

Towards Modeling of Traffic Demand of Node in Large Scale Network

Kensuke Fukuda
National Institute of Informatics
Tokyo, 101-8430, Japan
kensuke@nii.ac.jp

Abstract—Understanding actual network and traffic properties of the Internet is essential to determine network parameters in large-scale network simulations. However, there is little knowledge about the distribution of macroscopic traffic demand for each node, though the topological properties of the network have been focused on. This paper investigates the distribution of traffic volume to and from a node at an organization level. As traffic volume data, we used byte counter data of all interfaces in all backbone routers in a nation-wide research and education (R&E) network in Japan. First, we show that traffic volumes to and from a node in the network are characterized by a lognormal distribution, which has a slower decay than a normal distribution, but a faster decay than a power-law distribution. Thus, an assumption in which the traffic demand is uniformly random or Gaussian distributed is not appropriated to model the traffic demand in large-scale network simulation. This finding implies that one has more possibility to observe an increase of delay or packet drop in simulation, comparing to the result that uses uniformly-random or Gaussian traffic demand, because of the locality of traffic. Moreover, we observed that in 87% of nodes, a traffic volume from the backbone to the node is 1-10 times larger than that for the opposite direction. This is a similar usage pattern appeared in residential light-user broadband traffic. Finally, we introduce a simple model to explain the distribution of traffic demand, based on a multiplicative growth of traffic volume. We confirm that the multiplicative model can reproduce a lognormal distribution of traffic volume by simple numerical simulation.

I. INTRODUCTION

Simulation is still an essential method to verify usefulness of new network algorithms for routing, congestion control, and peer-to-peer traffic control, even though large-scale testbed environments are now available. For large-scale simulation, network traffic and topology models potentially affect the total performance of the proposed algorithm. In the worst case, the performance results may not make sense because inappropriate network and traffic models were chosen [5].

There have been many studies to characterize statistical properties of network topology and traffic over years. Specifically, characterization of the inter- and inner-AS level Internet topologies is a current hot topic in the network research fields [4], [16], [1], [19], [10], [14]. Recent studies showed that the Internet topology is far from a sort of random graph, and is basically characterized by scale-free and small-world properties [4], [16]. The former indicates the appearance of a power-law distribution of node degree, and the latter means the network has the small diameter and locally clustered property. These topological properties clearly affect the performance of

the network [1]. Although theoretical and empirical studies are revealing more and more topological characteristics, these studies might still be far from the perfect model.

Toward finding for more realistic parameters in large-scale network simulation, following issues at least must be clarified; (a) the statistics of the network topology, (b) the assignment of the link property (i.e., bandwidth), (c) the traffic demand of each node.

This paper focuses on issue (c). To quantify the traffic demand in a large-scale network, ideally, traffic distribution per application should be characterized spatially and temporally, then an appropriate mixture of those traffic distributions per node should be identify as the source and sink of traffic in a node. However, there is little knowledge about the traffic distribution among nodes currently, although many methods have been proposed for estimating a characteristics of OD pairs (origin-destination) by using packet trace [2], [11], [21], [8]. Actually, a packet trace from a single backbone link is still useful for the estimation of the traffic demand of a node. However, it is highly difficult to capture all packets to and from nodes at all the related links in a real large-scale inter- and inner-AS network, especially in a hierarchical network. Also, there is an alternative approach that quantifies the relationship between the population of countries (or region) and the geographic location of Internet resources (e.g., router and interface) [9]. However, this method cannot provide us actual traffic demand, because it is hardly to map the population into a finer scale than the country-scale, especially into AS-level.

Hence, we need to clarify what kind of distribution (i.e., power-law, lognormal, uniform, or normal distribution) should be assumed in simulation from macroscopic level. In this paper, as a first step toward a better understanding of a traffic demand distribution in a real network, we investigate the traffic demand and growth in a nation-wide R&E network, called SINET [18], which provides network connectivity to over 700 organizations in Japan. Our analysis is based on the traffic volume data measured by the interface counter of almost all routers in our network via SNMP, rather than a packet traffic trace in a backbone.

The main contributions of our paper are as follows; (1) We found that the distribution of traffic to and from a node in the network is widely spread, and well characterized by a lognormal distribution. Thus, one cannot set traffic demand of a node in a given network by uniformly randomly or in

a Gaussian manner. The direct implication of this result is that the traffic demand followed by the lognormal distribution will cause the increase of the delay and packet drop for an performance evaluation of a new algorithm, because of the locality of traffic. (2) We introduced that the asymmetric nature of the distribution of traffic demand distribution can be explained by a multiplicative growth model. Thus, this implies that one should use the lognormal model as long as the traffic growth of the network is multiplicative as widely observed in several networks.

II. DATASETS

We used the node traffic data measured at a Japanese nation-wide R&E network called SINET. SINET provides an Internet connectivity to over 700 organizations, corresponding to universities, colleges, and research laboratories. For organizations, it provides commodity network services as well as special-purpose services for “big science” such as astronomical research, high-energy physics research, and so on. We analyzed commodity traffic and special-purpose traffic separately, because dedicated lines were deployed for the latter purpose.

We monitored each network link between the backbone and an organization for administrative purposes (We refer to an organization as a node in the following sections). However, deployment of a flow export architecture (e.g., NetFlow) to all links is costliness, instead, we have been collecting interface counter values of almost all routers in our backbones via SNMP, and archiving the per-interface traffic logs with RRDtool [13]. Thus, while the data are too coarse-grained to analyze a dynamics of microscopic flows, the data allow us to fully investigate the macroscopic behavior of traffic in the large-scale network. Each raw traffic trace consisted of mean values of traffic volume whose bin size was 2 hours passing through the link between a node and its upstream for 5 weeks in May-Jun, 2006.

For the analysis, the mean traffic volume during 5 weeks was calculated for each node. Our data trace included only commodity traffic generated by end-users; thus, it did not count any dedicated lines for big science. We also excluded the traffic data of the edge routers between the network and the rest of the Internet (i.e., the international links, commercial links, and other R&E network links). This is because that we cannot identify any node information outside our network from these data. Thus, the traffic amounts of incoming and outgoing were different in our analysis. Furthermore, we normalized a mean traffic volume of node because of a security reason.

In the following sections, we use the terms “incoming” and “outgoing” from the viewpoint of backbone. The Incoming traffic signifies the traffic from a node to the backbone, corresponding to the upload traffic from the end-user’s perspective. The outgoing traffic flows from the backbone to a node.

III. ANALYSIS RESULTS

One of the main purposes of this research is to characterize the functional form of the distribution of traffic volume for

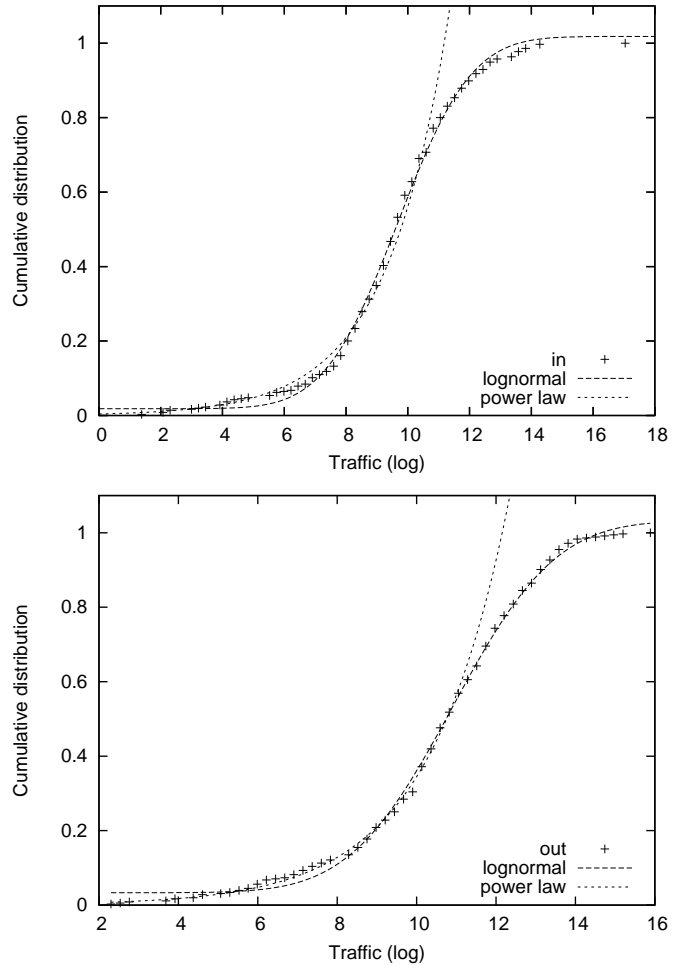


Fig. 1. Result of least square fitting of incoming traffic (top) and outgoing traffic (bottom). The x-axis indicates the log-scale traffic volume and the y-axis shows a cumulative distribution. The fitting functions are a lognormal distribution and a Pareto distribution

each node in the network. To identify the functional form, we fitted two typical functional forms having a long-tail to the data plots by least-squares fitting; a Pareto distribution and a lognormal distribution. The probability density function of a Pareto distribution is given by

$$p(x) = ax^{-(a+1)}, \quad (1)$$

for $a > 0$. a is called as a shape parameter controlling a burstiness. A Pareto distribution is a well-known distribution for modeling Internet traffic time series [15]. The probability density function of a lognormal distribution is given by

$$p(x) = \frac{1}{x\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(\log x - \mu)^2}{2\sigma^2}\right) \quad (2)$$

where μ is the mean of $\log(x)$, and σ is the standard deviation of $\log(x)$. Thus, the mean of $p(x)$ is $\exp(\mu + \sigma^2/2)$ and the variance is $\exp(2\mu + \sigma^2)(\exp(\sigma^2) - 1)$. A lognormal distribution decays faster than that of a Pareto distribution, and much slower than that of a normal distribution by definition.




Fig. 2. The growth of the backbone traffic volume

We did not fit the plots to a normal distribution because the plots clearly have a stretched shape.

The best fit among two functional forms was by the lognormal distribution, as shown in Figure 1. It is apparent that the fit of the lognormal distribution covers a wide spread of both incoming and outgoing traffic volumes. More interestingly, both traffic volumes can be modeled with the same functional form but with different mean and variances.

Next, we analyzed the growth rate of the traffic volume in the network. To calculate the total traffic volume in the network for a month, we summed up incoming mean traffic volumes of nodes and incoming mean traffic volumes of the border links between our network and the rest of the Internet. The same calculation was carried out for the outgoing mean traffic volume. These calculations guaranteed to avoid double counting of traffic volumes.

Figure 2 displays the growth of the mean traffic volume in the backbone over 13 months from 2005 to 2006. The volumes are normalized to the volume in May 2005. In principle, the incoming and outgoing traffic volumes would be the same; however, some traffic data were lost due to frequent changes in the network configuration. The traffic volume of the R&E network is characteristically affected by the behavior of students, who appear and disappear with the seasonal variation; traffic volume decreases during summer and spring vacations. However, ignoring the seasonal variations, the estimated traffic growth rate was 50% per year. This value is consistent with the traffic growth rate in IXEs and that for the residential broadband users in Japan described in [6], [3].

Finally, we plot the cumulative distribution $P(> \rho)$ of outgoing/incoming ratio ρ of mean traffic volume in a node in Figure 3. $\rho > 1$ indicates a node where an outgoing volume is larger than an incoming one, as an information consumer. On the other hand, a node with $\rho < 1$ corresponds to an information supplier. In the special case, $\rho = 1$ is a node having the same incoming and outgoing traffic volumes. We found that only 13% of nodes generates more incoming traffic volume than outgoing one. In other words, 87% of nodes are

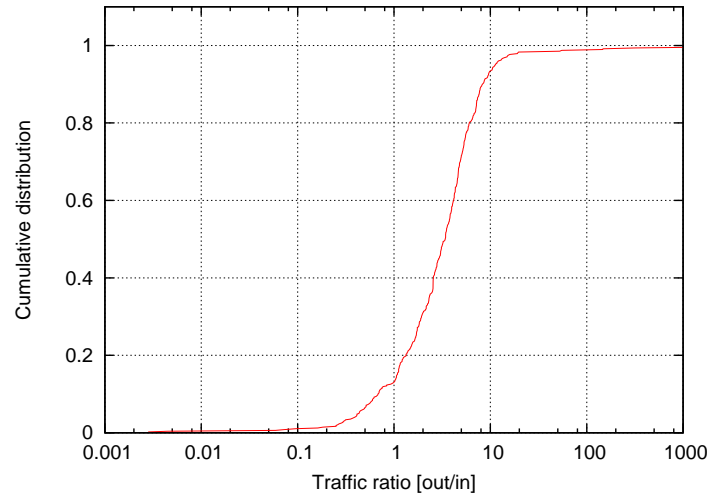


Fig. 3. Cumulative distribution of outgoing/incoming ratio of mean traffic volume.

more downloading than uploading. 5%- and 95%-percentiles are 0.4 and 11.3, respectively. Since our data set did not include big science data, those data were mainly generated by interactive end-users in an organization. Because the main usages of network such users are likely web browsing, streaming, video-conference and a few peer-to-peer software, these traffic patterns tend to indicate download-dominant asymmetric nature. Furthermore, we hardly confirmed a clear difference in the outgoing/incoming relationships among the types of data link media from a node to the backbone, except for the bandwidth limitation. These results supports an idea that a statistical behavior of both incoming and outgoing traffic volumes representing the behavior of aggregated users can be modeled by a unique functional form.

IV. REPRODUCTION MODEL

In the previous section, we obtained evidences to estimate the distribution of traffic volumes for each node in the network. This section discusses a reproduction mechanism of the traffic volume distribution characterized by a lognormal distribution.

A simple and well-known model to reproduce a lognormal distribution is the one proposed by R. Gibrat [7]. This model is based on a multiplicative stochastic process, and it was originally proposed for explaining the distribution of incomes for individuals/companies in the economic research field.

In terms of the traffic volume distribution, a traffic volume $x(t+1)$ of a node at time $t+1$ is given by

$$x(t+1) = b(t)x(t). \quad (3)$$

Here, $b(t)$ corresponds to the growth rate satisfying an ideally independently distributed stochastic variable with positive values. The main idea is that the random growth of the traffic volume is expressed as a percentage of the current weight, and is independent of its current actual value. Thus, if the process




Fig. 4. The numerical simulation result of the multiplicative model: (a) the plot after 20 iterations of multiplicative growth, (b) the plot after 40 iterations of multiplicative, and (c) the plot after 20 iterations of multiplicative growth and more 20 iterations of additive growth.

is iteratively used, one obtains

$$x(t+1) = b(t) * b(t-1) * \dots * x(0) = \left(\prod_{j=0}^t b(j) \right) * x(0), \quad (4)$$

where $x(0)$ is the initial traffic volume of a node selected from uniformly randomly distributed variables. One takes the logarithm of the above equation,

$$\log(x(t+1)) = \sum_{j=0}^t \log(b(j)) + \log(x(0)). \quad (5)$$

This means that $\log(x(t+1))$ follows the normal distribution by the central limit theorem. Therefore, $x(t+1)$ follows a lognormal distribution. In other words, the distribution of traffic volume for nodes in the network follows a lognormal distribution when each node increases its traffic according to Eq.(3). It should be emphasized that the multiplicative increase is essential for the reproduction of a lognormal distribution; the above equations clearly indicate that the additive increase (i.e., $x(t+1) = x(t) + b(t)$) only yields to a normal distribution.

To apply this model to our network data, we should discuss the growth rate of the traffic volume. The results shown in Fig.4 and Ref. [3] estimated that the growth rate of the macroscopic traffic is 30-50% per year. Additionally, Ref. [3] reported that traffic growth rates have been stable for recent four recent years. We have no information on the growth of traffic volume in the target network before 2004; however, it is plausible that the growth rate would be almost the same value as that of the residential broadband traffic. Also, another evidence for supporting the assumption is that a multiplicative traffic growth was appeared in Internet backbone traffic in US during 1990-2002 [12].

We performed simple numerical simulations of the process we explained above; There are 1000 nodes each of which has an initial traffic taken from a uniformly random value,

$1 \leq x_0 \leq 10^4$. In each time step t , the traffic volume in a node increases, according to Eq.(3), where $b(t)$ is a uniformly random value $1.0 \leq b(t) \leq 2.0$. Figure 4 shows the distributions of traffic volumes reproduced by the model. The plots for 20 and 40 iterations of the multiplicative growth model clearly show a lognormal distribution from the initial state with uniformly random distribution. Thus, we can reproduce the observed traffic volume distribution over nodes in the real network. Moreover, in order to investigate an effect of a change of traffic growth patterns, we carried out an additional simulation consisting of 20 iterations of the additive growth followed by 20 iterations of the multiplicative growth. The shape of the resulting plot still resembles that of 20 iterations in higher traffic volume, though there is a difference in smaller traffic volume due to the additive growth. In conclusion, the effect of the multiplicative growth still remains even if it changes to the additive one ¹.

V. CONCLUDING REMARKS

This paper analyzed the distribution of traffic demand of nodes in a network, for the purpose of determining suitable network parameters for large-scale simulations. By using interface byte counter information, we first showed that the distribution of the traffic demand can be modeled with a lognormal distribution quite unlike a uniformly random, or a Gaussian distribution. This is a pessimistic message; lognormal distributed traffic exhibits a highly asymmetric nature, so that it likely causes longer delays and more possibility of the congestion in comparison with the result of setting traffic to be randomly located. The asymmetric nature of the traffic demand of a user (or a flow) has also been reported in Refs [2], [17], [6]. They did not identify the functional form of the traffic demand distribution, although the data indicated a slower decay than a Gaussian distribution. Considering the fact that the distribution of the number of users among organizations cannot be so widely spread, the distribution of the traffic demand of a user has a high probability to follow a lognormal distribution.

Our result of the outgoing/incoming ratio of traffic volumes resembles the traffic usage of light users in residential broadband networks described in Ref. [3]; A light user uses fewer peer-to-peer applications, and most usages are web browsing. In our dataset, although we were not able to obtain the traffic breakdown, the network links are for commodity traffic (i.e., no big science traffic), and the network operators in an organization filter out peer-to-peer traffic by manual. These are likely a reason why our result was consistent with that for light-users. As Ref.[3] concluded that there is no clear boundary between light user and heavy-hitter, the distribution of traffic demand may change its parameters (i.e., mean and variance) as a result of the appearance of new applications.

¹Of course, if an increasing volume in the additive growth is assumed to be much larger than the current volume, the shape of the distribution would largely change. However, we have no evidence of such traffic evolution, currently.

We finally introduced a reproduction model of the observed behavior based on the multiplicative growth. Although the model itself is simple, the simulation result explained the actual distribution very well. Sustaining multiplicative growth might be a strong assumption for the future Internet environment. However, even though the network growth will saturate in the future, the distribution of the traffic demand is still likely to be closer to a lognormal distribution than a normal distribution, because (1) there already has been enough iterations which stand for the multiplicative assumption and (2) once the distribution of the traffic is spread over the several orders of magnitude, it is not easy to smooth out. Similarly, a multiplicative growth model is used for explaining the evolution of the network topology [20]. In this sense, the multiplicative growth might be a key function to characterize basic network properties.

In this paper, we clarified a basic property of the distribution of traffic demand, as a first step. However, there are many unsolved problems. One of them is to quantify the relationship between source and sink nodes. Our SNMP-based analysis cannot be directly applied to this problem. A deeper traffic breakdown is required to understand the traffic behavior. For each node, the network usage should be clarified in terms of applications, and we will analyze this by using a packet capture data passing through our backbone link. Also, the traffic matrix estimation is likely to be a useful approach to quantify the source and sink relationship from a more coarse-grained viewpoint. Next, we should quantify the relationship between the relative position of the network and the traffic demand. For this, we can use node specific information such as degree and clustering coefficient. This information is useful in locating the traffic demand in the network. The last question is the applicability of our results. Our result is likely consistent with one for an inner-AS network, although we would need to do a comparison of other networks to say for certain. Moreover, we currently have no evidence that our result applies to an inter-AS network. It is impossible to collect the traffic demand from many ASes by using our approach. Thus, we should use a sophisticated method like traffic matrix estimation [11], [21], [8] to estimate the traffic demand from other measurement points.

ACKNOWLEDGEMENT

This research was partially supported by JSPS Grant-in-Aid for Young Scientists (B).

REFERENCES

- [1] R. Albert, H. Jeong, and A.-L. Barabási. Attack and error tolerance of complex networks. *Nature*, 406:378–382, 2000.
- [2] N. Brownlee and K. Claffy. Understanding internet traffic streams: dragonfiles and tortoises. *IEEE Communications Magazine*, pages 110–117, Oct. 2002.
- [3] K. Cho, K. Fukuda, H. Esaki, and A. Kato. The impact and implications of the growth in residential user-to-user traffic. In *SIGCOMM*, pages 207–218, Pisa, Italy, Sept. 2006.
- [4] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM*, pages 251–262, Portland, OR, Aug. 1999.

- [5] S. Floyd and V. Paxson. Difficulties in simulating the internet. *IEEE/ACM Transactions on Networking*, 9(4):392–403, 2001.
- [6] K. Fukuda, K. Cho, and H. Esaki. The impact of residential broadband traffic on Japanese ISP backbones. *SIGCOMM CCR*, 35(1):15–21, Jan. 2005.
- [7] R. Gibrat. *Les inégalités économiques*. Paris, Sirey, 1931.
- [8] A. Gunnar, M. Johansson, and T. Telkamp. Traffic matrix estimation on a large ip backbone - a comparison on real data. In *SIGCOMM Internet Measurement Conference*, pages 149–160, Sicily, Italy, Oct. 2004.
- [9] A. Lakhina, J. Byers, M. Crovella, and I. Matta. On the geographic location of internet resources. *IEEE Journal on Selected Areas in Communications*, 21(6):934–948, Aug 2003.
- [10] L. Li, D. Alderson, W. Willinger, and J. Doyle. A first-principles approach to understanding the internet’s router-level topology. In *SIGCOMM*, pages 3–14, Portland, OR, Aug. 2004.
- [11] A. Medina, N. Taft, S. Bhattacharyya, and C. Diot. Traffic matrix estimation: Existing technology and new directions. In *SIGCOMM*, pages 161–176, Pittsburgh, PA, Aug. 2002.
- [12] A. M. Odlyzko. Internet traffic growth: Sources and implications. In *Optical Transmission Systems and Equipment for WDM Networking II*, pages 1–15, San Jose, CA, Jan. 2003.
- [13] T. Oetiker. RRDtool: Round Robin Database Tool. <http://ee-staff.ethz.ch/~oetiker/webtools/rrdtool/>.
- [14] R. V. Oliveira, B. Zhang, and L. Zhang. Observing the evolution of internet as topology. In *SIGCOMM*, pages 313–324, Kyoto, Japan, Aug. 2007.
- [15] K. Park and W. Willinger, editors. *Self-Similar Network Traffic and Performance Evaluation*. Wiley, 2000.
- [16] R. Pastor-Satorras, E. Smith, and R. V. Sole. Dynamical and correlation properties of the Internet. *Physical Review Letters*, 87:258701, 2000.
- [17] L. Plissonneau, J.-L. Costeux, and P. Brown. Analysis of peer-to-peer traffic on ADSL. In *PAM2005 (LNCS3431)*, pages 69–82, Boston, MA, Mar. 2005.
- [18] SINET: Science Information Network in Japan. <http://www.sinet.ad.jp>.
- [19] N. Spring, R. Mahajan, and D. Wetherall. Measuring isp topologies with rocketfuel. In *SIGCOMM*, pages 147–161, Pittsburgh, PA, Aug. 2002.
- [20] X. Wang and D. Loguinov. Wealth-based evolution model for the internet as-level topology. In *INFOCOM 2006*, pages 1–11, Barcelona, Spain, Apr. 2006.
- [21] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg. Fast accurate computation of large-scale ip traffic matrices from link loads. In *SIGMETRICS*, pages 206–213, San Diego, CA, June 2003.