

# On the characteristics of Internet traffic variability: Spikes and Elephants

Tatsuya Mori

NTT Service Integration Laboratories,  
NTT corporation

3-9-11 Midoricho Musashino-city Tokyo Japan,  
and

Dept. of Information and Computer Science,  
School of Science and Engineering,  
Waseda University

3-4-1 Okubo Shinjuku Tokyo Japan  
tatsuya@nttlabs.com

Shozo Naito

NTT Information Sharing Laboratories,  
NTT corporation

3-9-11 Midoricho Musashino-city Tokyo Japan  
naito.shozo@lab.ntt.co.jp

Ryoichi Kawahara

NTT Service Integration Laboratories,  
NTT corporation

3-9-11 Midoricho Musashino-city Tokyo Japan  
kawahara.ryoichi@lab.ntt.co.jp

Shigeki Goto

Dept. of Information and Computer Science,  
School of Science and Engineering,  
Waseda University

3-4-1 Okubo Shinjuku Tokyo Japan  
goto@goto.info.waseda.ac.jp

## Abstract

*Analysing and modeling of traffic play a vital role in designing and controlling of networks effectively. To construct a practical traffic model that can be used for various networks, it is necessary to characterize aggregated traffic and user traffic. This paper investigates these characteristics and their relationship. Our analyses are based on a huge number of packet traces from five different networks on the Internet. We found that: (1) marginal distributions of aggregated traffic fluctuations follow positively skewed (non-Gaussian) distributions, which leads to the existence of “spikes”, where spikes correspond to an extremely large value of momentary throughput, (2) the amount of user traffic in a unit of time has a wide range of variability, and (3) flows within spikes are more likely to be “elephant flows”, where an elephant flow is an IP flow with a high volume of traffic. These findings are useful in constructing a practical and realistic Internet traffic model.*

## 1 Introduction

Analysing and modeling of traffic play a vital role in designing networks and controlling traffic effectively. To con-

struct a practical traffic model, it is necessary to characterize both aggregated traffic and user traffic for the following reasons. Firstly, the characteristics of aggregated traffic are essential metrics for designing network bandwidth. This is because the fluctuations of aggregated traffic strongly affect the queuing behavior of routers. Knowledge of the fluctuations is also useful in the evaluation of QoS mechanisms such as differentiated services and queue scheduling algorithms. Secondly, the characteristics of user traffic are important for the purposes of traffic engineering. Several recent papers have proposed router algorithms to control per-flow bandwidth [7, 9, 17]. Exact knowledge of these characteristics is helpful in fixing the parameters for such algorithms or proposing more efficient algorithms. Characterising user traffic is also useful in estimating the performance experienced by individual users.

Up to now, a number of studies have investigated the characteristics of aggregated traffic and user traffic. The numerous papers on aggregated traffic have revealed that fluctuations of aggregated traffic in various networks exhibit self-similarity [8, 16]. Self-similarity means that traffic fluctuations exhibit long-term time correlation. It has also been found that the marginal distributions of aggregated traffic fluctuation play a key role in determining network performance [1, 4, 11]. Here, marginal distribution indicates a distribution of throughput per unit time (see lower

part of Fig. 1). On the other hand, several papers on user traffic have revealed that the duration of individual user traffic follows heavy-tailed distribution [3, 20]. The characteristics were modeled by the well known “Pareto ON/OFF” model. Recent measurement-based studies have also found that the size or average rate of user traffic exhibits quite a wide range of variability [14, 21].

Although the characteristics of aggregated and user traffic have been widely studied, few reports have discussed the relationship between them<sup>1</sup>, which is essential for consistently incorporating both characteristics into a traffic model. Since aggregated traffic is the superposition of user traffic, the two characteristics are related. Knowing the relationship is also useful for determining traffic engineering policy. For instance, let us consider the following case. Aggregated traffic sometimes exhibits extremely large values (i.e., bursts) which may cause severe packet loss or long queuing delays at routers. If we know the characteristics of user traffic when bursts occur, we can establish a per-flow traffic engineering policy to control these bursts.

In this study, we analyzed the characteristics of aggregated traffic and user traffic and investigated the relationship between the two. We obtained a huge number of packet traces from five different networks on the Internet. The rest of this paper is organized as follows. Section 2 describes how the traffic data was collected. Section 3 studies the characteristics of aggregated traffic fluctuations and discusses *spikes*. Here, a spike is an extremely large value of momentary throughput. In Section 4, we define *per-time-unit flow* and investigate the characteristics of user traffic. We also demonstrate the existence of *elephant flows*, which are IP flows with high volume traffic. Section 5 explains that elephant flows occur more frequently in spikes. Section 6 discusses the results and section 7 concludes this paper.

## 2 Traffic Data

We obtained packet traces from five different networks on the Internet. We captured packet data at a main communication link in each network. None of the links were severely congested, but all had sufficient traffic. As is well known, traffic exhibits a wide range variability during long-term measurement. For example, the average traffic volume is quite different during busy daytime periods than during less busy night periods. Therefore, we should capture traces during an adequate period to avoid the effect of non-stationarity. In this paper, all traces were captured within a period of 300 seconds during daily busy hours. We assumed

---

<sup>1</sup>The most well known work is the one by Willinger et al. [20]. They found that superposition of user traffic with heavy-tailed duration time makes aggregated traffic self-similar. However, as we later show, their traffic model cannot cover the characteristics of aggregated and user traffic on today’s Internet. We will discuss this issue later.

that traffic variability would be stationary within this time interval. Under these conditions, each trace had sufficient packets to be analyzed statistically.

In our analyses, we used one-way traffic as summarized in Table 1 because traffic volumes are asymmetrical. Investigating traffic with a larger volume is meaningful from the viewpoint of evaluating performance. The networks we used in our analyses are briefly described below.

**NTTlab** This link is used as the main external connection line of NTT R&D center (NTTlab). It is a 12-Mbps ATM circuit. We captured trace data at an Ethernet segment connected to the external link via the gateway router. The measurements were conducted during daily busy hours on various weekdays in July 2001. In total, we used 56 traces for this paper.

**Waseda** This link is the main external connection line for Waseda university, which is connected to the Inter-Ministry Research Information Network (IMnet) [5]. The link is a 100-Mbps Ethernet. The measurements were done during daily busy hours on various weekdays in April and July 2001. We used a total of 71 traces for this paper.


**OCN-SINET** This link connects NTT’s Open Computer Network (OCN) and the Science Information Network (SINET). OCN is the commercial Internet backbone network operated by NTT, and SINET is the largest Internet backbone network for scientific research institutes in Japan. The link is a 135-Mbps ATM circuit. The measurements were done during daily busy hours on various weekdays in January 2000. We obtained a total of 32 traces for this paper. For information about this data, see [6].

**APAN** This link is one of the international lines for the Asia-Pacific Advanced Network (APAN) [2]. It connects the APAN Tokyo Exchange Point in Japan and the STAR TAP in the US. The link is an OC3 circuit. We captured trace data using OC3mon/Coral [12]. The measurements were done during daily busy hours on various weekdays in August 1999. We used a total of 44 traces for this paper.

**WIDE** This is one of the international links for the Widely Integrated Distributed Environment (WIDE) project – the largest Internet research community in Japan [19]. The line is a 100-Mbps Ethernet with an 18-Mbps CAR (committed access rate). The trace data is available from the MAWI traffic archive [10]. All the traces were measured during daily busy hours (14:00–). We used traces measured on various weekdays in September to November 2001. In total, we used 66 traces for this paper. More detailed information on this data is given in [10].

**Table 1. Summary of data**

data	line	direction	bandwidth	# traces	avg packets	avg rate
NTTlab	ATM	incoming	12 Mbps	56	$1.55 \times 10^5$	3.75 Mbps
Waseda	Ethernet	incoming	100 Mbps	71	$1.84 \times 10^6$	23.31 Mbps
OCN-SINET	ATM	OCN-to-SINET	135 Mbps	32	$8.06 \times 10^5$	12.50 Mbps
APAN	OC3	US-to-JP	155 Mbps	44	$5.98 \times 10^5$	5.14 Mbps
WIDE	Ethernet	US-to-JP	100 Mbps	66	$1.61 \times 10^6$	18.00 Mbps



**Figure 1. Marginal distributions of traffic fluctuations**

**Summary** We used 269 traces from five different networks, and all of them were captured during a period of 300 seconds. Table 1 summarizes these.

### 3 Characteristics of Aggregated Traffic Fluctuations: Spikes

Let us analyze fluctuations of aggregated traffic for each link. We use throughput time series  $X(t_i)$  to represent traffic fluctuations.  $X(t_i)$  are defined as follows. Let  $t_i$  ( $1 \leq i \leq M$ ) be time intervals, where  $M$  is the total number of time intervals in a trace. For all  $i$ , the length of  $t_i$  is the same time interval  $\tau$ . For each  $t_i$ , we define variable  $X(t_i)$  as the throughput in  $t_i$ . It is calculated with formula  $X(t_i) = b(t_i)/\tau$ , where  $b(t_i)$  is the total number of bits sent during the time interval  $t_i$ . In this paper, the length of  $\tau$  is set to 0.1 seconds (100 milliseconds). Since each trace covers 300 seconds, the number of time intervals  $M$  for each trace is 3000 ( $3000 = 300/0.1$ ). It should also be pointed out that, while we used a time interval of 0.1 seconds, the following results also held for various time intervals such as 0.01 seconds, 1 seconds, and 10 seconds.

Figure 1 shows the time series of throughput variables

$X(t_i)$  (upper) and their marginal distributions (lower) for randomly chosen traces in the five networks described in Section 2. For the comparison, we draw the curve of Gaussian distributions having the same mean and standard deviations (dashed lines). These figures clearly show that for all five traces, the marginal distributions of traffic fluctuations are not Gaussian. They are all positively skewed.


To characterize the distributions, we use *skewness*, defined as

$$\text{skewness} = \frac{\langle (X(t_i) - \langle X(t_i) \rangle)^3 \rangle}{\sigma_X^3}, \quad (1)$$

where  $\langle X(t_i) \rangle$  stands for the *mean* of  $X(t_i)$  and  $\sigma_X$  is the *standard deviation* of  $X(t_i)$ . If the distribution is skewed positively (negatively), skewness becomes positive (negative). If the distribution is strictly Gaussian, the shape of the distribution is symmetric and the skewness must be 0. For the five traces in Fig. 1, we calculate the skewness of throughput variables. Table 2 reveals that all skewnesses take positive values. This quantitatively indicates that all the distributions in Fig. 1 are positively skewed. In Table 2, the traffic fluctuations of the NTTlab trace is positively skewed most in these five traces.

**Table 2. Skewness of five traces**

	NTTlab	Waseda	OCN-SINET	APAN	WIDE
skewness	1.31	0.92	0.65	0.58	0.69



**Figure 2. Skewness of all traces**

We calculate the skewness of throughput variables for all traces. Figure 2 illustrates the skewness for all traces in the five networks. Each point corresponds to one trace. While there are a few traces with a skewness close to zero, most have positive skewness. The average skewness is at least 0.39 for the five networks.

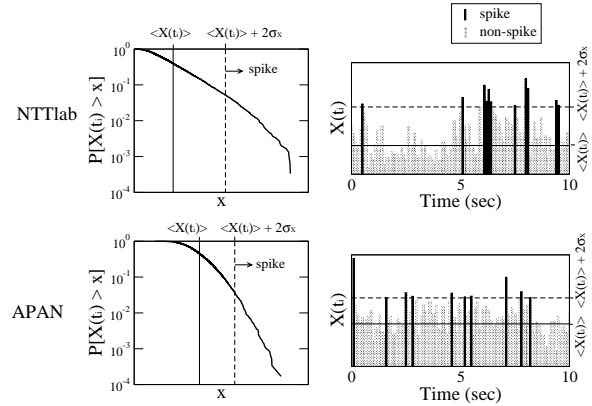
From the above, we conclude that marginal distributions of traffic fluctuations are positively skewed in various networks. Some previous studies such as those by Addie et al. [1], and Grossglauser and Bolot [4] also suggest that marginal distributions of traffic fluctuations are positively skewed and do not follow Gaussian distributions. However, the amount of traffic data used in their studies was quite a few. Our analyses are based on vast number of traces. This supports the results widely for various networks — several Internet uplinks and backbone lines.

The positively skewed marginal distribution of traffic fluctuations suggests that there are throughput variables  $X(t_i)$  whose values are extremely large. These extremely large throughput values may cause severe packet loss or long queueing delay at routers. That is, they are strongly related to network performance. Therefore, it is meaningful to characterize them. Here, we define an extremely large throughput variable  $X(t_i)$  as a *spike*.

**Definition 1:** We call a throughput variable  $\hat{X}(k)$  a spike if it satisfies

$$\hat{X}(k) \geq \langle X(t_i) \rangle + 2\sigma_X, \quad 1 \leq k \leq N_S, \quad (2)$$

where  $\langle X(t_i) \rangle$  stands for the mean of  $X(t_i)$ ,  $\sigma_X$  is the standard deviation of  $X(t_i)$ , and  $N_S$  is the num-



**Figure 3. Examples of spikes**

ber of spikes in a trace. The set of spikes  $\Omega_S \ni \hat{X}(k)$  ( $1 \leq k \leq N_S$ ) satisfies  $\Omega_S \subset \Omega_X$ , where  $\Omega_X$  is the set of throughput variable  $X(t_i)$  ( $1 \leq i \leq M$ ).

We set a mean value plus two standard deviations for the threshold value of large throughput variables — spikes. Figure 3 has examples of spikes for the NTTlab and APAN traces used in Fig. 1. In Fig. 3, the graphs at the top are NTTlab traces and those at the bottom are APAN traces. These two traces have the maximum and minimum skewness among Table 2. The graphs at the left are the marginal distributions of traffic fluctuations (throughput), and the ones at right are actual time series for throughput (part of them). The solid lines indicate  $\langle X(t_i) \rangle$  and the dashed lines indicate  $\langle X(t_i) \rangle + 2\sigma_X$ . In the figures at right, the black zone corresponds to spikes and the gray zone corresponds to non-spikes.

Let us investigate the relationship between spikes and skewness. We define the *occurrence ratio* of spikes in each trace. It is calculated as  $N_S/M$ , where  $M$  is the number of throughput variables in a trace and  $N_S$  stands for the number of spikes in the trace. It should also be noted that in a Gaussian distribution, the occurrence ratio of spikes (i.e., the probability a random value will exceed the mean plus two deviations) is about 0.023.

Figure 4 is a scatter plot of the occurrence ratio and the skewness for all the traces in the five networks. Here, each point corresponds to one trace. The solid lines represent the linear regression, and variable  $R$  indicates the value of the




Figure 4. Scatter plot of the occurrence ratio of spikes and skewness

correlation coefficients. It is clear that the occurrence ratio and the skewness are positively correlated for all five networks. The correlation coefficients are at least 0.63. If the occurrence ratio of spikes increases, the skewness also increases. We can also see that in most traces, the occurrence ratios of spikes are higher than the Gaussian distribution (i.e., most of occurrence ratios are greater than 0.023).

These observations can be summarized as follows. In all traces, marginal distributions of traffic fluctuations are positively skewed. According to positively skewed distributions, there are a larger number of spikes than those calculated assuming Gaussian distribution. Since spikes may cause packet loss or queueing delays, it is necessary to know their characteristics and causal mechanisms.

#### 4 Characteristics of User Traffic: Elephant

This section analyzes the characteristics of user traffic. We investigate an IP flow as user traffic. Let us first introduce *per-time-unit flow* which is an IP flow defined in a given time unit. We can then calculate the size distributions of per-time-unit flows for all traces. We found that there were quite large per-time-unit flows with non-negligible probability. We call these *elephant flows*.

##### 4.1 Per-time-unit flow

We introduce a per-time-unit flow to divide the traffic into flows according to the time interval. We see how the behavior of each flow contributes to the characteristics of aggregated traffic.

Figure 5 is a diagram of per-time-unit flow. In this figure, each rectangle corresponds to one IP flow. An IP flow is a group of IP packets having the same combination of source IP address, destination IP address, source port, destination port, and the protocol field. We divide an IP flow into  $M$  pieces, where  $M$  stands for the number of time intervals defined in Section 3. The length of time interval  $\tau$  is again set to 0.1 seconds. For each time interval  $t_i$ , we define per-time-unit flow  $fl(t_i, j)$ ,  $1 \leq j \leq N_{t_i}$  (shaded regions in




Figure 5. Diagram of per-time-unit flow.

Fig. 5).  $N_{t_i}$  is the number of per-time-unit flows in time interval  $t_i$ . The per-time-unit flow  $fl(t_i, j)$  should contain at least one packet in time interval  $t_i$ .

For each per-time-unit flow  $fl(t_i, j)$ , we define its size as  $b(fl(t_i, j))$  (bits). Throughput variable  $X(t_i)$  satisfies


$$X(t_i) = \frac{\sum_{j=1}^{N_{t_i}} b(fl(t_i, j))}{\tau}. \quad (3)$$

For all traces in Fig. 1, we calculate  $N_{t_i}$  for each  $X(t_i)$  belonging to spikes ( $X(t_i) \in \Omega_S$ ) and non-spikes ( $X(t_i) \in (\Omega_X - \Omega_S)$ ) separately. We then take the average values of  $N_{t_i}$  for spikes and non-spikes. Table 3 lists the results. Although the average value of  $N_{t_i}$  is slightly larger in spikes than in non-spikes, the difference is not large. Since the numbers of  $N_{t_i}$  are almost the same, we may conclude that a large value of  $X(t_i)$  derives from large values of  $b(fl(t_i, j))$ . That is,  $b(fl(t_i, j))$  is considered to be the major factor determining whether  $X(t_i)$  is a spike or not, and  $N_{t_i}$  is not relevant.

In the following sections, we focus on the size distribution of per-time-unit flow  $b(fl(t_i, j))$ .

**Table 3. Average of  $N_{t_i}$  for spikes and non-spikes**

	NTTlab	Waseda	OCN-SINET	APAN	WIDE
spike	17.8	80.9	191.7	50.9	253.6
non-spike	15.2	78.2	174.3	48.3	213.7



**Figure 6. LLCD plots of  $b(fl(t_i, j))$ .**

## 4.2 Size distribution of per-time-unit flow — occurrence of elephant flows —

We investigate the size distribution of per-time-unit flow  $fl(t_i, j)$  (i.e., the distribution of  $b(fl(t_i, j))$ ).

We calculate the following complementary cumulative distribution.

$$P[b(fl(t_i, j)) > x] \quad (4)$$

Figure 6 shows the log-log complementary cumulative distribution (LLCD) plots of  $b(fl(t_i, j))$  for the five traces in Fig. 1. The distributions show that  $b(fl(t_i, j))$  varies over several orders of magnitude. That is, there are large per-time-unit flows with non-negligible probability, while most per-time-unit flows are relatively small. We call a large per-time-unit flow an *elephant flow*.

**Definition 2:** We call a per-time-unit flow  $\hat{fl}(m) = fl(t_i, j)$  an elephant flow if it satisfies

$$b(\hat{fl}(m)) \geq \langle b(fl(t_i, j)) \rangle + 2\sigma_b, \quad 1 \leq m \leq N_E, \quad (5)$$

where  $b(\hat{fl}(m))$  is the size of the flow,  $\langle b(fl(t_i, j)) \rangle$  is the mean of  $b(fl(t_i, j))$  and  $\sigma_b$  is the standard deviation of  $b(fl(t_i, j))$ , and  $N_E$  is the total number of elephant flows in the trace. The set of elephant flows  $\Omega_E \ni \hat{fl}(m)$  ( $1 \leq m \leq N_E$ ) satisfies  $\Omega_E \subset \Omega_F$ , where  $\Omega_F$  is the set of per-time-unit flows  $fl(t_i, j)$  ( $1 \leq i \leq M, 1 \leq j \leq N_{t_i}$ ).

Similar to the definition of spikes, we set the mean value plus two standard deviations for the threshold values of large per-time-unit flows — elephant flows. The dashed lines in Fig. 6 indicate the threshold values.

We define the *occurrence ratio* of elephant flows (number) and *occupation ratio* of elephant flows (size). They are defined by

$$\text{occurrence ratio} = \frac{N_E}{\sum_i N_{t_i}} \quad (6)$$

and

$$\text{occupation ratio} = \frac{\sum_m b(\hat{fl}(m))}{\sum_{i,j} b(fl(t_i, j))}. \quad (7)$$

We calculate the above two ratios for the traces in Fig. 1. Results are summarized in Table 4 and these indicate that for all five traces, a small number of elephant flows occupy a large part of the entire aggregated traffic<sup>2</sup>.

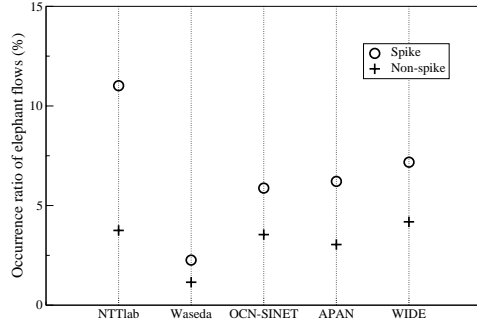
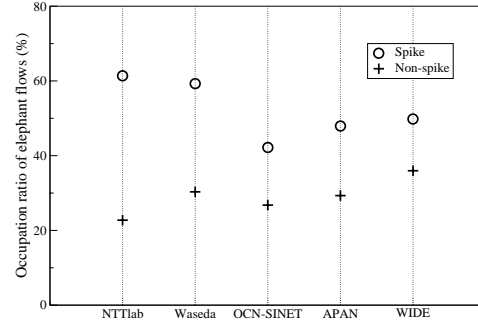
This kind of phenomenon is well known as the *elephant and mice* phenomenon or *the vital few and the trivial many* rule. Some papers have reported that the size or average rate of IP flows follows the elephant and mice phenomenon [14, 21]. In this study, we found that the size of a per-time-unit flow (i.e., momentary rate of IP flow) also have this tendency. Thus, a small percentage of flows occupies a large part of the whole aggregated traffic in a unit of time.

While researchers such as Papagiannaki et al. [14] and Zhang et al. [21] focused on the behavior of average flow rate (i.e., average flow rate is equal to flow size divided by flow duration time), this paper uses per-time-unit flows. This is because they represent the behavior of IP flows more accurately during a short time interval. As the rate of an IP flow is not constant but is highly variable, our approach and findings can directly address the relationship between flow-level behavior and the characteristics of aggregated traffic.

<sup>2</sup>While we used the five traces in Fig. 1, the results were the same for the rest of the traces as well.

**Table 4. Occurrence and occupation ratios of elephant flows**

	NTTlab	Waseda	OCN-SINET	APAN	WIDE
Occurrence ratio	3.01 %	1.15 %	2.35 %	3.22 %	4.65 %
Occupation ratio	35.72 %	39.07 %	25.59 %	41.22 %	41.30 %

**Figure 7. Occurrence ratio of elephant flows (average)****Figure 8. Occupation ratio of elephant flows (average)**

## 5 Spikes and Elephants

As we have discussed in the previous subsections, there are elephant flows, which have a large number of packets during a given short time interval. We presume that the existence of elephant flows is closely related to spikes. In other words, we presume that within a spike, the occurrence of elephant flows would be higher than in other areas (e.g., in non-spikes).

To prove this, we calculate the following statistics. For per-time-unit flows belonging to spikes and non-spikes, we separately calculate the occurrence and occupation ratios of elephant flows. Their definitions are the same as those given in formulas (6) and (7).

Figures 7 and 8 show the occurrence and occupation ratios of elephant flows for all five traces. We used the average value for traces belonging to each network. The figure shows that both ratios of elephant flows are higher in the spikes. This means that:

1. Within spikes, the occurrence probability of elephant flows is higher than that in non-spikes (about 1.7 – 2.9 times higher).
2. A large part of the spikes are constructed by elephant flows (for spikes, about 42 – 61% of per-time-unit flows are elephant flows).

We thus verified our presumption.

The most significant result here is that there are elephant flows in spikes at a higher probability than in non-spikes. In other words, when there are a larger number of elephant flows in a time interval, the throughput variable of the time interval tends to be a spike. Our flow level analyses reveal that the occurrence of elephant flows is an essential factor in generating spikes.

## 6 Discussion

The relationship between the characteristics of aggregated traffic and user traffic plays a key role in constructing a practical and realistic traffic model. However, little work has focused on this issue. The heavy-tailed ON/OFF source traffic model [20] is the best well known and most valuable one, which takes both the characteristics into account. This traffic model reflects the characteristics of (1) self-similarity (aggregated traffic) and (2) heavy-tailed duration (user traffic). A number of measurement analysis have shown the generality of these characteristics. The model explains both characteristics in a consistent fashion. That is, the superposition of heavy-tailed ON/OFF user traffic leads to the self-similarity of aggregated traffic. Willinger et al. [20] proved the above relationship within a rigorous mathematical framework, and Park et al. [15] also verified it by simulation. The model has also been implemented on such widely used network simulators as NS [18] or OPNET [13].

However, the model has two big weaknesses. Firstly, the

marginal distribution of aggregated traffic fluctuations is assumed to be *Gaussian*. As we have shown, this Gaussian assumption does not agree with the characteristics of real Internet traffic. Secondly, traffic is assumed to be aggregated from a number of flows having the same fixed rate. Therefore, each flow has the same size on a certain time scale when the flow is in an ON-period. As we have seen, flow rates (i.e., size of per-time-unit flow) vary over several orders of magnitude.

Thus, it is essential that we include positively skewed marginal distributions of traffic fluctuations (aggregated traffic) and the great variability in flow rate distribution (user traffic) in the traffic model. Our measurement-based analyses revealed these characteristics and the relationship between them. By considering these characteristics, we can construct a practical and realistic Internet traffic model. Our findings will also be useful in the setting of traffic engineering policies for user traffic. For instance, since elephant flows make a large contribution to spikes (i.e., bursts), controlling user traffic that takes the form of elephant flows will effectively suppress bursts. We will extend our research to a theoretical investigation of the relationship to better understand these phenomena in network traffic.

## 7 Conclusion

In this paper, we investigated the characteristics of aggregated and user traffic, and the relationship between them. We collected a huge number of packet traces from five different networks on the Internet. Our main findings are as follows. (1) Marginal distributions of aggregated traffic fluctuations followed positively skewed (non-Gaussian) distributions, which led to the existence of spikes, (2) the size of user traffic per unit time revealed a wide range of variability, and (3) elephant flows were more likely within spikes. These findings will be useful in the construction of practical and realistic Internet traffic models.

## References

- [1] R. G. Addie, T. D. Neame, and M. Zukerman. Modeling superposition of many sources generating self similar traffic. In *Proceedings of ICC*, pages 387–391, June 1999.
- [2] Asia-Pacific Advanced Network. <http://www.apan.net>.
- [3] M. Crovella and A. Bestavros. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. In *Proceedings of ACM SIGMETRICS*, May 1996.
- [4] M. Grossglauser and J. C. Bolot. On the relevance of long-range dependence in network traffic. *IEEE/ACM Transactions on Networking*, pages 629–640, October 1999.
- [5] Inter-Ministry Research Information Network. <http://www.imnet.ad.jp>.
- [6] R. Kawahara, K. Ishibashi, T. Hirano, H. Saito, H. Ohara, D. Satoh, S. Asano, and J. Matsukata. Traffic measurement and analysis in an ATM-based internet backbone. *Computer Communications*, 24:1508–1524, 2001.
- [7] R. Kawahara and N. Komatsu. A scalable IP traffic control method for weighted bandwidth allocation per flow. *IEICE Trans. Commun.*, E84-B(10):2815–2829, October 2001.
- [8] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic. *IEEE/ACM Transactions on Networking*, 2(1), February 1994.
- [9] D. Lin and R. Morris. Dynamics of Random Early Detection. In *ACM SIGCOMM*, September 1997.
- [10] Measurement and Analysis on the WIDE Internet. <http://www.wide.ad.jp/wg/mawi>.
- [11] T. Mori. A study on the difference of characteristics of self-similar traffic with similar Hurst parameter. In *IEICE General Conference*, March 2001. (in Japanese).
- [12] OC3mon/Coral. <http://www.caida.org/tools/measurement/coralreef>.
- [13] OPNET. <http://www.opnet.com>.
- [14] K. Papagiannaki, N. Taft, S. Bhattacharyya, P. Thiran, K. Salamatian, and C. Diot. A pragmatic definition of elephants in internet backbone traffic. In *ACM SIGCOMM Internet Measurement Workshop*, 2002.
- [15] K. Park, G. Kim, and M. Crovella. On the effect of traffic self-similarity on network performance. In *Proceedings of the SPIE International Conference on Performance and Control of Network Systems*, November 1997.
- [16] K. Park and W. Willinger. *Self-similar network traffic and performance evaluation*. Wiley Interscience, 2000.
- [17] I. Stoica, S. Shenker, and H. Zhang. Core-stateless fair queueing: Achieving approximately fair bandwidth allocations in high-speed networks. In *ACM SIGCOMM*, September 1998.
- [18] The Network Simulator - ns-2. <http://www.isi.edu/nsnam/ns>.
- [19] Widely Integrated Distributed Environment project. <http://www.wide.ad.jp>.
- [20] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson. Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking*, 5(1):71–86, 1997.
- [21] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker. On the characteristics and origins of internet flow rates. In *ACM SIGCOMM*, August 2002.