Scaling Laws in Network Traffic

Ovid C. Jacob
ovid.jacob@sun.com

Sun Microsystems Computer Company
Network and Security Products Group
901 San Antonio Road, MS NWK02-130
Palo Alto, CA 94303

1 Abstract

In this paper we propose a new scaling law for wide area network traffic, based on some observation of Paxson and coworkers [1] Leland and coworkers [2] Crovella and Bestavros [3] as well as on the work of West and co-workers [4, 5].

Paxson and coworkers found that network arrivals are modeled very well by self-similar processes rather than the usual Poisson processes used so far. West and coworkers have shown that biological systems which have a network flow in them display a remarkable property in that many biological quantities can be described by a simple scaling law.

We propose then that network traffic should be modeled by a scaling law where network traffic dependent quantities scale as integer multiples of $\frac{1}{3}$ power of size of the network.

2 Introduction

Several papers have studied recently the nature of network traffic, both local area network (LAN) and WAN. The primary focus of these papers has been the nature of the distribution of traffic. Specifically, the papers which will be discussed here question the received notion that this traffic has a Poisson distribution, and have found instead that the traffic has a self-similar distribution.

We put forward here a new hypothesis, regarding the functional form of this traffic. Specifically, we propose that this traffic depends on the number of nodes or hosts present in the network, and we propose that this functional form is

$$Y = cN^{\beta}$$

where Y is some network dependent variable, c is some constant, N is the number of hosts (IP addresses) on the net, and β is the scaling exponent. Basing our work on that of West and co-workers, we advance the hypothesis that in WAN traffic, this exponent is $\frac{1}{3}$.

The outline of the paper is as follows: first, we will discuss some of the earlier work on LAN traffic initiated by Leland and co-workers [2]; next, we present the work on WAN traffic of Paxson and Floyd [1]; after that, we present the work on allometric scaling of West and co-workers [4, 5]; we then discuss recent work on self-similarity in WWW traffic of Crovella and Bestavros [3]; finally, we put all this together and see how this leads us to formulate the scaling law mentioned above.

3 LAN Traffic Studies

The first work we will study is that of Leland and co-workers [2] on the self-similar nature of Ethernet traffic. The work was based on data collected between August 1989 and February 1992 at Bellcore Morristown Research and Engineering Center.

Let us first direct the reader to a visual representation of this self-similarity: Figure 4 of [2] show that plotting packets/time unit versus time in time units of 0.01 sec, 0.1 sec, sec, 10 sec and 100 sec show all the same shape, being quite bursty. This differs strongly from the shape of packets/time unit versus time had the traffic a Poisson, distribution when we would expect a rather different behavior for different time units.

We introduce now a more formal, mathematical definition of self-similarity. We introduce a stochastic process

$$X = X_t : t = 0, 1, 2, ...$$

with an auto correlation of the form

$$r(k) = k^{-\beta} L(t), k - - > \infty$$

where $0 < \beta < 1$ and L is slowly varying at infinity. One implication of self-similar processes is that if there is a spectral function f(.), it would obey a power law near the origin

$$f(\lambda) = const.\lambda^{\gamma}$$

with $0 < \gamma < 1$ and $\gamma = 1 - \beta$.

The authors use two formal mathematical models to represent self-similar processes: fractional Gaussian noise and fractional autoregressive integrated moving-average (ARIMA) processes. The latter has an auto correlation of the form

$$r(k) = \frac{1}{2}\{|k+1|^{2H} - |k|^{2H} + |k-1|^{2H}\} k > 0$$

where the self-similarity Hurst parameter H is $\frac{1}{2} < H < 1$. Calculations show that $H = 1 - \frac{\beta}{2}$.

Leland and co-workers found that the Ethernet traffic data is best described by $H = 0.84$ and $\gamma = 0.66$. This corresponds to $\beta = 0.33$. Figure 5 of [2] shows a more detail measurement of H, with its value fluctuating very little about 0.84.

The authors of [2] don't give a mechanism or hypothesize why the traffic might be self-similar, but they find that Ethernet LAN is statistically self-similar; the degree of self-similarity is well measured by Hurst's parameter H, with burstier traffic corresponding to higher H; external LAN and external TCP traffic also seem to have a self-similar characteristics as the overall LAN traffic; the packet traffic models currently considered in literature are not able to capture the self-similar property - perhaps they should?!

4 WAN Traffic Studies

We now proceed to look at the work of Paxson and Floyd [1], which looked at wide area network traffic. Taking as their starting point the work of Leland and co-workers [2] described above, the authors looked at TELNET and FTP sessions from 24 traces involving data gathered at LBL, DEC WRL, Bellcore, UCB, coNCert, as well as data from UK. The authors find that for WAN traffic, connection arrivals are well-modeled by Poisson processes, but that for packet arrivals, self-similar processes model the data much better.

In more detail, TELNET connection arrivals are well-modeled by Poisson models, but TELNET packets are better modeled by self-similar processes. FTP connection arrivals is not well modeled by Poisson models and FTP packets arrivals is also better modeled by self-similar processes.

Paxson and Floyd use large-scale correlations as a measure of self-similarity. They define a signal as having long-range dependence if the auto correlation function $r(k)$ diverges

$$\sum_k r(k) = \infty$$

The authors propose several methods for the production of self-similar traffic: multiplexing On/OFF that have a fixed rate in the ON period and ON/OFF period lengths that are heavy-tailed; $M/G/\infty$ queue model, where customers arrive according to a Poisson process and have service time drawn from a heavy-tailed distribution with infinite variance. This discussion is beyond the scope of this paper.

Unfortunately, the authors do not give semi-analytic forms of the traffic, against which we could test our hypothesis. But a look at Figure 5 of [1] gives a value for the variance of the data of -0.33 rather than -1 as expected from Poisson processes. We get a similar value from the data in Figures 7, Figures 12, and 13 of [1].

The authors conclude that we should abandon Poisson models of WAN for most traffic except user session arrivals. They also suggest that multiplexed TCP and all-protocol traffic might be described very well by self-similar processes.

5 Allometric Scaling Laws in Biology

West and collaborators [5] study a large body of data from biology which shows that a biological variable Y depends on the body mass M in an allometric scaling way

$$Y = cM^{\beta}$$

where β is a scaling exponent and the constant in front depends on characteristics of the organism. If these relations reflect geometric constraints, β is a simple multiple of $\frac{1}{3}$. But much biological data is more consistent with β being a multiple of $\frac{1}{4}$.

The authors of [5] proceed to show that with three rather simple assumptions, we can obtain this result. They introduce these assumptions: One, a space-filling fractal branching pattern is required for the network to supply the entire volume of the organism; Second, the final branch of the network is a size-invariant unit; and Third, the energy required to distribute resources is minimized.

Let us follow the authors' arguments. Following them, we introduce typical lengths l_k and radii r_k for the branches of the network. The volume rate flow is defined by

$$\frac{\partial Q_k}{\partial t} = \pi r_k^2 u_k$$

where u_k is the average flow velocity. Each level k has n_k branches, so at this level the total number of branches is $N^k = n_0 n_1 ... n_k$. The special terminal units are size invariant, and are defined by radius r_c, length l_c and average flow velocity u_c.

They show that

$$\frac{\partial Q_0}{\partial t} = B \ = M^a = N_k \frac{\partial Q_k}{\partial t} = N_k \pi r_k^2 u_k = N_c \pi r_c^2 u_c$$

so that $N_c = M^a$.

They introduce then scale factors $\delta_k = \frac{r_{k+1}}{r_k}$ and $\gamma_k = \frac{l_{k+1}}{l_k}$.

For self-similar fractals, $\delta_k = \delta$, $\gamma_k = \gamma$ and $n_k = n$. Thus, $N_k = n^k$, so $N_c = n^N$.

Let us study now the behavior of total volume of fluid in the network.

$$V_b = \sum_k N_k V_k = \sum_k \pi r_k^2 l_k n_k$$

$$V_b = \frac{V_0}{(1 - n\gamma\delta^2)} = \frac{V_c(\gamma\delta^2)^{-N}}{(1 - n\gamma\delta^2)}$$

It can be shown that $V_b = M$ so that $(\gamma\delta^2)^{-N} = M$ so that

$$a = \frac{-ln(n)}{ln(\gamma\delta^2)}$$

Let us look now at what can we learn about δ and γ . The first hypothesis, the existence of a space-filling implies that the size of the network is best described by $l_k >> r_k$. with a volume at the k-th level of $\frac{4}{3}\pi(\frac{l_k}{2})^3 N_k$. The third condition, energy minimization, can imply that the flow if volume preserving or area preserving. Volume preservation for space-filling means that

$$\frac{4}{3}\pi(\frac{l_k}{2})^3 N_k = \frac{4}{3}\pi(\frac{l_{k+1}}{2})^3 N_{k+1}$$

so $\gamma_k^3 = (\frac{l_{k+1}}{l_k})^3 = \frac{N_k}{N_{k+1}} = \frac{1}{n}$, so $\gamma_k = n^{-\frac{1}{3}} = \gamma$.

Area preserving flow has quite different implications. We have $\pi r_k^2 = n\pi r_{k+1}^2$, or $\delta_k = \frac{r_{k+1}}{r_k} = n^{-\frac{1}{2}} = \delta$, $\gamma = n^{-\frac{1}{2}}$. This implies, with $\gamma = n^{-\frac{1}{3}}$, that $a = \frac{3}{4}$ and

$$B = M^{\frac{3}{4}}, r_0 = M^{\frac{3}{8}}, l_0 = M^{\frac{1}{4}}$$

.

The authors of this work argue that in biological systems we have area preserving flow, so that we get scaling with multiples of $\frac{1}{4}$. Table 1 of [5] supports this hypothesis amply, with many cardiovascular and respiratory quantities explained well by the quarter law. They also point out that this is a fractal of dimension $D = 3$, where $a = \frac{D}{D+1} = \frac{3}{4}$.

6 WWW Traffic Studies

Crovella and Bestavros [3] pick up where Leland and co-workers left off, and examine WAN traffic on the World Wide Web (WWW). Their data was gathered between January and February of 1995. Table 1 of [3] describes the data used. They use the same formalism discussed in this paper already.

The data they gathered seems to imply a value of $H = 0.83$, or a value of the auto correlation exponent $\beta = 0.33$. They also analyze the specific times associated with file transfers. For all file transfers, they obtain $\alpha = 1.27$ or $H = 0.87$ so $\beta = 0.27$. For text files only, though, they get $\alpha = 1.59$, or $H = 0.70$ or $\beta = 0.59$. Here, $H = \frac{(3-\alpha)}{2}$, so $\alpha = 1 + \beta$.

The model they use for explaining the nature of the WWW traffic is that introduced by Paxson and Floyd [1], discussed above, regarding multiplexed ON/OFF sources. For ON times, they observe $\alpha = 1.0 - 1.3$ or $\beta = 0 - 0.3$, while for the OFF times, $\alpha = 1.5$ or $\beta = 0.5$, so it seems that the ON times is responsible for the self-similar traffic.

7 Scaling Laws in Network Traffic

We are ready now to propose our model. We use as our foundation the analysis of West and co-workers [5].

We will model our computer network as a network of switches/hubs with ports p_k. Signals are transported by this network from port to port.

We will make several assumptions about this network: First, to supply all the ports with signal, it has a fractal branching pattern; Second, the final branch of the network, like the end port on a desktop or other computer, is of roughly the same network throughput; and Third,the energy required to distribute these signals is minimized.

More specifically, we introduce these assumptions: First, the factual branching is linear preserving of ports or hops h_k; and Third, the energy required to distribute the signal is minimized and is linear preserving in hops.

The first hypothesis, the existence of a fractal pattern implies that the the network is best described by $h_k >> p_k$. We will require that this fractal is linear preserving in ports. The total number of ports at the k-th level is $p_k N_k$.

This implies that

$$p_k N_k = p_{k+1} N_{k+1}$$

so $\gamma_k = \frac{p_{k+1}}{l_k} = \frac{N_k}{N_{k+1}} = \frac{1}{n}$, so $\gamma_k = n^{-1} = \gamma$.

Line preserving flow which minimizes energy transmission of signal has certain implications. We have the following for hops $h_k = nh_{k+1}$, or $\delta_k = \frac{h_{k+1}}{h_k} = n^{-1} = \delta$. This implies, with $\gamma = n^{-1}$, $a = \frac{1}{3}$ and

$$Y = N^{\frac{1}{3}} p_0 = N^{\frac{1}{3}}, h_0 = N^{\frac{1}{3}}$$

.

Thus, a quantity Y depends on the number of hosts in the network proportional to the $\frac{1}{3}$ power of those hosts. Likewise, the average amount of traffic at port p_0 is proportional to a similar quantity of total number of hosts to the $\frac{1}{3}$ power.

There seems to be ample support for our hypothesis. If we look at the data of Paxson and Floyd [1], we note that the variance data is well described by $-\frac{1}{3}$. Likewise, if we look at the work of Crovella and Bestavros [3], $\beta = \frac{1}{3}$. For all file transfers, $\beta = \frac{1}{3}$ while for text files, $\beta = \frac{2}{3}$.

As an aside, we remark that LAN traffic is also well described by this hypothesis. Leland and co-workers' [2] data is well supported by a auto correlation with $\beta = \frac{1}{3}$.

8 Application

Consider a layer 4 device (in the OSI model). This device receives packets from N nodes on the network. It sends the traffic received to a list of servers, depending on the load. Suppose these servers are designed to do load balancing, round robin for instance.

From the present work, we know that if the nodes N double, the load will only grow by $2^{\frac{1}{3}}$, or 1.26. Depending on the load on the existing servers, we might not need to add new servers to be used for load balancing.

Namely, if the CPU utilization is less than 79then we do not need to add servers. But if the utilization is greater than 79

More genereally, if the number of load increases by a factor F, we expect the load to increase only by $F^{\frac{1}{3}}$. As long as the CPU utilization on the servers is less than $F^{-\frac{1}{3}}$, we do not need to add more servers.

9 Conclusions

We proposed a novel scaling law for WAN traffic, which predicts that network dependent quantities scale as integer multiples of $\frac{1}{3}$.

References

[1] Vern Paxson and Sally Floyd, 'Wide Area Traffic: The Failure of Poisson Modeling', IEEE/ACM Trans. Networking, Vol. 3, pp. 226-244, June 1995.

[2] Will E. Leland, Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson, 'On the Self-Similar Nature of Ethernet Traffic (Extended Version)', IEEE/ACM Trans. Networking, Vol. 2, pp. 1-15, January 1994.

[3] Mark E. Crovella and Azer Bestavros, 'Self-Similarity in WWW traffic: Evidence and Possible Causes,' IEEE/ACM Trans. Networking, Vol. 5, pp.835- 846, December 1997.

[4] Diane Banegas , 'The Tree of Life: Scientists Model Nature's System of Fractal-branching Networks,' SFI Bulletin, pp. 7-10, Summer 1997.

[5] Geoffrey B. West, James H. Brown, Brian J. Enquist, 'A General Model of Allometric Scaling in Biology,' Science, Vol. 276, p.122, 1997.