Packet interarrival modeling for efficient detection of elephant flows on the Internet

J Rivillo and I W Phillips†

† Loughborough University

Abstract: In the last years the research community has been studying the Internet behaviour to have a better understanding of what could allow a refinement for its protocols, management as well as components to improve the Internet's performance. A great effort has been done in providing suitable models in this direction. This paper is a review of the Internet modeling with special focus on the flow's packet interarrival distribution and its application in the efficient early detection of elephant flows.

1. Introduction

In the old days of traffic engineering, modeling in telecommunications networks was simpler than today. There was only one kind of traffic: voice. It presented well known characteristics; Poisson interarrival rates (time between calls reaching the central office switch) and exponential call length distribution. There was no need to worry about things such as network layers, as they did not exist. It was easy to measure critical values of the important parameters. Queuing theory permitted analysis of voice networks to meet any desired performance characteristics; for example, call-blocking probability.

Recent measurements of traffic network have shown that it is often bursty over a wide range of time scales, a much wider range than is captured by traditional traffic models. The variability over wide time scales implies that bursts do not average out over long enough time period; no matter what is the scale being considered on the number of packets crossing a network link, always a similar amount of packet bursts will be resulted. Statistically, this phenomenon is referred to as *Self-Similarity* and has received significant attention in the networking community [1][2][3][4][5][6][7]. These studies were resulted because traditional models were too inaccurate which leads to wrong network behavior simulations and under optimized elections of different network design parameters.

These advances in modeling can be in different network's invariants were summarized by Fischer [8]**Error! Reference source not found.**. In this research, special interest lies in the statistical behavior of the flow's packet interarrivals, which has been shown to be heavy-tailed [9]. Detailed information about it is provided in the next two sections where the heavy tailed distributions and the packet interarrival behavior are explained.

2. Heavy-tailed distributions

These are distributions whose tails follow a power-law with low exponent, in contrast to traditional distributions (e.g., Exponential, Poisson) whose tails decline exponentially. In the late 80's and early 90's, accumulate experimental evidence began to show that some properties of computer systems and networks displayed distributions with very long tails, and thus attention turned to heavy-tailed distributions particularly in the mid 90's [1][9].

The definition of a heavy tailed distribution is the following: Let X be a random variable with cumulative distribution function F(x) = P[X ≤ x] and its complement $\overline{F}$ (x) = 1 - F(x) = P[X > x]. If X is heavy tailed, then the function will be

$$P[X > x] \sim cx^{-\acute{a}} \; ; 0 < \alpha < 2 \qquad (1)$$

Where c is a positive constant, and a ~ b means:

$$\lim_{x \to \infty} \frac{a(x)}{b(x)} = 1 \qquad (2)$$

This definition restricts our attention to distributions with strictly polynomial tails, although broader classes such as the subexponential distributions [10] can be defined and most of the qualitative remarks are similar.

The traditional modeling methods have focused on distributions with "light" tails which declines exponentially fast (or faster) and arbitrarily large observations are vanishingly rare. On the other hand, heavy tailed models behave quite differently because arbitrarily large observations have non-negligible probability. In fact, large observations, although rare, can dominate a system's performance characteristic. This is what happens in the elephant and mice phenomenon, where a few elephant flows carry most of the traffic data. Other examples of heavy tail distributions are the sizes of data objects in computer systems and the process/job lifetimes.

Special care must be paid to deal with heavy tailed distributions, because the workload metrics following heavy tailed distributions are extremely variable. In fact, their variance is infinite and for $\alpha<2$, the mean is infinite as well. In practice, empirical moments are slow to converge or are non convergent. To characterize system performance, either attention must shift to distribution itself, or attention must be paid to the timescale of analysis.

Although there are difficulties in dealing with heavy tailed distributions, they exhibit some properties that can be exploited in systems design: the Expectation Paradox and the Mass-Count Disparity.

2.1. The Expectation Paradox

The Expectation Paradox implies that in a heavy tailed event, the longer the wait of an event, the longer it is expected to wait. Poisson distribution presents the opposite phenomenon, take phone call generation as an example, the longer the wait, the sooner a call is expected. This is also another example telling why the traditional modeling is unsuitable for detection application in current networks.

Assuming a heavy-tailed distribution is a Pareto one. This distribution has the following definition:

$$p(x) = \alpha k^{\alpha} x^{-(\alpha+1)}; 1 < \alpha \le 2, x > k \qquad (3)$$

Then,

$$E[X \mid X > k] = \frac{\alpha}{\alpha - 1} k \qquad (4)$$

In this formula, α is a constant, and the expected time for the next event ($E[X|X>k]$) is greater when k (the waiting time) increases.

2.2. Mass-Count Disparity

The second useful property of heavy tailed distributions is the mass-count disparity [11]. This property can be stated formally as [10]:

$$\lim_{x\to\infty} \frac{P[X_1 + .. + X_N > x]}{P[\max(X_1,..,X_N) > x]} = 1, \forall n > 2 \qquad (5)$$

which is the case when the X_i are positive random variables drawn from a heavy-tailed distribution. This property states that when considering collections of observations of a heavy-tailed random variable, the aggregated mass contained in the small observations is negligible when comparing to the largest observation in determining the likelihood of large values of the sum. In practice this means that the majority of the mass in a set of observations is concentrated in a very small subset of the observations.

3. Flow's packet interarrival time modeling

The packet interarrival time (*IA*) is defined as the difference of the arrival times (*AT*) of the i^{th} packet and the $(i-1)^{th}$ packet:

$$IA = AT(i) - AT(i-1) \qquad (6)$$

The measured traces have shown a heavy-tailed distribution in the flow's packet interarrival time. An example of this behaviour are the HTTP applications, where the HTTP requests and responses for one web page that contain several objects are typically captured with very short interarrival times and between each web page loading, a user needs reading time or thinking time, which may last up to tens of minutes and cause the tail of the distribution to be significant.

Many heavy-tailed distributions have been tested by varying relevant parameters and comparing them to traces: Chi-squared distribution, exponential distribution, inverse Gaussian (Wald) distribution, lognormal distribution, Pareto distribution and Rayleigh distribution. Cruickshank [11] concludes that among all this heavy tailed distributions, the function which best fits to the HTTP traces observed for packet interarrival time is the inverse Gaussian distribution.

The National Laboratory for Applied Network Research (NLANR) [12] has been studying the intra-properties of flows in different traces and they state that it is possible to distinguish a high volume flow from the others through the packet interarrival time distribution and this may be useful to router designers need to design high performance forwarding paths for certain high volume flows.

4. Elephant and mice phenomenon

Today's Internet carries traffic from a wide range of applications, each of them with different requirements and constraints on network resources. Applications with special constraints such as video-conferencing, Internet telephony, on-line gaming, multimedia streaming and many more, have gradually appeared and the suitability of Internet’s best-effort policy and the need for providing different qualities of service to each type of application are being proposed.

The two types of flows referred to are those which make an extensive use of network resources (high bandwidth long-lived), the elephant flows, and those which do not consume so many resources, the mice flows. Typically, elephant flows consist of low-priority applications, that is large data transfer transactions and peer to peer file sharing. On the contrary, mice flows tend to be sensitive to delay jitter and high loss rates, which are mainly experienced in on-line gaming, small-sized web requests, multimedia broadcasting and voice over IP.

The accurate identification and special treatment of elephant flows, either by rerouting, throttling, priority management, is crucial to guarantee a better performance of the global network and a higher user satisfaction.

5. Elephant and mice packet interarrival time distributions

In order to illustrate what is noted in the previous paragraph, a 70-second traffic trace collected by NLANR has been considered. This is a backbone traffic trace that is collected at the output of the Indianapolis router towards Cleveland. Such trace contains about 395,000 different flows carrying around 3.74Gbytes. Figure 1A shows the total contribution of the largest traffic flows. The x-axis represents percentage of flows, and the y-axis shows the percentage of traffic volume. For only 0.1% of the traffic flows, that is 395 flows, which contribute to nearly 83% of the total traffic. The next 0.9% of the biggest traffic flows (another 2,555) contributes to an extra 13% of the traffic. The striking feature is that the remaining 99% flows (391,050) carry only 4% of the total traffic.

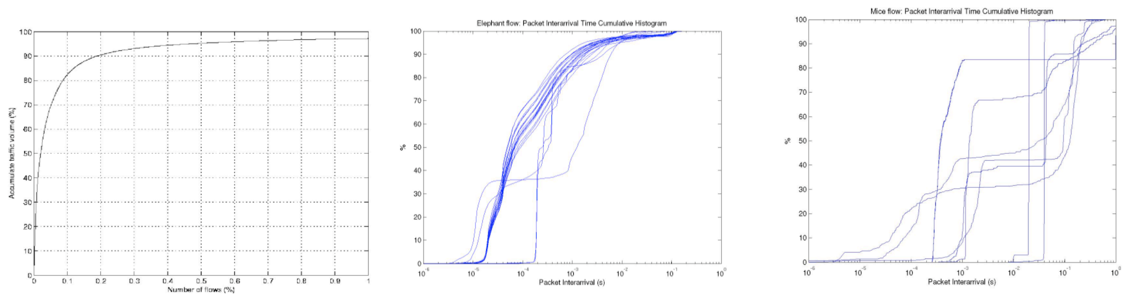

Figure 1. A) Accumulated traffic volume. B) Packet interarrival time distributions of the biggest elephants. C) Packet interarrival time distributions of the biggest mice flows

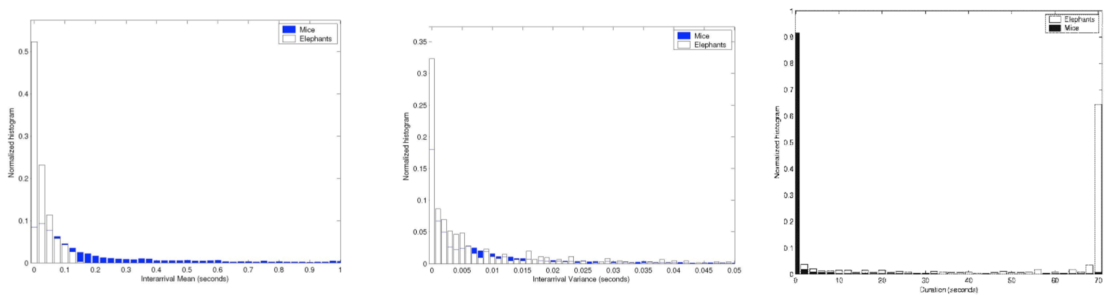

Figure 2. A) Mean packet interarrival time. B) Packet interarrival time variance. C) Flow duration

Packet interarrival times of the elephant and mice distributions from the trace observer are illustrated in this section. The threshold chosen to define the elephant flows is the 0.15% biggest flows, which according to figure 1A, carry over 87% of the total traffic.

Figures 1-2 show some statistics on this elephant traffic. The mean and variance in the packet interarrival time for both elephant and mice flows can be observed in figures 2A and 2B. As expected, elephant flows have short average packet interarrival time, typical from intensive applications. Figure 1B and 1C shows the packet interarrival time distributions of the biggest elephant and mice flows respectively in the trace. All the packet interarrival distributions of the elephants reach the 100% before 0.2 seconds, showing that the elephant flows doesn't have important pauses. Also it is observed that most of the elephant packet interarrival time values are concentrated in lower values than in the mice ones. Finally, Figure 2C shows the flow duration where black and white bars represent mice and elephant flows respectively. As observed, most of elephants are long-duration flows and mice flows rarely exceeds 5 seconds. All these flow's features explained could be exploited in designing an accurate detection strategy.

7. Conclusions

Firstly, this paper reviewed the current modeling status in the flow packet interarrival time. Secondly, the packet interarrival time distributions obtained in an Internet trace for both elephant and mice flows was discussed. Finally, the differences in the packet interarrival time found in the elephant and mice flows suggest that it can be used to detect elephant flows in aggregated traffic in an efficient way.

Acknowledgments

This work is undertaken under the project MASTS (Measurement and Analysis at all Scales in Time and Space), which has brought together efforts from Loughborough University, Cambridge University and the University College of London.

References

[1] 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," in IEEE/ACM Transactions on Networking, 5(1):71{86, February 1997.

[2] R. Addie, M. Zukerman, and T. Neame, "Fractal Traffic: Measurements, Modeling and Performance Evaluation," in Proc. IEEE INFOCOM '95, pp. 977-984, 1995.

[3] S. Andradottir, K.J. Healy, D. H. Withers, and B. L. Nelson, "Long-lasting Transient Conditions in Simulations with Heavy-tailed Workloads," in Proc. Of the 1997 Winter Simulation Conference.

[4] M. E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," in Proceedings of the 1996 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems. pp. 160--169.

[5] V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling," IEEE/ACM Transactions on Networking, Vol. 3 No. 3, pp. 226-244, June 1995.

[6] W. E. Leland, M. S. Taqqu, W. Willinger, and D.V. Wilson. "On the self-similar nature of Ethernet traffic (extended version)," IEEE/ACM Transactions on Networking, 2:1-15. 1994

[7] A. Feldmann, A. C. Gilbert, W. Willinger, and T. G. Kurtz, "The Changing Nature of Network Traffic: Scaling Phenomena," in ACM Computer Communication Review, vol. 28, pp. 5-29, Apr. 1998.

[8] M. J. Fischer, T. B. Fowler, "Fractals, Heavy Tails and the Internet", in Sigma, Mitretek Systems, McLean, VA, Fall 2001.

[9] K. Park, G. Kim and M. Crovella, "On the Effect of Traffic Self-similarity on Network Performance," in Proceedings of the 1997 SPIE International Conference on Perfomance and Control of Network Systems.

[10] M. E. Crovella, "Performance evaluation with heavy tailed distributions ". In Job Scheduling Strategies for Parallel Processing, D. G. Feitelson and L. Rudolph (eds.), pp. 1–10, Springer Verlag, 2001. Lect. Notes Comput. Sci. vol. 2221.

[11] H. Cruickshank, "Internet QoS Measurement and Traffic Modeling," in ATS Conference, Stuttgart. 2003

[12] National Laboratory for Applied Network Research, "Tracing on intra-flow patterns", http://www.nlanr.net/NA/Iat/ , accessed November 2006.