $\lambda^2 > 1$ gives a significant difference. The estimator λ^2 is given by:

$$\lambda^2 = \frac{X^2 - K - df}{n - 1},$$

where n is the number of observations in the data set; df represents the number of degrees of freedom of the test given by $df = n - 1 - np_i$ where np_i is the number of parameters that were used to estimate the analytical distribution; X^2 and K are given by the following equations:

$$X^2 = \sum_{i=1}^N \frac{(Y_i - np_i)^2}{np_i}, \text{ and } K = \sum_{i=1}^N \frac{Y_i - np_i}{np_i},$$

where Y_i is the number of items in bin i of the empirical data which is partitioned into N bins. To solve the possible divide-by-zero problem with deterministic distributions, in [23] the following alternative equations are used:

$$X^2 = \sum_{i=1}^N \frac{(np_i - Y_i)^2}{Y_i}, \text{ and } K = \sum_{i=1}^N \frac{np_i - Y_i}{Y_i}.$$

In our analysis, rather than using the original raw data, we use it by applying a logarithmic transformation. This can reduce the range of data and can be used to stabilise the variance of a sample [27]. From the analytical models found in the log-transformed domain, it is easy to derive the analytical expressions of the distribu-

tions into the original domain. Once we get an analytical model, we compare them to the empirical data by means of graphical representations of the corresponding probability density functions (PDFs) and quantile-quantile plot (Q-Q plot). The Q-Q plot is powerful for examining whether the dataset fits well with the analytical model based on the estimated parameters. These are obtained by plotting the quantiles of the two cumulative distributions against each other. The closer to the identity function (*i.e.* a straight line passing through points (0, 0) and (1, 1)) the better the match between the two distributions.

We model the data sets with Gaussian Mixture Models (GMMs) to improve the approximation results rather than using a single distribution, such as the Extreme distribution, which is widely being used as an analytical model for approximating the cluster traffic generation behaviour [13]. The related application of finite GMMs is receiving increasing attention, which is flexible enough to capture various shapes of continuous distributions that do not fall into any of the known parametric probability distributions. With some mixtures, the peaks of the original empirical model may not be obvious. In general, for finite mixtures of k -GMMs, the mixture density is defined as:

$$f(x, p_1, p_2, \dots, p_k, \mu_1, \mu_2, \dots, \mu_k, \sigma_1^2, \sigma_2^2, \dots, \sigma_k^2) =$$

$$\sum_{i=1}^k p_i N(\mu_i, \sigma_i^2),$$

with mixing proportion p_1, p_2, \dots, p_k and such that $\sum_{i=1}^k p_i = 1$.

The mixing proportion denotes the probability that a given observation arises from the i -th mixture component by $x \sim N(\mu_i, \sigma_i^2)$ with probability p_i , where $1 \leq i \leq k$.

Our results show that adopting GMMs and λ^2 discrepancy measures are good and flexible in modelling the shape of the traffic of a grid site. It is worth mentioning that we adopted the Expectation Maximisation (EM) algorithm to estimate the parameters of GMMs. This is because although MLE is the most popular method used to estimate parameters for the distributions being tested, it has certain limitations. For instance, it is required that the likelihood function must be bounded [33]. However, this cannot always be achieved. In our case, the Gaussian mixture likelihood function is unbounded for some measured traffic metrics. Therefore MLE fails to converge, but the convergence of EM algorithm is mathematically proven. Therefore, we adopted the EM algorithm applied to parameter estimation of GMMs.

4.3.2 Traffic Analysis

We are interested in finding an accurate model to characterise packet-level traffic generation behaviour of this grid site. We applied the statistical techniques discussed in Section 4.3.1 to each of our traffic metrics on a separate daily basis. This allowed us to reduce the non-stationary effect of the traffic as well as to show that our adopted modelling approach is adaptive to traffic characteristics.

We show the plots in a comparison of the empirical and estimated analytical results for each day, along with a Q-Q plot. A summary of our findings is given below. To save space, we only show a representative on one day for PDFs and Q-Q plots of the best approximated GMMs. We obtain very similar results for all the other days

(summarised in Table 3 and Table 4) where the GMM is compared with adopting Extreme distribution as the underlying model.

Throughput: Figure 3(a) presents comparisons between empirical and analytical models. It is clear that the analytical data from GMMs approximate the actual traffic data quite well, confirming the goodness of the approximation. The Q-Q plot in Figure 3(b) also shows a good approximation. The GMMs are presented in Table 3, including their parameters and the Extreme distribution discrepancy analysis. For each day, the first row always shows a low value for the discrepancy measure by GMMs, *i.e.* Day 1, $\lambda^2 = 0.1165$ and the second row displays the results for Extreme distribution, *i.e.* Day 1, $\lambda^2 = 0.3690$, relatively poorer than GMMs.

Packet Rate: Figure 4 and Table 4 shows a good approximation between the empirical and GMMs. Note that the Q-Q plot shows a good match, apart from where the difference in the tail. This, however, is a small and insignificant fraction of the overall data set. In most cases, an extreme tail is ignored when constructing a model or, where necessary, the technique of distribution splitting is applied to obtain separate approximations [10].

Packet Loss Rate: Our analysis results show that for almost every day the percentage of packet loss is stable around the range of 2% to 5%. It is interesting to note that on day 1 there is a serious loss of packets, around 13% loss for unknown reasons, which is worth investigating further. Due to space limitations, we only show the PDF of packet loss rate in Figure 5(a). Our GMM-based model approximates closely to the real traffic data with $k = 3$, which confirms the goodness of the approximation.

RTT: Figure 5(b) demonstrates a good agreement between the empirical traffic data and the analytical model of packet RTT, which is shown by GMMs with $k = 3$ as well .

Day	Distribution	Parameters	λ^2
1	GMMs	$\mu_1 = 3.3752, \sigma_1 = 1.6906,$ $\mu_2 = 4.3200, \sigma_2 = 0.4841,$ $p_1 = 0.5862, p_2 = 0.4138$	0.1165
	Extreme	$a = 4.4744, b = 1.5082$	0.3690
2	GMMs	$\mu_1 = 4.2445, \sigma_1 = 0.5087,$ $\mu_2 = 3.8926, \sigma_2 = 1.9592,$ $p_1 = 0.4264, p_2 = 0.5736$	0.1820
	Extreme	$a = 4.8642, b = 1.7398$	0.3874
3	GMMs	$\mu_1 = 2.6092, \sigma_1 = 1.2126,$ $\mu_2 = 4.3127, \sigma_2 = 0.5786$	0.0057
	Extreme	$p_1 = 0.3805, p_2 = 0.6195$ $a = 4.2168, b = 0.9453$	0.0293
4	GMMs	$\mu_1 = 3.1786, \sigma_1 = 1.5985,$ $\mu_2 = 4.2493, \sigma_2 = 0.5442,$ $p_1 = 0.4810, p_2 = 0.5190$	0.0657
	Extreme	$4.3583, b = 1.3391$	0.5001
5	GMMs	$\mu_1 = 2.4609, \sigma_1 = 1.2223,$ $\mu_2 = 4.2681, \sigma_2 = 0.6039,$ $p_1 = 0.2981, p_2 = 0.7019$	0.0111
	Extreme	$a = 4.2502, b = 0.8869$	0.0390
6	GMMs	$\mu_1 = 2.5319, \sigma_1 = 1.2147,$ $\mu_2 = 4.2773, \sigma_2 = 0.5893,$ $p_1 = 0.2998, p_2 = 0.7002$	0.0107
	Extreme	$a = 4.2628, b = 0.8589$	0.0440
7	GMMs	$\mu_1 = 3.4463, \sigma_1 = 1.8693,$ $\mu_2 = 4.2540, \sigma_2 = 0.5243,$ $p_1 = 0.5065, p_2 = 0.4935$	0.0986
	Extreme	$a = 4.5698, b = 1.5671$	0.4584

Table 3: Summary of Results for Throughput

Day	Distribution	Parameters	λ^2
1	GMMs	$\mu_1 = 3.9813, \sigma_1 = 1.4541,$ $\mu_2 = 5.3666, \sigma_2 = 1.8694,$ $p_1 = 0.9250, p_2 = 0.0750$	0.0719
	Extreme	$a = 5.2674, b = 0.8375$	0.1596
2	GMMs	$\mu_1 = 4.3293, \sigma_1 = 0.4277,$ $\mu_2 = 3.7855, \sigma_2 = 1.8523,$ $p_1 = 0.2962, p_2 = 0.7038$	0.1498
	Extreme	$a = 4.7664, b = 1.6706$	0.3466
3	GMMs	$\mu_1 = 3.0776, \sigma_1 = 1.4402,$ $\mu_2 = 4.3741, \sigma_2 = 0.4623,$ $p_1 = 0.6073, p_2 = 0.3927$	0.0791
	Extreme	$a = 4.2260, b = 1.4113$	0.1339
4	GMMs	$\mu_1 = 3.0275, \sigma_1 = 1.4838,$ $\mu_2 = 4.3454, \sigma_2 = 0.4531,$ $p_1 = 0.6203, p_2 = 0.3797$	0.0953
	Extreme	$4.1879, b = 1.3615$	0.9929
5	GMMs	$\mu_1 = 1.9560, \sigma_1 = 0.8671,$ $\mu_2 = 4.2370, \sigma_2 = 0.5892,$ $p_1 = 0.3347, p_2 = 0.6653$	0.0053
	Extreme	$a = 4.0388, b = 0.7667$	0.1249
6	GMMs	$\mu_1 = 3.4476, \sigma_1 = 1.5975,$ $\mu_2 = 4.3833, \sigma_2 = 0.4285,$ $p_1 = 0.6967, p_2 = 0.3033$	0.0209
	Extreme	$a = 4.4222, b = 1.4333$	0.1051
7	GMMs	$\mu_1 = 4.3256, \sigma_1 = 0.5783,$ $\mu_2 = 3.3965, \sigma_2 = 1.9609,$ $p_1 = 0.3019, p_2 = 0.6981$	0.0775
	Extreme	$a = 4.5583, b = 1.6371$	0.2573

Table 4: Summary of Results for Packet Rate

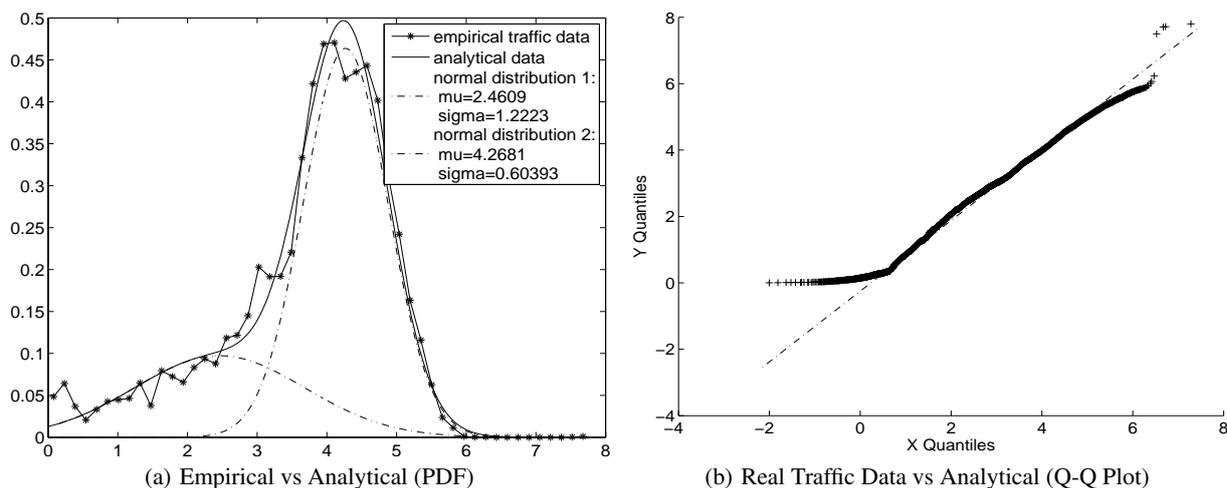


Figure 3: Throughput Day 5

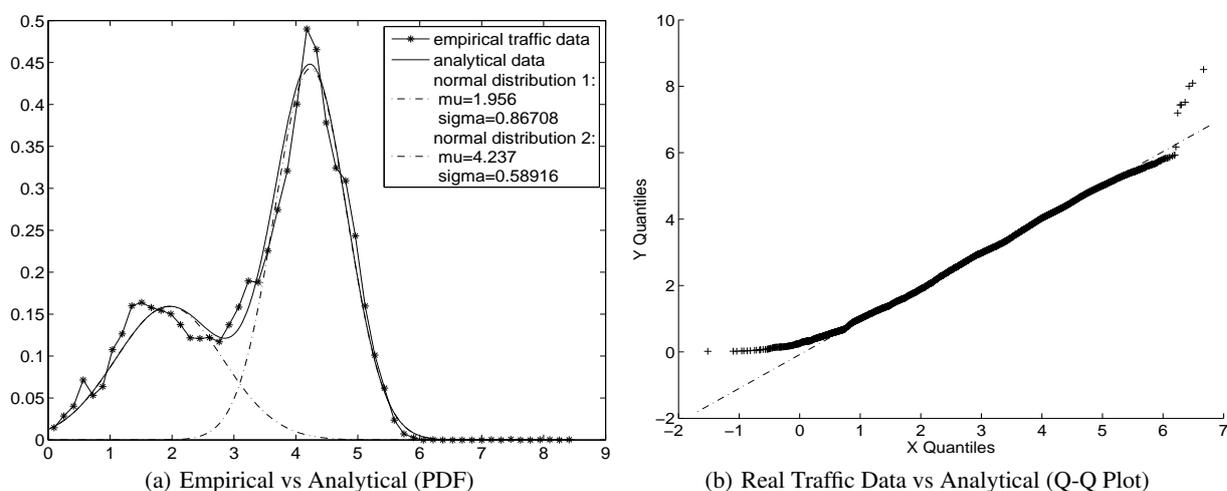


Figure 4: Packet Rate Day 5

Packet Size: For the packet size analysis, our evaluation shows that a large portion of the packets are in the 1024-2047 Byte range, with a relatively small variation. The two dominant packet sizes are 1514 and 66 Bytes, claiming 55.1% and 36.6%, respectively, of all captured packets. Packets of such sizes typically represent the maximum transmission unit (MTU)-worth of data and an acknowledgement (ACK) respectively (including TCP, IP and MAC headers).

Finally, we analyse the correlation between packet size, packet rate and throughput. Intuitively, the packet rate metric should be independent of the packet size and have a linear relationship with the network throughput. These dependence relationships could be interpreted quantitatively using the correlation coefficient value between the metrics. When two sets of data are completely independent the correlation coefficient value should be 0, whereas when two sets of data are completely dependent the correlation coefficient value should be 1. In our case, the correlation coefficient value between packet size and packet rate is 0.001319813, which indicates these two metrics are actually almost independent of each other. However, we found that packet rate has a positive correlation

with throughput as the correlation coefficient approaches 0.4797655.

4.3.3 Summary

Overall, all of these statistical analysis results help us identify the synthetic traffic behaviour of the grid site. Such information is important for understanding network utilisation and efficiency of grid sites, thus allowing optimal design and fine-tuned management of the underlying network. The GMM-based modelling approach can be reused as a foundation for “on-the-fly” analytical model approximation with continuously updated parameters. Such an approach can achieve an accurate model approximation adaptively as well as avoid the non-stationary effects of long term observation. The results presented in this paper can also be used to build approximate data generation models for grid sites and develop simulators.

5. IMPLICATIONS & FUTURE WORK

The work presented here is merely a starting point. Our analysis is from one grid site which might be not representative enough of the LCG or other grid sites. However, this work reveals a number of interesting questions and research directions, which would be

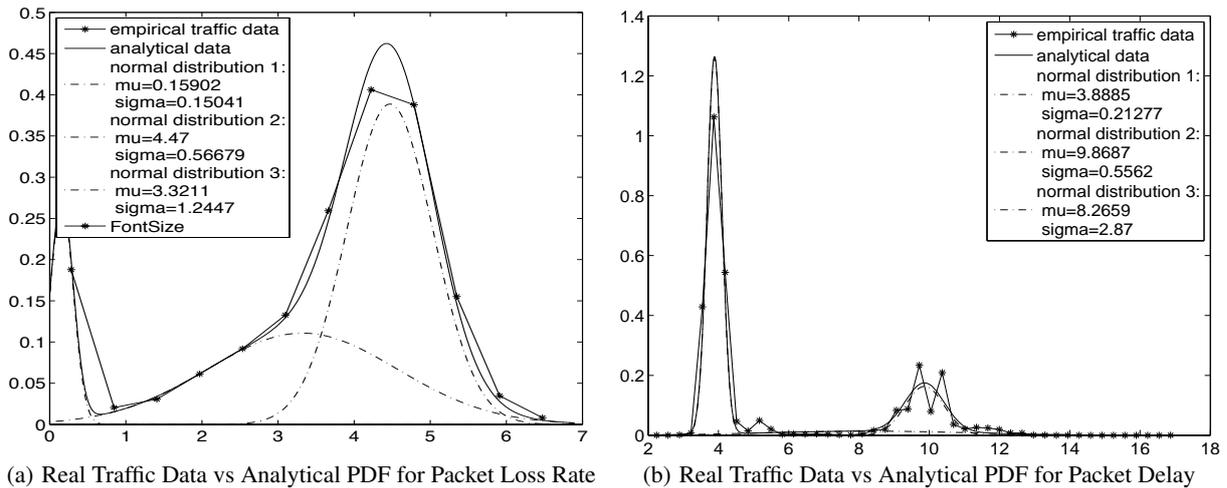


Figure 5: Packet loss rate and RTT Day 5

important for network designers and researchers who wish to incorporate a realistic representation of grid site traffic in their work. Thus, these results are still quite valuable as, to our knowledge, there is little work for grid site traffic analysis at this level.

The characteristics of the data traffic generated by the UIBK grid site is useful for several reasons. Although the observation duration is only seven days, the traffic traces contain more than half a million records at the “second” level contributing to highly fine-grained traffic analysis. These traces are long enough to discover many fine degree traffic behaviours, such as self-similarity, long-range dependencies and traffic correlations etc., all of which are important to be taken into consideration for future grid sites network design, performance enhancement and traffic prediction. For example, as pointed out in [11], the Hurst parameter H , used to measure the self-similarity property, provides a better satisfactory measure of the “burstiness” of self-similar traffic. The nature of congestion produced by self-similar traffic models differs from that predicted by standard formal models and displays a more complicated behaviour.

In addition, in this paper we reveal that both BDII and NFS consume a large proportion of this grid site traffic. For NFS traffic, it might lead us to rethink the design and network management of grid sites or more generally high performance clusters. The question arising here is how should we optimise the design of future grid sites to avoid such a big chunk of NFS traffic to storage servers generated from the grid headnode, which might be undesirable. For BDII traffic, it brings up the question of how should we optimise the frequency of the grid resource publishing behaviour to avoid excessive advertisement traffic? For example, when the workload of a grid site is high, should a grid site publish its current status more frequently to avoid being selected or it should publish resource information less frequently to free some of its network resources? Such questions need to be answered for optimising grid systems.

Furthermore, as we see from our analytical results, most metrics can be approximated reasonably well by using GMMs throughout the seven days duration. These results imply that even if we collect more data, we probably can still use GMMs as the underlying model with different parameters. Therefore, GMMs can be a good candidate as an empirical traffic modelling tool. In our case, two or three Gaussian process mixtures ($k = 2, 3$) can achieve a reasonable approximation. This result shows two benefits for future work

on modelling or simulation of grid site traffic generation. First, with relatively low dimensional GMMs ($k = 2, 3$), the number of parameters for GMMs are not large, which balances the complexity of the model and the accuracy of traffic approximation. Second, since GMMs fit better than most other empirical distributions for all the analysed metrics, we can avoid the use of many other distributions for model approximation in the future. These single distribution-based models are still used in a significant number of works on traffic analysis, *e.g.* [13]. This can reduce the computation overhead significantly. Since GMMs only contain Gaussian processes with a lower computation overhead and the simplicity of the underlying mathematics form, traffic analysis can be embedded as a module of a real-time traffic monitoring process. This allows “on-the-fly” accurate traffic model extraction and traffic performance prediction in real-time. This approach would be promising for building future grid network monitoring frameworks as well as for high-level applications, which need such real-time information for application specific purposes (*e.g.* server selection in terms of the predicted network performance).

The models developed in this paper can also be used for optimising the performance of network protocols or grid applications (such as TCP, GridFTP). For example, Yildirim *et al.* [34] show how to improve TCP throughput performance by constructing a balance among tuned buffer sizes and parallel data streams. In their algorithms, network metrics such as packet RTT and packet loss rate are needed for optimisation. The traffic characterisation results introduced in this paper can be adopted to achieve a better prediction of these network metrics. Our approaches can be used to further improve such optimisations, which also directly implies a performance improvement of GridFTP-based parallel data transfer behaviour and its workload as introduced in [21].

For future work, we hope to analyse multiple grid site traffic with a finer degree of granularity and a longer duration. This might allow us to confirm our findings or discover new traffic characteristics that differ from the results of this paper. Either outcome would be valuable. We also aim to perform detailed analysis for some critical grid network interactions, such as BDII-based resource advertisements. We also hope to adopt more sophisticated analysis methods, such as fractional Brownian motion (fBm) models to approximate the self-similar and non-stationary behaviour of the data traffic. Finally, we plan to build accurate grid site traffic generators based on

our empirical observations.

6. CONCLUSION

Understanding the nature of traffic in grid sites is essential for engineering, operations, and performance optimisation. To achieve such goals, it is important to know the expected traffic behaviour of the major contributors to the grid traffic. In this paper, we introduced the evaluation results for the traffic behaviour observed at the UIBK site, an LCG site running real-world production-level grid services. We investigated the proportion of traffic generation by different applications and found that traditional HTTP and other Web applications, which generate a large proportion of the traffic on the Internet, have only a minor impact on grid sites. Instead, grid applications, such as BDII for grid resource advertisement, generate a significant proportion of traffic. Such results implicate different traffic behaviour between a grid site and Web or data centres.

Furthermore, we revealed that the traffic behaviour of this grid site exhibits self-similarity and long-range dependency. Such behaviour has serious implications for the design, management and optimisation of high-speed networks, such as in congestion avoidance and traffic load balancing. It indicates that an optimal data traffic solution based on traditional memoryless Poisson traffic models will be suboptimal when the traffic exhibits self-similarity and long-range dependencies properties.

We performed a packet-level statistical analysis of the generated traffic. Using the χ^2 discrepancy measure we developed the analytical models by means of a mixture of Gaussian processes. Such models have a good agreement between the actual traffic data and analytical modelling as well as keep a reasonable model complexity, which is not too difficult to implement. This is better than using a single distribution presented in many existing results. The simplicity of using only Gaussian processes also enables the possibility of “on-the-fly” traffic model approximation with reasonable computational overhead, which would allow us to build accurate “on-the-fly” traffic predictions. This can be applied to real-time traffic monitoring and higher-level optimisation of data intensive applications.

Overall, our results in this paper cover a number of important aspects of the network data traffic characteristics of a grid site. To our knowledge, the traffic analysis work we presented in this paper is the only analysis of a real-world production level grid site. We believe that such an evaluation is valuable for future grid site design, development, performance improvement, network management and data-aware application design.

7. ACKNOWLEDGEMENT

We thank Gregor Mair and Peter Oettl at the University of Innsbruck and Michael Welzl at the University of Oslo for helping us collect traffic traces. We also thank the European Commission for funding the EC-GIN project [4] under which this work was completed.

8. REFERENCES

- [1] Berkeley Database Information Index. <https://twiki.cern.ch/twiki/bin/view/EGEE/BDII>.
- [2] Globus Resource Allocation Manager (GRAM). <https://www.globus.org/toolkit/docs/2.4/gram>.
- [3] Ganglia. <https://ganglia.sourceforge.net>, April 2003.
- [4] Europe-china grid internetworking (ec-gin), fp6 strep project, contract number 045256. <http://www.ec-gin.eu/>, 2006.
- [5] W. Allcock, J. Bester, J. Bresnahan, A. Chervenak, L. Liming, S. Meder, and S. Tuecke. GridFTP Protocol Specification. *GGF GridFTP Working Group Document*, Sep. 2002.
- [6] E. Altman, K. Avrachenkov, and C. Barakat. A stochastic model of tcp/ip with stationary random losses. *IEEE/ACM Trans. Netw.*, 13(2):356–369, 2005.
- [7] S. Androzzzi, S. Burke, L. Field, S. Fisher, B. Konya, M. Mambelli, J. Schopf, M. Viljoen, and A. Wilson. Glue schema specification version 1.2. 2005.
- [8] P. Barford and M. Crovella. Generating representative web workloads for network and server performance evaluation. *ACM SIGMETRICS Performance Evaluation Review*, 26(1):151–160, 1998.
- [9] J. Beran, R. Sherman, M. Taqqu, and W. Willinger. Long-range Dependence in Variable-bit-rate Video Traffic. *IEEE Trans. on Comm.*, 43(234):1566–1579, 1995.
- [10] M. Borella. Source models of network game traffic. *Computer Communications*, 23(4):403–410, 2000.
- [11] M. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic: Evidence and Possible Causes. *IEEE/ACM Transactions on networking*, 5(6):835–846, 1997.
- [12] Y. El-khatib and C. Edwards. A survey-based study of grid traffic. In *Proceedings of the International Conference on Networks for Grid Applications (GridNets '07)*, 2007.
- [13] D. Ersoz, M. Yousif, and C. Das. Characterizing Network Traffic in a Cluster-based, Multi-tier Data Center. In *Distributed Computing Systems, 2007. ICDCS'07. 27th International Conference on*, pages 59–59, 2007.
- [14] W. Feng and P. Tinnakornsrisuphap. The Failure of TCP in High-performance Computational Grids. In *Supercomputing, ACM/IEEE 2000 Conference*, pages 37–37, 2000.
- [15] A. Field, U. Harder, and P. Harrison. Network Traffic Measurements in a Switched Ethernet Environment. In *UKPEW2002*, pages 47–58, June 2002.
- [16] A. Field, U. Harder, and P. Harrison. Measurement and Modelling of Self-similar Traffic in Computer Networks. *IEE Proceedings-Communications*, 151(4):355–363, 2004.
- [17] D. Figueiredo, B. Liu, A. Feldmann, V. Misra, D. Towsley, and W. Willinger. On TCP and Self-similar Traffic. *Performance Evaluation*, 61(2-3):129–141, 2005.
- [18] L. Gleser and D. Moore. The effect of dependence on chi-squared and empiric distribution tests of fit. *The Annals of Statistics*, 11(4):1100–1108, 1983.
- [19] T. Karagiannis, M. Molle, and M. Faloutsos. Long-range Dependence: Ten Years of Internet Traffic Modeling. *IEEE Internet Computing*, 8(5):57–64, 2004.
- [20] T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido. A nonstationary Poisson view of Internet traffic. In *IEEE INFOCOM*, volume 3, pages 1558–1569. Citeseer, 2004.
- [21] N. Kourtellis, L. Prieto, A. Iamnitchi, G. Zarrate, and D. Fraser. Data transfers in the grid: workload analysis of globus gridftp. In *DADC '08: Proceedings of the 2008 international workshop on Data-aware distributed computing*, pages 29–38, New York, NY, USA, 2008. ACM.
- [22] A. Kumar. Comparative performance analysis of versions of TCP in a local network with a lossy link. *IEEE/ACM Transactions on Networking (TON)*, 6(4):485–498, 1998.
- [23] J. Lakkakorpi, A. Heiner, and J. Ruutu. Measurement and characterization of Internet gaming traffic. In *Research Seminar on Networking, Helsinki University of Technology*,

Networking Laboratory, Espoo, Finland. Citeseer, 2002.

- [24] Z. Liu, N. Niclause, and C. Jalpa-Villanueva. Traffic model and performance evaluation of web servers. *Performance Evaluation*, 46(2-3):77–100, 2001.
- [25] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling tcp throughput: a simple model and its empirical validation. *SIGCOMM Comput. Commun. Rev.*, 28(4):303–314, 1998.
- [26] R. Pang, M. Allman, M. Bennett, J. Lee, V. Paxson, and B. Tierney. A First Look at Modern Enterprise Traffic. In *Proc. Internet Measurement Conference*, pages 15–28, 2005.
- [27] V. Paxson. Empirically derived analytic models of wide-area tcp connections. *IEEE/ACM Trans. Netw.*, 2(4):316–336, 1994.
- [28] V. Paxson. End-to-end routing behavior in the Internet. *ACM SIGCOMM Computer Comm. Review*, 36(5):56, 2006.
- [29] V. Paxson and S. Floyd. Wide Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM Transactions on Networking (TON)*, 3(3):226–244, 1995.
- [30] S. Pederson and M. Johnson. Estimating model discrepancy. *Technometrics*, 32(3):305–314, 1990.
- [31] W. Stallings. *High-speed Networks: TCP/IP and ATM Design Principles*. Prentice-Hall, Inc., NJ, USA, 1997.
- [32] W. Willinger, M. Taqqu, R. Sherman, and D. Wilson. Self-similarity Through High-variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level. *IEEE/ACM Transactions on Networking (TON)*, 5(1):71–86, 1997.
- [33] D. Xu. The applications of mixtures of normal distributions in empirical finance: A selected survey. Working Papers 0904, Univ. of Waterloo, Dept. of Economics, Sept. 2009.
- [34] E. Yildirim, D. Yin, and T. Kosar. Balancing tcp buffer vs parallel streams in application level throughput optimization. In *DADC '09: Proceedings of the second international workshop on Data-aware distributed computing*, pages 21–30, New York, NY, USA, 2009. ACM.