Internet Traffic: Statistical Multiplexing Gains

Jin Cao, William S. Cleveland, Dong Lin, and Don X. Sun
{cao,wsc,dong,dxsun}@research.bell-labs.com

I. Introduction

This note describes recent results on the effect of statistical multiplexing on the long-range dependence (LRD) of Internet packet traffic. Details and bibliographies may be found in a series of papers at http://cm.bell-labs.com/stat/InternetTraffic/webpapers.html.

Let a_j, for $j = 1, 2, \ldots$ be the arrival times of packets on an Internet link, let $t_j = a_{j+1} - a_j$ be the inter-arrival times, and let q_j be the packet sizes. Suppose we divide time up into equal-length, consecutive intervals. Let p_i be the packet counts in interval i. The t_j and q_j are studied as time series in j, and the p_i as a time series in i.

The packet traffic on a link is the result of the statistical multiplexing of packets from different active connections. Let c be the mean number of active connections over an interval of time during which link usage is stationary. c serves as a measure of the magnitude of statistical multiplexing over the interval. This note addresses the effect of an increasing c on the LRD of the three traffic variables t_j, q_j, and p_i.

Results are based on the following: (1) the mathematical theory of marked point processes; (2) empirical study of live packet traces from 15 interfaces whose 5-min average traffic rates range from about 2 kbps to 250 mbps; (3) empirical study of synthetic packet traces from network simulation with NS; (4) simple statistical models, FSD models and FSD-MA(1) models, for the traffic variables.

II. Theory

Theory is easy and convincing for a link on an over-provisioned network where c is never so big that there is more than minor queueing. As c increases, a_j tends *toward* a Poisson process, so t_j tends *toward* independence. As c increases, q_j also tends *toward* independence. We use the word *toward* because t_j and q_j each always has an LRD component, but the contribution of the component to the variability of the time series gets less and less.

The correlation structure of p_i, however, does not change, so LRD is preserved in the counts. (This might seem like a contradiction to the results for t_j; later we describe how the two results fit together.) The coefficient of variation (standard deviation divided by the mean) of p_i goes to zero like $1/\sqrt{c}$. This means that the excursions of p_i above or below the mean, which last for long periods of time because of the LRD, get smaller and smaller in magnitude relative to the mean. While the LRD of the p_i is unchanging in the sense that the autocorrelation is unchanging, eventually the LRD ceases to be salient because the variability of the p_i gets small.

The over-provision theory applies to a link for the range of values of c that are small enough that queueing at the input router and nearby upstream devices is not substantial. Upstream queueing of a sufficiently large magnitude invalidates the assumptions that lead to the over-provision results. A more sophisticated theory to handle the queueing does not exist. Rough arguments suggest that the LRD of t_j and q_j should still tend *toward* independence, that the LRD of the p_i is eventually altered, and that the coefficient of variation eventually stabilizes to a small positive constant. But we need to appeal to empirical study of live traffic from the Internet, and synthetic link traffic from simulations, to resolve the uncertainty.

III. FSD and FSD-MA(1) Statistical Models

We found that very simple statistical time series models, which we call fractional sum difference (FSD) models and FSD-MA(1) models, provide an excellent fit to the t_j, q_j, and p_i for the live and synthetic link packet traces.

The first step in the modeling is to transform t_j and p_i to bring their marginal distributions closer to Gaussian. We transformed the t_j by sixth roots and the p_i by logs. The three time series, two transformed and one not, are then normalized by subtracting the sample mean and dividing by the sample standard deviation. Let t_j^*, q_j^*, and p_i^* denote the normalized time series. (We carried out our analyses without the transformation, and the results were similar, but the transformation puts the statistical modeling on a more rigorous basis.)

Let z_u be an FSD time series, normalized to have mean 0 and variance 1. Then

$$z_u = \sqrt{1-\theta}\; s_u + \sqrt{\theta}\; n_u,$$

where s_u and n_u are independent of one another and each has mean 0 and variance 1. n_u is white noise, that is, an uncorrelated time series. s_u is a fractional ARIMA model

$$(I - B)^d s_u = \epsilon_u + \epsilon_{u-1}$$

where $Bs_u = s_{u-1}$, $0 < d < 0.5$, and ϵ_u is white noise with mean 0 and a variance $\sigma_\epsilon^2 = ((1-d)\Gamma^2(1-d))/(2\Gamma(1-2d))$ to make the variance of s_u equal to 1.

Let $r_s(k)$ and $r_z(k)$ be the autocorrelation functions of of s_u and z_u respectively for lags $k = 0, 1, 2, \ldots$. s_u has LRD, and $r_s(k)$ falls off like k^{2d-1} and increases at all positive lags as d increases. The autocorrelation function of z_u is

$$r_z(k) = (1-\theta) r_s(k).$$

Thus z_u, whose variance is 1, is the sum of the correlated component $\sqrt{1-\theta} s_u$, whose variance is $1-\theta$, and the uncorrelated component $\sqrt{\theta} n_u$, whose variance is θ. The dependence decreases as θ increases toward 1; if $1-\theta$ decreases by a certain factor, then all $r_z(k)$ for $k > 0$ decrease by the factor. Finally, when $\theta = 1$, z_u is white noise.

The power spectrum of z_u is

$$p_z(f) = (1-\theta)\sigma_\epsilon^2 \frac{4\cos^2(\pi f)}{\left(4\sin^2(\pi f)\right)^d} + \theta$$

where the frequency f has units cycles/inter-arrival for t_j, cycles/packet for q_j, and cycles/interval-length for p_i, and where $0 \leq f \leq 0.5$. $p_z(f)$ decreases monotonically as f increases. $p_z(f)$ goes to infinity at $f = 0$ like $\sin^{-2d}(\pi f)$, so if $\theta < 1$, no matter how close θ gets to 1, $p_z(f)$ gets arbitrarily large near $f = 0$, but its ascent begins closer and closer to 0 as θ gets closer to 1.

The FSD-MA(1) model is

$$z_u = \sqrt{1-\theta}\, s_u + \sqrt{\theta}\, n_u,$$

similar to the FSD model, but where n_u instead of white noise is a first order moving-average

$$n_u = \zeta_u + \beta\zeta_{u-1},$$

where ζ_u is Gaussian white noise with mean 0 and variance $(1+\beta^2)^{-1}$, which makes the variance of n_u equal to 1. If $\beta = 0$, the moving-average component is white noise so the model is simply an FSD.

IV. Empirical and Simulation Study

For the live and synthetic packet traces, the FSD model provides an excellent fit to q_j^* and p_i^*, and the FSD-MA(1) to the t_j^*. Values of c have a range of 20–8200 connections for the live traces and 1–215 connections for the synthetic traces. The t_j^* and p_i^* for very small c can have effects not captured by the modeling, but even here the models serve as an excellent approximation for studying LRD. We found that a d of about 0.4 provides an excellent fit for all three traffic variables and for all c. Estimates of θ are less than 1 for all three traffic variables. In other words, there are elements of LRD in all variables at all rates.

As the trace value of c increases, θ tends to 1 for t_j^* and q_j^*. so there is a clear reduction in LRD. This means the q_j^* tend toward independence. For t_j^* from some interfaces, the FSD-MA(1) often has a β significantly greater than zero, so the t_j^* tend to short-range dependent in these cases. This is likely caused by upstream queueing at these interfaces. For other interfaces, $\beta = 0$ provides a good fit, so the model is an FSD, and the t_j^* tend to independent.

For p_i^*, θ shows no consistent change, so the autocorrelation of therefore the LRD of the p_i^* shows no consistent change with c; The coefficient of variation at all interfaces declines quite close to the rate $1/\sqrt{c}$ predicted by the over-provision theory. The unchanging LRD and the decline of the coefficient of variation occur, surprisingly, even at the interfaces with substantial queueing. However, the count interval length is 100 ms, and it is possible that for smaller intervals an alteration would occur.

V. t_j vs. p_i

Theory, empirical study, and simulation study show LRD dissipates for the inter-arrival times t_j but does not change for the arrival counts p_i in fixed-length intervals of time. The appearance is a contradiction. This is a case where the formal mathematics yields an unequivocal proof, but where we need an heuristic argument for better understanding. We will do this using the over-provision theory.

First, we fix the interval length ℓ used for the definition of p_i. Instead of considering p_i, we study p_i/ℓ, just a change of units, which are now packets per sec or p/s. Consider the process

$$t_j^{(b)} = b^{-1}(a_{bj+1} - a_{(bj+1-b)}),$$

for $j = 1, 2, \ldots$. These are the inter-arrival times per packet of blocks of b packets in units of s/p. To relate the $t_j^{(b)}$ to p_i/ℓ, we need to make the $bt_j^{(b)}$ vary around the interval length ℓ. Suppose there are m traffic sources, each with a mean inter-arrival time μ, that are multiplexed to form the link traffic. Then the mean of t_j is μ/m, and the mean of of $bt^{(b)}$ is $b\mu/m$. Thus we take $b = m\ell/\mu$. This means b increases with the magnitude of the multiplexing. We could take $1/t_j^{(b)}$, with units p/s, to be another measure of packets per second and consider its properties as a surrogate for p_i/ℓ to resolve our contradiction. But the math would be too hard. Instead we take $t_j^{(b)}$ itself as the surrogate.

Our surrogate has a very simple dependence on the t_j:

$$t_j^{(b)} = b^{-1} \sum_{v=bj}^{v=bj+1-b} t_v.$$

The operation of going from t_j to $t_j^{(b)}$ is a low-pass linear filtering of t_j followed by taking every b-th value. As m increases, so does b, and the frequency band that is passed gets closer and closer to 0. But as we have seen, t_j has an LRD component that affects the power spectrum over a band closer and closer to zero as m increases. The decreasing pass band of the operation increasingly filters out more of the noise process, $\sqrt{\theta}n_u$, than the LRD component $\sqrt{(1-\theta}s_u$, and effectively boosts the smaller LRD component for $t_j^{(b)}$ for a larger m back to where it was for a smaller value of m.

VI. Outcomes

Recent and current work is showing, as one would expect, that these results have important implications for traffic engineering. If we fix a buffer size and fix an amount traffic as measured by the connection load c, then the link speed needed to achieve a fixed QoS criterion, such as 0.5% packet loss, results in a utilization that increases dramatically with c. In other words, utilizations can increase dramatically from the edges to the core. More broadly, the results show that engineering studies that are meant to apply to the Internet as a whole, and that use synthetic or live packet traffic to assess performance, need to consider packet traces varying across a wide range of magnitudes of statistical multiplexing in order to achieve generality.