

A measurement study of correlations of Internet flow characteristics

KUN-CHAN LAN, JOHN HEIDEMANN^{a,b}

^a*National ICT Australia Ltd, Bay 15, Australian Technology Park, Eveleigh NSW 1430, Australia*

^b*Information Sciences Institute, University of Southern California, 4676 Admiralty Way, Marina Del Rey, CA 90292*

Abstract

Previous studies of Internet traffic have shown that a very small percentage of flows consume most of the network bandwidth. It is important to understand the characteristics of such flows for traffic monitoring and modeling purposes. Several prior researchers have characterized such flows using different classification schemes: by size as elephant and mice; by duration as tortoise and dragonfly; and by burstiness as alpha and beta traffic. However, it is not clear how these different definitions of flows are related to each other. In this work, using data recorded from two different operational networks, we study these “heavy-hitter” flows in four different dimensions, namely size, duration, rate and burstiness, and examine how they are correlated. This paper makes three contributions: First, we systematically characterize prior definitions for the properties of such heavy-hitter traffic. Second, based on our datasets, we observe that there are strong correlations between some combinations of size, rate and burstiness. Finally, we provide a plausible explanation for the observed correlations. We show that these correlations could be explained by transport and application-level protocol mechanisms.

Key words: Network Measurement, Heavy Hitter, Internet Traffic

1 Introduction

Recent studies have shown that a very small percentage of flows carry the majority of the bytes [26,10,2]. It is important to understand the properties of

Email address: Kun-chan.Lan@nicta.com.au, johnh@isi.edu (KUN-CHAN LAN, JOHN HEIDEMANN).

such traffic for traffic monitoring and modeling purposes. In this paper, we refer to such flows as “heavy-hitter” flows. By studying these heavy-hitter flows, one can understand a large portion of the overall traffic. Potential applications for employing such knowledge include anomaly and attack detection [15], scalable differentiated services [20,5], usage-based pricing and accounting [9,27]. However, while important and several efforts have looked at characterizations of heavy hitters by size [30,6,10,2,33,21], duration [4], and burstiness [25], there has been no systematic effort to study how these characteristics interact.

Several researchers previously have characterized Internet flows using different classification schemes: size (mice and elephant) [30,6,10,2,33,21], duration (dragonfly and tortoise) [4] and burstiness (alpha and beta traffic) [25]. While each of these studies provides different insights into understanding the characteristics of Internet flows, it is not clear how they relate to each other. For example, are most elephant flows long-lived (The answer depends on what types of links are used to transfer large file.)? On the other hand, are most long-lived flows due to the download of large files? Various applications can benefit from understanding the relationship between different characterizations of heavy-hitter flows. For instance, understanding of the relationship between long-lived and large-size flows might help one evaluate different pricing schemes (e.g. usage-based vs. duration-based). Knowledge of the correlation between high volume and bursty traffic could shed some insight into distinguishing large file transfer from malicious traffic.

Previously, Zhang et al. [32] showed that there is a strong correlation between flow size and rate. They hypothesized that users might have chosen the size of their transfer strongly based on the available bandwidth. In this work, based on datasets from two different sources, we propose another plausible explanation for the strong correlation between flow size and rate. While user behavior might have an effect on flows with a larger size, our data suggests that the strong correlation between size and rate might be better explained by *protocol reasons* for small- or medium-size flows. Our observation has some important implications for application and protocol design. For example, we show that, for small/medium flows, the strong correlation between rate and size is likely a pervasive artifact due to different timeout mechanisms. Such an observation might argue for the use of a larger packet size or a larger initial window to improve TCP performance (so that more data can be sent in one RTT before the timeout occurs).

The contribution of this paper is threefold. First, to our knowledge, our work is the first to systematically characterize the properties of these heavy-hitter flows (Section 5). Second, based on data collected from two different sources, we observe that there are strong correlations between some combinations of size, rate and burstiness. Finally, we provide a plausible explanation for the observed correlations. We show that these correlations can be explained by

transport and application-level protocol mechanisms (Section 6).

Note that, in this study, due to time constraints, the results of this paper are based on only a limited set of traces. However, since our data are recorded from two different levels of operational networks (one regional network and one backbone link), we believe that our work still provides some useful insights and a first step toward understanding the relationship between different characterizations of Internet flows.”

2 Flow characterization

We define a *flow* as an unidirectional series of IP packets with same source and destination addresses, port numbers and protocol number. Similar to previous studies [8,32], we use a 60 second timeout to decide that if an idle flow has terminated. In this work, we characterize and study Internet flows in four different dimensions, namely size, duration, rate and burstiness. Size is the total number of bytes sent in a flow (including headers). Duration is the time elapsed between the first packet and the last packet of a flow. Rate is size divided by duration. However, to the best of our knowledge, currently there is no consensus on definitions for burstiness. While all other characteristics of a flow are defined over the entire flow duration, burstiness is a property of *part* of the flow. Previous work on traffic self-similarity has identified this problem in characterizing burstiness at some timescale [14]. We propose three definitions of burstiness which are described shortly. We ignore very short flows, particularly flows with a duration of less than 100 millisecond, based on similar reasons to prior work [32].

Previous studies showed that distributions of flow sizes in the Internet traffic have a long tail. In this work, we focus on flows that are in the tail of the distribution and term them as heavy-hitter flows. Heavy-hitter flows typically account for only a small percentage of total flows but consume most of the network bandwidth. Specifically, we define and classify flows in four different dimensions: size (elephant and mice), duration (tortoise and dragonfly), rate (cheetah and snail) and burstiness (porcupine and stingray). We use a threshold-based scheme to define heavy-hitter flows in each category. We compute the mean plus three standard deviations of the sampled data to set the particular threshold¹. For example, an *elephant* flow is defined as a flow with

¹ We are in the process of looking at the results using median instead of mean, since mean might not be a good metric for heavy-tailed distributions. In addition, some distributions may not have well-defined second moments, or even first moments. Hence, we also look at using percentiles (e.g. the largest 1% of all flows) as breakpoints, as described later in Section 7.

a size larger than the mean plus three standard deviations of all flows.

Size (s): We define *elephants* as flows with a size larger than or x kB and *mice* as flows with a size less than or equal to x kB. For readability, we use the notation s to stand for *size* for the rest of the paper. For example, $flow_s$ means the size of a flow. Other notations (d , r and b) are used for duration, rate and burstiness respectively.

Duration (d): We define *tortoises* as flows with a duration longer than y minutes and *dragonflies* as flows with a duration less than or equal to y minutes.

Rate (r): we define *cheetahs*² as flows with a rate greater than z kB/sec and *snails* as flows with a rate less than or equal to z kB/sec.

Burstiness (b): In this work, three different definitions of burstiness are proposed. Our first definition of burstiness is based on the variation of traffic at a time-scale of T . Given a flow, we first divide it into bins b_i of duration T . Assuming that s_i is the number of bytes sent in b_i , *variance burstiness* of that flow is then defined as the standard deviation of all s_i .

The problem of using such a definition is that the result typically depends on the choice of T . In particular, a larger T tends to bias against small-size flows which have less data to sent in each T . However, small-size flows can still be bursty by sending most of their data in a very short period. Another weakness of this definition is the relationship between T and the flow duration. For flows shorter than T , variance is undefined, and boundary effects add error for flows shorter than 3–5 T . In addition, this definition does not consider network conditions, leading us to explore to alternative definitions.

Second, we consider *RTT burstiness*. We first define burst size as the number of bytes sent in each RTT of a flow. We then characterize RTT burstiness as the product of the mean burst size and the average RTT. That is,

$$burst_s \stackrel{\text{def}}{=} \text{bytes sent in each RTT}$$

$$burstiness \stackrel{\text{def}}{=} \text{mean}(burst_s) \times RTT_{avg}$$

Such a definition avoids the drawback of defining burstiness based on one particular fixed time scale. However, in practice, it is non-trivial to measure the RTTs of an uni-directional flow.

² We borrow the terms “elephant”, “mice”, “tortoise” and “dragonfly” from previous work. We use the term “cheetah” for its swiftness and “porcupine” for its sharp bristles which are visually similar to the shape of the bursts in the traffic.

Trace	Date	Duration	# of packets	# of flows	TCP	UDP
Los Nettos	Apr, 2003	2 hours	168 million	2.8 million	82%	15%
NLANR	Aug, 2002	20 minutes	42 million	0.1 million	94%	4%

Table 1

Characteristics of packet traces

Our third definition, *train burstiness*, defines a burst as a train of packets with a packet inter-arrival time less than a threshold t . Burst size is the number of bytes sent during each burst. Burst duration is the time elapsed between the first packet and the last packet of a burst. Burst rate is the burst size divided by burst duration, excluding any one-packet train. Inter-burst is the inter-arrival time between two bursts. Train burstiness is then defined as the product of mean burst rate and mean inter-burst. In other words,

$$burst \stackrel{\text{def}}{=} \text{packets with interarrival time} < t$$

$$burst_s \stackrel{\text{def}}{=} \text{bytes sent in each burst}$$

$$burst_d \stackrel{\text{def}}{=} \text{duration of a burst}$$

$$burst_r \stackrel{\text{def}}{=} \frac{burst_s}{burst_d}$$

$$burst_i \stackrel{\text{def}}{=} \text{gap between bursts}$$

$$burstiness \stackrel{\text{def}}{=} \text{mean}(burst_r) \times \text{mean}(burst_i)$$

We evaluate both *variance burstiness* and *train burstiness* and find that the results are qualitatively similar. For brevity, in this paper we present only the results based on train burstiness. We define *porcupines* as flows with burstiness greater than m kB and *stingrays* as flows with burstiness less than or equal to m kB.

In this paper, we present the results based on the analysis using $x=152$ kB, $y=12$ minutes, $z=101$ kB/s, $m=48.7$ MB and $t=1$ ms. We calculate the mean plus three standard deviations in each category to obtain these values (i.e. 152 kB, 12 minutes, etc.). In Section 7, we look at three other ways of defining heavy-hitters. We first define heavy-hitters as the top 1% of all flows. Second, we set the threshold as the cutoff point in the heavy-tailed distribution. We select the cutoff point by employing the *aest* test as proposed in [7]. Finally, we define heavy-hitters as the largest flows that together contribute 50% or more of the aggregated traffic. We find that the results do not change significantly in all three cases.

3 Related Work

Prior work has classified Internet flows based on several different schemes: size (elephant and mice), duration (tortoise and dragonfly) and burstiness (alpha and beta traffic). In our work, we study how these classifications relate to each other. Additionally, previous studies showed that there is a strong correlation between size and rate of Internet flows. They hypothesized that such a correlation between size and rate might be due to user behavior. In this paper, based on the data we collected, we demonstrate that the correlation between size and rate for small- or medium-size flows could be better explained by protocol reasons.

3.1 Elephant and Mice

While the sizes of most Internet flows are small, the majority of packets and bytes of Internet traffic are carried by a small percentage of large flows. This property persists across several levels of aggregation [30,6,10,2,33], and is known as the “elephant and mice phenomenon”.

Several previous studies tried to identify *elephant* flows. Estan et al. [8] defined *elephant* as any flow whose rate that is larger than 1% of the link utilization. Papagiannaki et al. [21] proposed a more sophisticated two-feature classification scheme to identify elephant flows. According to their definition, flows are characterized as “elephant” based on both their volume and their persistence in time. Note that the definition of “flow” in Estan’s work is similar to ours (as described in Section 2), but the flow granularity chosen by Papagiannaki et al. is at the network prefix level.

Our definition of *elephant* is closer to Estan’s work. We define *elephant* flows as flows with a size larger than the mean plus three standard deviations of the sampled data. Specifically,

Prior: elephant := $flow_s > 1\%$ of link bandwidth

Ours: elephant := $flow_s > (\text{mean} + 3 * \text{std})$ of all flows

Note that our work does not focus on how to choose the criterion for defining an elephant. Instead, given a fixed criterion, we focus on the correlation between elephant flows and other dimensions (i.e. duration, speed and burstiness). However, to understand if different choices of the threshold would affect our results, we also look at the effects from using different criteria in Section 7. We find that our results do not significantly change due to different choices of the threshold

3.2 Tortoise and Dragonfly

Brownlee et al. [4] studied Internet flows from a different aspect. They classified Internet flows based on their durations. They found that 45% of flows have a duration of less than two seconds (dragonflies), and less than two percents of the flows last longer than fifteen minutes and carry more than 50% of the total bytes on a link (tortoises).

Our definition of long-lived flows are flows with a duration larger than the mean plus three standard deviations of the sampled data. That is,

Prior: tortoise := $flow_d > 15\text{minutes}$

Ours: tortoise := $flow_d > (\text{mean} + 3 * \text{std})$ of all flows

Additionally, we look at the other properties of these long-lived flows (i.e. size, rate and burstiness).

3.3 Alpha and Beta traffic

Sarvotham et al. [25] showed that traffic bursts typically arise from just a few high-volume connections that dominate all the others. They named such flows as *alpha* traffic and define them as any flow whose peak rate exceeds certain threshold. Specifically, they identified the connection(s) that transmits the largest number of bytes in each 500ms time bin and labeled it as an *alpha* flow if its rate exceeds the mean (Agg_μ) plus three standard deviations (Agg_{std}) of the aggregate traffic. They defined the remaining flows as *beta* traffic and found that the beta component of the aggregate traffic carries the same fractal scaling exponent as the aggregate traffic.

In this paper, we propose three different definitions of burstiness, namely *variance burstiness*, *RTT burstiness*, *train burstiness*, as described in Section 2. We have evaluated both variance burstiness and train burstiness and found that the results are qualitatively similar. Our definition of bursty flows are flows with a burstiness larger than the mean plus three standard deviations of the sampled data. In other words,

Prior: alpha := $burst_{peak} > Agg_\mu + 3 * Agg_{std}$

Ours: porcupine := $flow_b > (\text{mean} + 3 * \text{std})$ of all flows

Surprisingly, as shown later in Section 5, our results are consistent with the observation from Sarvotham’s work (where they found that most bursty flows are due to transfer of large files over fast links) even when we define *burstiness* differently.

3.4 Flow analysis

Previously Zhang et al. [32] looked at flows with a duration longer than 30 seconds and found that there is a strong correlation between flow size and rate. They hypothesized that, for large flows, the strong correlation between size and rate might be due to user behavior. In other words, users tend to choose the size of their transfer based on the available bandwidth. While user behavior might introduce some correlation between rate and size, we find that the strong correlation between size and rate for small- or medium-size flows might be better explained by protocol reasons. Additionally, we show that using flow duration as a metric to separate large- and small-size flows could be misleading. As described later in Section 6.1, our data suggests that most of the flows longer than 30 seconds actually only have a medium or small size.

3.5 Multi-dimensional traffic characterization

Estan et al. [9] proposed a traffic characterization scheme that automatically groups traffic into minimal clusters of conspicuous consumption. They analyzed traffic along multiple different dimensions (source address, destination address, protocol, source port and destination port) at once, and then compressed the results into a concise report. While our work is also based on a multi-dimensional classification scheme, we focus on understanding the relationship between different dimensions. Additionally, we look at a different multi-dimensional space (size, duration, rate and burstiness). One possible extension of our work is to apply similar technique like theirs on the multi-dimensional space we study to detect interesting/important traffic clusters.

4 Traces

The datasets we utilize in this study are from two different sources. The first set of traces were collected at Los Nettos [19], a regional area network in Los Angeles. Los Nettos has peering relationships with several ISPs and the LA-Metropolitan Area Exchange, and serves a diverse clientele that includes academic institutes and corporations around the Los Angeles area. The second set of traces were from the NLANR site [18]. The NLANR traces were previously collected on an Abilene OC48 backbone link that lies between Indianapolis and Cleveland. The characteristics of the traces are summarized in Table 1. Although the duration of NLANR trace is much shorter than Los Nettos trace, its mean flow size is significantly larger. Because the short duration of NLANR trace will inevitably introduce a bias against long-lived flows,

Category	% of no. of bytes	% of no. of flows
Los Nettos		
Elephant	71%	1%
Tortoise	43%	4%
Cheetah	16%	2%
Porcupine	39%	0.9%
NLANR		
Elephant	82%	4%
Tortoise	45%	4%
Cheetah	36%	2%
Porcupine	40%	1%

Table 2

Fraction of Internet traffic for each category in terms of bytes and number of flows

our results are mainly based on Los Nettos traces. We utilize NLANR traces for comparison and validation. Note that Los Nettos data has a larger percentage of UDP traffic due to the presence of a DNS root name server. Figure 1 shows the distributions of different flow metrics in Los Nettos data. The scaling exponents α for distributions of flow size, duration, rate and burstiness are 1.2, 1.8, 1.3 and 1.5 respectively.

5 Relationships between different characterizations of heavy-hitter flows

In this section, based on the data we collected, we present some properties of different characterizations of heavy-hitter flows.

As shown in Table 2, although accounting for only a very small percentage of total flows, these heavy-hitter flows carry a large portion of network traffic. In particular, porcupine flows carry almost as much traffic as tortoises although they account for less flows.

Table 3 shows relationships between different characterizations of heavy-hitter flows. The first column in the upper table indicates that, in Los Nettos traces, “20% of elephants are also tortoises, 7% of elephants are cheetahs, and 19% of elephants are porcupines”. The last column shows that about 68% of porcupines are also elephants, which indicates that bursty flows are strongly correlated with large-size flows. The third column shows that cheetah flows are less correlated with other categories in Los Nettos traces when comparing against

given				
expect	Elephant	Tortoise	Cheetah	Porcupine
Los Nettos				
Elephant	-	6%	3%	68%
Tortoise	20%	-	0.007%	8%
Cheetah	7%	0.004%	-	3%
Porcupine	19%	1%	4%	-
NLANR				
Elephant	-	29%	72%	75%
Tortoise	17%	-	12%	9%
Cheetah	39%	8%	-	80%
Porcupine	28%	5%	57%	-


Table 3
Relation between elephant, tortoise, cheetah and porcupine flows

Rank	Elephant	Tortoise	Cheetah	Porcupine
1	web (67%)	DNS (51%)	web (53%)	web (71%)
2	kazaa (5%)	web (15%)	DNS (28%)	smtp (10%)
3	telnet (3.5%)	telnet (9.1%)	ftp (5%)	ftp (6%)
4	gnutella (2%)	ftp (5%)	smtp (3.3%)	nntp (2.1%)
5	nntp (2%)	smtp (4.5%)	WinMX (1.3%)	pop (1.3%)


Table 4
Top five applications in terms of total number of flows in different categories in Los Nettos trace

NLANR traces. As described in Section 4, Los Nettos data has a larger percentage of DNS traffic due to the presence of a root name server. A large portion (about 60%) of cheetah flows in Los Nettos data are comprised of small bursts of DNS traffic. The mean size of cheetah flows consequently is significantly larger in NLANR data than in Los Nettos traces. As a result, there are more cheetah flows that are also elephants in NLANR traces than in Los Nettos data.

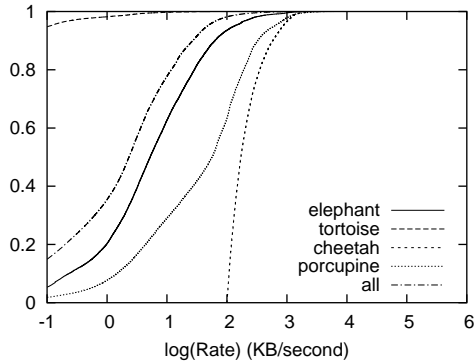
Table 4 shows five of the most popular applications in each category of Los Nettos data. We identify the applications based on their well-known port numbers (e.g. port 80 for web traffic). Note that this approach might introduce bias against some applications such as P2P traffic which commonly uses random port numbers to avoid the blocking of firewall. Overall, web and P2P applications account for most of Internet traffic in terms of the number of bytes,



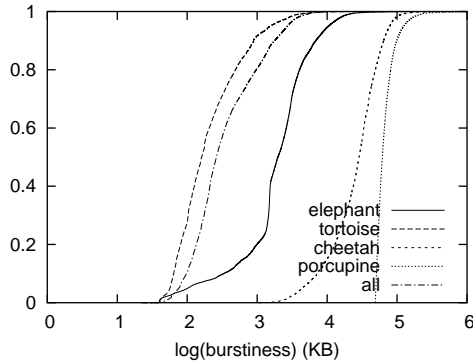
(a) Distributions of flow size



(b) Distributions of flow duration



(c) Distributions of flow rate



(d) Distributions of flow burstiness

Fig. 1. The relationship between different characterizations of heavy-hitter flows

which is consistent with prior work [16]. In particular, web traffic accounts for most of the fast and bursty traffic. More than 50% of long-lived flows are DNS traffic (We classify any flow that uses port 53 as DNS traffic, and do not distinguish zone transfers from standard queries. A closer examination of our traces, however, shows that most of these long-lived flows are comprised of DNS zone transfers.). Surprisingly, DNS traffic is also responsible for significant portion of high-rate traffic. A closer look at our traces shows that a large number of DNS flows consist of burst of packets due to repeated DNS queries originated from the same host. Similar results were also reported in a previous study of CAIDA [31]. Note that some of elephant flows are contributed by telnet traffic. Such an observation suggests that interactive traffic like telnet could still consume significant amount of network bandwidth due to its persistence in time.

Figure 1 shows the relationship between different characterizations of flows in Los Nettos data. The results for NLANR traces are similar. For brevity, we do not show the same plots for NLANR data here.

First, we look at the flow size distribution for different classifications of flows, as shown in Figure 1(a). One interesting insight that one can infer from Figure 1(a) is the origin of long-lived flows. There are two possibilities for the cause of long-lived flows. The first case is due to the user/application behavior. For example, a long-live flow might occur when an application repeatedly sends some amount of traffic and then pauses for a long period (such as periodic DNS updates or telnet). Another possibility is the transfer of a big file over slow links. Based on our traces, we find that the former explanation is more plausible. As shown in Figure 1(a), only about 6% of tortoises are flows with a size greater than 100 kB and around 80% of tortoises are smaller than 10 kB, which does not support the second case. Hence, we conjecture that the majority of long-lived flows in our traces are most likely due to application/protocol reasons. Furthermore, about 70% of cheetah flows are smaller than 10 kB, which indicates that a large number of fast flows contain only a small burst of packets (such as the bursty DNS queries described previously). Finally, the distributions of porcupine flows and elephant flows share some similarities, which again suggests that they might have some correlation.

Next, we look at the distribution of flow durations for different types of flows. As shown in Figure 1(b), more than 70% of Internet flows have a duration of less than 10 seconds, which is consistent with prior work [4] that reported that most Internet flows are short-lived. More than 95% of cheetah flows are short (less than one second), which confirms that most cheetah flows consist of just a small burst of packets. About 50% of elephant flows have a duration longer than two minutes and 20% of elephants last longer than 15 minutes, which suggests that most elephant flows are long-lived. Note that the last 10% of elephant flows have a similar duration, which is due to the boundary effect of our fixed-length traces. Finally, about 65% of porcupine flows have a duration less than 10 seconds and more than 95% of porcupines last less than 2 minutes. Since most porcupine flows are also elephants, this observation suggests that most of the bursty traffic might be due to the transfer of large files over fast links. Note that our observation is consistent with prior work [25] even though we define “bursty flow” differently (as described in Section 3.3).

Figure 1(c) shows the flow rate distribution for different types of flows. About 80% of porcupine flows have a rate greater than 10 kB/sec and 30% of porcupines have a rate greater than 100 kB/sec, which suggests that most bursty flows are also fast. Around 30% of elephant flows are faster than 10 kB/sec and about 5% of elephants are faster than 100 kB/sec, which implies that most elephant flows are not fast. Lastly, we find that around 80% of Internet flows have a rate less than 10 kB/sec.

Finally we look at the distribution of flow burstiness for different types of flows, as shown in Figure 1(d). Based on our definition of burstiness, tortoises, elephants and most Internet flows are comparatively less bursty than

Category	large-size	long-lived	fast	bursty
Elephant	Y	Y	N	N
Tortoise	N	Y	N	N
Cheetah	N	N	Y	Y
Porcupine	Y	N	Y	Y

Table 5
Taxonomy of heavy-hitter traffic

metrics	Correlation coefficient	
	Los Nettos	NLANR
(rate,burstiness)	0.83 *	0.82 *
(size,rate)	0.81 *	0.87 *
(size,burstiness)	0.80 *	0.77
(size,duration)	0.21	0.23
(duration,burstiness)	-0.17	-0.07
(duration,rate)	-0.32	-0.04

Table 6
Correlation between different categories

porcupines and cheetahs. Specifically, there are only around 5% of elephant flows are burstier than 10MB. More than 80% of Internet flows and more than 90% of tortoise flows are less bursty than 1MB. More than 80% of cheetah flows are burstier than 10MB, although most cheetahs only consist of a small number of packets.

Table 5 shows a taxonomy that characterizes the “heavy-hitter” traffic. In summary, elephant flows are long-lived, but neither fast nor bursty. Tortoise traffic is slow and not bursty. Individual tortoise flows in general do not use up a lot of network bandwidth although aggregatively they consume significant amount of bandwidth, as shown previously in Table 2. Cheetah flows are typically small but bursty. Finally, porcupine flows are likely due to the download of big files over fast links. These results obviously depend on our definitions of heavy-hitter traffic. We also look at other ways of defining heavy-hitters, as later described in Section 7, and find that the results do not change significantly.

6 Origin of correlation between different flow statistics

Zhang et al. [32] showed that there is a strong correlation between flow rate and size. Motivated by their work, in this paper we study the physical explanation for the observed phenomena of correlations between different flow statistics. Table 6 shows six pairs of correlations: rate and size, rate and duration, rate and burstiness, size and duration, size and burstiness, and duration and burstiness. We computed correlations of the log of these data because of the large range and uneven distribution. To compute the correlation between different flow statistics, we use rank-based Kendall’s τ method, which is less sensitive to outliers and non-normality than the standard Pearson estimate [28]. As shown in Table 6, we find that size, rate and burstiness are strongly correlated. In this section, based on our data, we present some plausible explanation for the reason of strong correlations between flow size, rate and burstiness.

Note that one might expect that there is a stronger correlation between size and duration than what is shown in Table 6. Since small flows account for more data points in our traces, one possibility for the observed weak correlation between size and duration might be that our results are bias toward small flows. To verify such a hypothesis, we look at the correlation between size and duration for large-size flows alone. However, we do not find a strong correlation between size and duration for large-size flows. One plausible reason could be that, for large-size flows, users might choose the size of their transfer based on the link speed, as suggested by prior work [32]. For example, one might decide not to download big files (or abort after a long wait) when browsing the web via a slow modem link. Hence, most of larger flows might tend to be seen on faster links. As a result, a larger-size flow might not have a longer duration if such a flow is sent over a faster link.

6.1 High correlation between rate and size

Previous work [32] showed that there is a strong correlation between flow rate and size. They hypothesized that the observed strong correlation is due to user behavior: users choose the size of their transfer based on available bandwidth. In this section, based on our data, we provide another plausible explanation for the observed correlation. We suspect that, while user behavior could have some effect on large-size flows, the origin of the observed correlation might be better explained by transport and application-level protocol mechanisms for small- or medium-size flows.

To systematically investigate the cause of correlation between flow size and

types	Correlation coefficient	
	Los Nettos	NLANR
size less than 10K	0.17	0.41
size between 10K and 100K	0.13	0.47
size greater than 100K	0.16	0.32
duration greater than 1 sec (ALL)	0.57	0.79
duration greater than 5 sec (ALL)	0.65	0.81 *
duration greater than 30 sec (ALL)	0.81 *	0.87 *
duration greater than 1 sec (TCP)	0.71	0.81 *
duration greater than 5 sec (TCP)	0.83 *	0.87 *
duration greater than 30 sec (TCP)	0.92 *	0.96 *
duration greater than 1 sec (UDP)	0.34	0.70
duration greater than 5 sec (UDP)	0.61	0.77
duration greater than 30 sec (UDP)	0.74	0.82 *

Table 7

Correlation between size and rate for different protocol, flow sizes and duration

rate, we first group flows based on their protocols, size and duration, as shown in Table 7.

As shown in Table 7, there is a strong correlation between rate and size for flows longer than 30 seconds³ (the correlation coefficients are greater than 0.8 for both traces). We do not see similar results for flows with a larger size (for example, the correlation coefficient for flows with a size larger than 100K is only 0.16 for Los Nettos traces). However, if the strong correlation between size and rate is due to that users choose the file to transfer based on the available bandwidth, as suggested by prior work, we expect to see a strong correlation between size and rate for large-size flows as well. After taking a closer look, surprisingly, we find that most of the flows longer than 30 seconds actually only have a medium or small size. As shown in Figure 2, 70% of such flows have a size of less than 10 kB and 90% of them are smaller than 60 kB.

While indicating that there is a strong correlation between size and rate, Table 6, however, does not provide enough information for understanding the cause of such a correlation. To visually examine at what range of size and rate where this correlation arises, we plot size against rate on a density plot. Figure 3 shows the density plots of TCP flows for Los Nettos and NLANR

³ As shown in the second row, the correlation between flow size and rate becomes stronger as we increase the threshold from 1 to 30 seconds




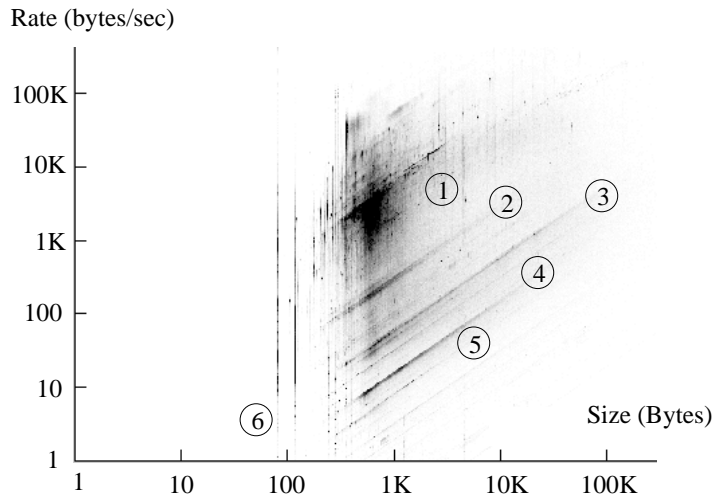
Fig. 2. Distribution of flow size for flows with a duration longer than 30 seconds in Los Nettos trace

traces. To generate each graph, the area is divided into a 1000x1000 grid. We then place each of the millions of flows from the traces into a grid cell, sum the number of flows in that cell and map it to a gray-scale value, with cells from 0 to 8192 flows representing white to pure black. The density plot therefore highlights which combinations of size and rate are most “popular”. In other words, a darker point on the plot indicates that there are more flows with that particular combination of size and rate⁴.

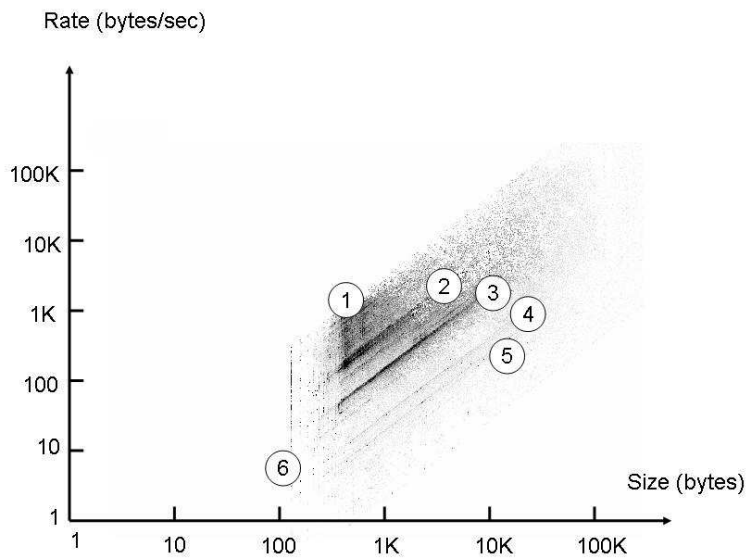
There are a few distinct features on both plots of Figure 3: several slanted bands on the right (region 2,3,4,5) and a few vertical lines on the left (region 6). The diagonal bands on the right indicate that the rate of flows is proportionally increasing to the size at a log scale (i.e. positively correlated). The vertical lines suggest that these are flows with the same size but different rates. Finally, there are more and also darker points in region 1 of Figure 3, which indicates that there are more TCP flows fall in this region (with sizes between 1 kB and 10 kB and rates from 1K/sec to 10K/sec).

A closer look at the flows in each diagonal band indicates that these flows have similar durations. The reason that these flows have similar durations can be explained by the protocol mechanism. Figure 4 demonstrates a typical flow in region 2 of Figure 3(a). While the actual transfer of a flow requires only a few hundred milliseconds, the timeout for SYN retransmission stretches the flow duration to about three seconds [3]. Flows in region 3, 4 and 5 of Figure 3(a) also have similar flow durations respectively. These similar durations are due to different timeout mechanisms. Specifically, most of the flows in region 3 last about 15 seconds and are mainly due to HTTP persistent connection timeout [11]. The flows in region 4 last about 60 seconds and flows in region 5 last about two minutes. The durations of these flows are mainly stretched out

⁴ Note that since there are more flows in Los Nettos traces than in NLANR traces, for presentation purposes, we reduce the number of flows required to represent the same gray-scale by a factor of 8 for NLANR traces. In other words, for NLANR traces, cells with from 0 to 1024 flows are represented by white to pure black.



(a) Los Nettos trace



(b) NLANR trace

Fig. 3. Density plots of size and rate for TCP flows (at log-log scale)

```

28.477 10.0.0.1.2355 > 10.0.0.2.80: S 12193306:12193306(0) win 8192
31.460 10.0.0.1.2355 > 10.0.0.2.80: S 12193306:12193306(0) win 8192
31.784 10.0.0.1.2355 > 10.0.0.2.80: . ack 3335637810 win 8760
31.792 10.0.0.1.2355 > 10.0.0.2.80: P 0:758(758) ack 1 win 8760
31.852 10.0.0.1.2355 > 10.0.0.2.80: . ack 124 win 8638
31.852 10.0.0.1.2355 > 10.0.0.2.80: R 12194065:12194065(0) win 0

```

Fig. 4. retransmission timeout in a small flow

```

13.893 10.0.0.1.1183 > 10.0.0.2.80: S 167211179:167211179(0) win 16384
16.878 10.0.0.1.1183 > 10.0.0.2.80: S 167211179:167211179(0) win 16384
23.119 10.0.0.1.1183 > 10.0.0.2.80: S 167211179:167211179(0) win 16384

```

Fig. 5. SYN retransmission in a small flow

```

8.865 1.0.0.1.12474 > 1.0.0.2.4308: S 3856:3856(0) ack 53513 win 9152
8.925 1.0.0.1.12474 > 1.0.0.2.4308: F 1:1(0) ack 1 win 9152
8.977 1.0.0.1.12474 > 1.0.0.2.4308: . ack 2 win 9152

```

Fig. 6. TCP connection with no data




Fig. 7. Distribution of TCP flow duration of Los Nettos traffic

by the TCP TIME_WAIT delay (2MSL wait). RFC 793 [23] specifies the MSL as 2 minutes. However, common implementation values typically range from 30 seconds to 2 minutes [29]. In summary, each diagonal band (region 2,3,4,5) consists of a group of flows with similar flow durations but varying amount of data. The spacing between different diagonal bands is due to variable flow durations which in turns are caused by different protocol mechanisms.

Another distinct feature of the plots is the existence of vertical lines on the left (region 6) for both Los Nettos and NLANR traces. These vertical lines mainly consist of flows with only three or four packets. They account for 8% of total flows in NLANR traces, and 9% in Los Nettos traces. After manually examining a large number of such flows in both traces, we find that they mainly consist of two types of flows.

The first type of flows, as shown in Figure 5, consist of a sequence of SYN retransmissions. We suspect that these flows are either some particular implementation of TCP (that stops re-transmitting after sending three SYN packets) or some kind of port scanning. The second type of flows, as shown in Figure 6, only transmit SYN and FIN with no data packets in between. We

suspect that these flows might be due to some kind of port scanning. Finally, as shown in Figure 7, more than 5% of flows have durations less than a couple of seconds and account for the concentration of flows in region 1. A detailed examination shows that these flows are normal web traffic.


The density plot of NLANR traces, as shown in Figure 3(b), is similar to that of Los Nettos traces. Note that we do not see a significant number of flows with SYN retransmission (region 2 in Figure 3(a)) in NLANR traces as in Los Nettos traces.

There is a strong correlation between flow rate and size for UDP traffic as well, as shown in Table 7. We also look at the density plots of UDP traffic for both Los Nettos and NLANR traces, as shown in Figure 8. The majority of UDP traffic is contributed by DNS flows (which account for 78% of all UDP flows in Los Nettos traces and 81% in NLANR traces). A common feature between Figure 8(a) and Figure 8(b) is the existence of several diagonal bands. Similar to the analysis for TCP flows, we find that these bands are also due to flows with similar durations and varying amount of traffic.


The diagonal line on the bottom (region 1 on both plots) consists of long-lived server-to-server DNS flows that last across the entire duration of our traces. The diagonal lines on the top of both plots consist of flows with durations ranging from 1 to 9 seconds. These flows are mainly contributed by DNS flows with repeated transmissions. Consistent with prior work [31], such flows account for a significant percentage of DNS traffic in our traces (52% in Los Nettos traces and 38% in NLANR traces). Figure 9 demonstrates one of such flows. The duration of such flows is a function of the number of retransmission and the length of timeout. Since DNS retransmission timeouts are typically some fixed values [17], the durations of these flows resultingly concentrate on certain lengths.

There are some vertical lines on the upper left-hand side of the plot for Los Nettos traces (region 2 in Figure 8(a)). A closer look shows that these small flows mainly consist of probe packets of game traffic. Finally, there are some dark dots on the bottom of Los Nettos plot (region 3 in Figure 8(a)). A careful examination shows that these flows are contributed by a number of extraordinarily busy sources sending repeated “A?” queries. These DNS flows account for about 4% of total DNS queries in our traces. A similar observation of such busy sources was also previously reported by CAIDA [31].

The flows in region 2 (the vertical line) and region 3 (the dark slanted line) of Figure 8(b) mainly consists of probe packets of Kazaa traffic. The typical duration for flows in region 3 is about 9 seconds.



(a) Los Nettos trace



(b) NLANR trace

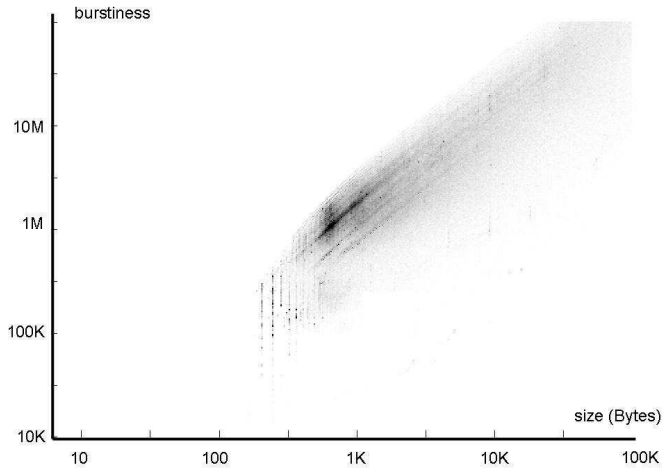
Fig. 8. Density plots of size and rate for UDP flows

```

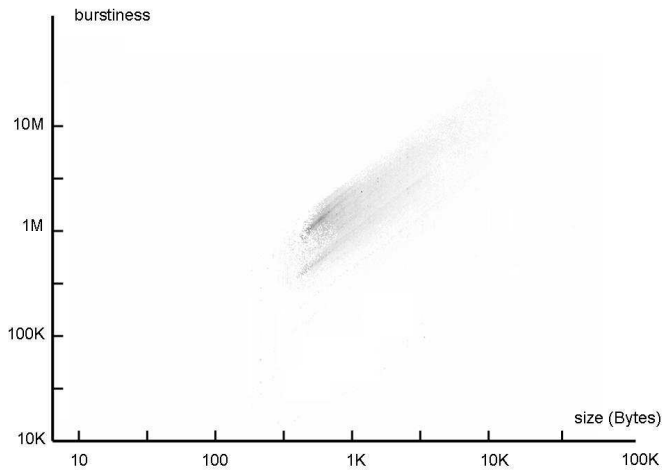
3.019 1.0.0.1.711 > 1.0.0.2.53: 643 PTR? 7.3.2.1.in-addr.arpa. (42)
4.046 1.0.0.1.711 > 1.0.0.2.53: 644 PTR? 7.3.2.1.in-addr.arpa. (42)
6.038 1.0.0.1.711 > 1.0.0.2.53: 645 PTR? 7.3.2.1.in-addr.arpa. (42)
10.037 1.0.0.1.711 > 1.0.0.2.53: 646 PTR? 7.3.2.1.in-addr.arpa. (42)

```

Fig. 9. DNS repeated query



(a) Los Nettos trace



(b) NLANR trace

Fig. 10. Density plots of size and burstiness for TCP flows

6.2 High correlation between burstiness and size

Table 6 shows that flow size and burstiness are also highly correlated. In this section, using similar analysis as described in Section 6.1, we show that the correlation between size and burstiness can also be explained by protocol reasons.

The density plots of size versus burstiness for TCP flows are shown in Figure 10. A common feature in Figure 10 for both traces is the existence of

```

6.072 10.0.0.1.1300 > 10.0.0.2.80: S 647474:647474(0) win 4288
6.145 10.0.0.1.1300 > 10.0.0.2.80: . ack 55646058 win 4288
6.146 10.0.0.1.1300 > 10.0.0.2.80: P 0:645(645) ack 1 win 4288
9.061 10.0.0.1.1300 > 10.0.0.2.80: P 0:645(645) ack 1 win 4288
9.116 10.0.0.1.1300 > 10.0.0.2.80: . ack 50 win 4240
9.117 10.0.0.1.1300 > 10.0.0.2.80: F 645:645(0) ack 50 win 4240

```

Fig. 11. Flow A

```

2.221 10.0.0.1.3784 > 10.0.0.2.80: S 972052848:972052848(0) win 3392
2.261 10.0.0.1.3784 > 10.0.0.2.80: . ack 3276864582 win 50000
2.262 10.0.0.1.3784 > 10.0.0.2.80: P 0:297(297) ack 1 win 50000
5.210 10.0.0.1.3784 > 10.0.0.2.80: P 0:297(297) ack 1 win 50000
5.258 10.0.0.1.3784 > 10.0.0.2.80: . ack 1588 win 50000
5.405 10.0.0.1.3784 > 10.0.0.2.80: . ack 2921 win 49666
5.406 10.0.0.1.3784 > 10.0.0.2.80: . ack 3376 win 50000
5.610 10.0.0.1.3784 > 10.0.0.2.80: F 297:297(0) ack 3376 win 50000

```

Fig. 12. Flow B




Fig. 13. Bandwidth-time plots of flow A and B

diagonal lines. Similar to the observation from previous section, we find that each diagonal line consists of a group of flows with similar duration.

Figure 11 and Figure 12 show two typical flows in the darkest diagonal line at the center of Figure 10(a) (region 1 on the plot). The corresponding time series plots of both flows are shown in Figure 13. The durations of both flows are stretched out by TCP retransmission timeout to around three seconds. As a result, the burstiness of the flow depends on the volume of traffic it transmits. Flow A is burstier than flow B because that flow A has a larger HTTP transfer (Specifically, flow A has a size of 1530 Bytes, a duration of 3.04 second, a rate of 503 Bytes/sec and a burstiness of 2055 kB, while flow B has a size of 914 Bytes, a duration of 3.39 second, a rate of 267 Bytes/sec and a burstiness of 950 kB.). Note that there are a few vertical lines at the left-hand side of Figure 10(a) (region 2 on the plot). These vertical lines are mainly due to SYN retransmissions and probing packets as described previously in Section 6.1.

Similarly, the correlation between rate and burstiness in Table 6 could also be explained by the above reasoning. For brevity, we do not show the corresponding plots here.

expect	given			
	Elephant	Tortoise	Cheetah	Porcupine
Elephant	-	1%	0.5%	83%
Tortoise	19%	-	0.003%	3%
Cheetah	9%	0.001%	-	2%
Porcupine	22%	1%	9%	-

Table 8

Relation between elephant, tortoise, cheetah and porcupine flows when defining heavy-hitters as the largest 1% of the flows

7 Sensitivity of results

Our work does not focus on choosing the criterion for defining *elephant*, *tortoise*, *cheetah* and *porcupine*. Instead, given a fixed criterion, we focus on the statistical properties of heavy-hitter flows in different dimensions, namely size, duration, rate and burstiness. However, to understand if different choices of thresholds (as defined Section 2) would affect our results, we investigate the effects from using three different criteria in this section.

First, instead of using the thresholds defined in Section 2, we define heavy-hitters traffic as the largest 1% of all flows. After applying such a definition to Los Nettos data, the resulting thresholds are equivalent to the choices of $x = 110, y = 23, z = 368, m = 56742$ (x, y, z, m are defined in Section 2). The second criterion that we employ is to apply the *aest* test (previously proposed by Crovella et al. [7]) to our data, and choose the threshold as the cutoff point in the heavy-tailed distribution (the scaling exponents α for size, duration, rate and burstiness are 1.2, 1.8, 1.3 and 1.5 respectively). The resulting thresholds are equivalent to the choices of $x = 138, y = 20, z = 121, m = 50111$. Finally, we define heavy-hitters traffic as the largest flows that together carry 50% or more of the total bytes. The resulting thresholds are equivalent to the choices of $x = 145, y = 21, z = 113, m = 51887$. As shown in Table 8, Table 9 and Table 10, although the numbers are slightly different, overall the results are similar to Table 3.

8 Discussion

Modeling and simulating Internet traffic is difficult due to its scale, heterogeneity and dynamics [13]. It is important to understand the causal root of traffic characteristic so that one can determine what is fundamental and what is just an artifact. Knowledge of fundamental correlations of traffic charac-

	given			
expect	Elephant	Tortoise	Cheetah	Porcupine
Elephant	-	3%	2%	74%
Tortoise	19%	-	0.006%	4%
Cheetah	8%	0.006%	-	3%
Porcupine	20%	1%	5%	-

Table 9

Relation between elephant, tortoise, cheetah and porcupine flows when defining heavy-hitters as flows beyond the cutoff point in the heavy-tailed distribution

	given			
expect	Elephant	Tortoise	Cheetah	Porcupine
Elephant	-	2%	0.8%	59%
Tortoise	22%	-	0.001%	3%
Cheetah	15%	0.01%	-	9%
Porcupine	24%	2%	4%	-

Table 10

Relation between elephant, tortoise, cheetah and porcupine flows when defining heavy-hitters as flows that consume 50% of total traffic

teristics allows one to limit the number of cases needed to be considered. Complementary to previous studies that characterized Internet flows based on different metrics (e.g. size, duration, etc.), this paper emphasizes on understanding the relationship between different characterizations of flows in order to get a better insight of traffic dynamics.

In this work, we show that some of traffic metrics are strong correlated (e.g. size and burstiness) while the others are relatively independent of each other (e.g. size and duration). Our results have some important implications in protocol design and network modeling/simulation.

As implied by the last row of Table 5, by paying more attention to bursty flows, one could captures most of high-rate and large-size traffic. Based on this observation, it seems reasonable to explicitly take burst traffic into account in protocol and router design. The recent proposal of Optical Burst Switching [24] is such an example. Additionally, such an insight can be applied to reduce the complexity of simulation. Instead of modeling and simulating traffic in details, one can focus on bursty traffic and still capture most of the traffic dynamics required in a large simulation. Furthermore, one can utilize the burstiness of traffic as another metric to identify elephant flows in addition to the use of size and duration [21].

On the other hand, as discussed in Section 6.1, using the duration of a flow as an indication of the volume of traffic sent could be misleading in some cases. Flow size and duration might need to be treated as different and independent dimensions.

In this work, we study characteristics of heavy-hitter flows in four different dimensions: size, duration, rate and burstiness. However, based on the root of traffic characteristics, one can still identify different classes of traffic within each individual metric. For example, in Section 5 we show that long-lived flows can be due to either transfer of large files or the effect of application/user behavior. Instead of treating them indiscriminately as one single class of long-lived flows, it seems more reasonable to separate them as different classes of traffic for network modeling and traffic monitoring purposes.

9 Future Work

In this work, we show that the durations of a large number of small/medium flows are stretched out by various protocol timeout mechanisms. The cause of timeouts might be due to either application/user behavior or network congestion. It is important to characterize these timeouts and understand their prevalence for performance and modeling purposes. For example, we observed a significant number of packet retransmissions in our traces. It would be interesting to understand what fractions of them are due to network congestion, software flaws [12], malicious attack, etc. Some recent work [1] based on measurements from NIMI [22] has shown that a significant number of TCP retransmissions in their data are not caused by congestion-induced packet losses.

Prior work [32] showed that the most frequent cause for limiting the rate of a flow is network congestion. Our data suggests that the origin of some long-lived flows are likely due to application behavior instead of download of big files. In addition, we confirm that most bursty traffic might be due to transfer of large files over fast links. However, relatively little study has been done to understand the cause of burstiness in Internet traffic. The burstiness of a flow can be due to either application/protocol behavior or network congestion. For TCP traffic, one way to infer the occurrence of queuing is to compare the observed burstiness of the flow with its congestion window size. As future work, we plan to study what fractions of burstiness in Internet traffic are due to network congestion.

In this study, due to time constraints, the results of this paper are based on a limited set of traces. We plan to collect more traces from other places, particularly traces from backbone links of a large ISP to further compare and

validate our results.

10 Conclusion

Previous studies of Internet traffic have shown that a small percentage of flows carry most of the network traffic. It is important to understand the characteristics of such flows for traffic monitoring and modeling purposes. Several prior studies have characterized such flows using different definitions: elephant and mice, tortoise and dragonfly, and alpha and beta traffic. However, it has not been clear how these different classifications of flows relate to each other. In our work, using data from different traffic sources, we study these “heavy-hitter” traffic in four different dimensions, namely size, duration, rate and burstiness, and examine how they are correlated. We first systematically characterize prior definitions for the properties of these heavy-hitter traffic. In our datasets, we observe that a significant percentage of long-lived flows are comprised of DNS traffic. Our data suggests that the bursty traffic is likely due to the transfer of big files over fast links, which is consistent with the observation from previous work even when we define bursty flows differently. We also observe that there are strong correlations between flow size, rate and burstiness. Additionally, we show that using the duration of a flow as an indication of the volume of traffic sent could be misleading in some cases. Flow size and duration might need to be treated as different and independent dimensions. Finally, we present a plausible physical explanation for the observed correlations between size, rate and burstiness.

References

- [1] M. Allman. Estimating loss rates with tcp, under submission. <http://roland.grc.nasa.gov/mallman/papers/least-submit.ps>.
- [2] Supratik Bhattacharyya, Christophe Diot, Jorjeta Jetcheva, and Nina Taft. Pop-level and access-link-level traffic dynamics in a tier-1 POP. In *Proceeding of ACM SIGCOMM Internet Measurement Workshop 2001*, pages 39–54, San Francisco Bay Area, November 2001.
- [3] Robert Braden. Requirements for Internet Hosts – Communication Layers. RFC 1122, Internet Request For Comments, October 1989.
- [4] Nevil Brownlee and kc claffy. Understanding internet traffic streams: Dragonflies and tortoises. *IEEE Communications Magazine*, 2002.
- [5] Wu chang Feng, Dilip Kandlur, Debanjan Saha, and Kang Shin. Stochastic fair blue: A queue management algorithm for enforcing fairness. In *Proceedings of the IEEE Infocom*. IEEE, April 2001.

- [6] Kimberly C. Claffy, Hans-Werner Braun, and George C. Polyzos. A parameterizable methodology for internet traffic flow profiling. *IEEE Journal of Selected Areas in Communications*, 13(8):1481–1494, 1995.
- [7] M. Crovella and M. Taqqu. Estimating the heavy tail index from scaling properties. *Methodology and Computing in Applied Probability*, 1(1):55–79, 1999.
- [8] C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *Proceeding of ACM SIGCOMM Internet Measurement Workshop 2001*, San Francisco Bay Area, November 2001.
- [9] Cristian Estan, Stefan Savage, and George Varghese. Automatically inferring patterns of resource consumption in network traffic. In *Proceedings of the ACM SIGCOMM*, pages 301–313, Karlsruhe, Germany, August 2003. ACM.
- [10] WenJia Fang and Larry Peterson. Inter-AS traffic patterns and their implications. In *Proceeding of IEEE GLOBECOM 99*, pages 1859–1868, Rio de Janeiro, Brazil, 1999.
- [11] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. Hypertext Transfer Protocol – HTTP/1.1. RFC 2616, Internet Request For Comments, June 1999.
- [12] Sally Floyd. Tcp and successive fast retransmits. *Lawrence Berkeley Laboratory Technical Report*, May 1995.
- [13] Sally Floyd and Vern Paxson. Difficulties in simulating the Internet. *ACM/IEEE Transactions on Networking*, 9(4):392–403, February 2001.
- [14] Leland, W.E.; Taqqu, M.S.; Willinger, and D.V. W.; Wilson. On the self-similar nature of Ethernet traffic (extended version). *ACM/IEEE Transactions on Networking*, 2(1):1–15, February 1994.
- [15] Ratul Mahajan, Steven M. Bellovin, Sally Floyd, John Ioannidis, Vern Paxson, and Scott Shenker. Controlling high bandwidth aggregates in the network. *ACM Computer Communication Review*, 32(2):62–73, 2002.
- [16] Sean McCreary and Kimberly C. Claffy. Trends in wide area IP traffic patterns - a view from ames Internet exchange. *ITC Specialist Seminar*, 2000.
- [17] P. Mockapetris. DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION. RFC 1035, Internet Request For Comments, November 1987.
- [18] NLANR. PMA Long Traces Archive. <http://pma.nlanr.net/Traces/long/>.
- [19] Los Nettos-Passing packets since 1988. <http://www.ln.net>.
- [20] R. Pan, L. Breslau, B. Prabhakar, and S. Shenker. Approximate fairness through differential dropping, January 2002.

- [21] D. Papagiannaki, N. Taft, S. Bhattacharyya, P. Thiran, K. Salamatian, and C. Diot. A pragmatic definition of elephants in internet backbone traffic. In *Proceeding of ACM SIGCOMM Internet Measurement Workshop 2002*, pages 175–176, Marseille, France, November 2002.
- [22] Vern Paxson, A. Adams, and M. Mathis. Experiences with NIMI. In *Proceedings of Passive and Active Measurement (PAM) 2000.*, 2000.
- [23] Jon Postel. TRANSMISSION CONTROL PROTOCOL. RFC 793, Internet Request For Comments, September 1981.
- [24] Chunming Qiao and Myungsik Yoo. Optical Burst Switching (OBS) - A New Paradigm for an Optical Internet. *Journal of High Speed Networks (JHSN) on WDM networks*, 8(1):69–84, 1999.
- [25] Shriram Sarvotham, Rudolf Riedi, and Richard Baraniuk. Connection-level analysis and modeling of network traffic. In *Proceeding of ACM SIGCOMM Internet Measurement Workshop 2001*, pages 99–104, San Francisco Bay Area, November 2001.
- [26] Anees Shaikh, Jennifer Rexford, and Kang Shin. Load-sensitive routing of long-lived IP flows. In *Proceedings of the ACM SIGCOMM*, pages 215–226. ACM, August 1999.
- [27] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: reshaping the research agenda. *ACM Computer Communication Review*, 26(2):19–43, April 1996.
- [28] Splus. Splus user’s guide. <http://www.insightful.com/products/default.asp>.
- [29] W. Stevens. TCP/IP illustrated, volume 1: The protocols. *Addison-Wesley*, 1994.
- [30] Kevin Thompson, G. Miller, and R. Wilder. Wide area internet traffic patterns and characteristics. *IEEE Network Magazine*, 11(6):10–23, November 1997.
- [31] Duane Wessels and Marina Fomenkov. Wow, that’s a lot of packets. In *Proceedings of Passive and Active Measurement Workshop (PAM) 2003*, San Diego, CA, April 2003.
- [32] Yin Zhang, Lee Breslau, Vern Paxson, and Scott Shenker. On the characteristics and origins of internet flow rates. In *SIGCOMM*, Pittsburgh, PA, USA, August 2002. IEEE.
- [33] Yin Zhang and Lili Qiu. Understanding the end-to-end performance impact of RED in a heterogeneous environment. *Cornell CS Technical Report TR2000-1802*, 2000.