THE FAILURE OF "*THE FAILURE OF POISSON MODELING*" FOR INTERNET BANDWIDTH MEASUREMENTS

Svante Ekelin[a] *and Martin Nilsson*[b]

[a]Ericsson Research, Torshamnsgatan 23, SE-164 80 Kista, Sweden, svante.ekelin@ericsson.com
[b]SICS, POB 1263, SE-164 29 Kista, Sweden, from.rvk@drnil.com

ABSTRACT

There is a classic paper, which claims that many parameters of network traffic are best modeled by stochastic distributions with infinite variance.

We believe that although this paper is correct, it has commonly been overinterpreted. It is true that some distributions have infinite variance, but not all. There are important cases where distributions of traffic parameters have finite variance.

As an example, we show that the end-to-end available bandwidth must have finite variance, outline a method for measuring it, and present some experimental data.

Keywords: Modeling, network, distribution, variance, tail-heavy, available bandwidth, measurement, Kalman filter.

1. INTRODUCTION

A now classic paper entitled *Wide-area traffic: The failure of Poisson modeling* [1] argues that network traffic cannot be accurately modeled by probability distributions with finite variance, but are better described by tail-heavy distributions. Although this is certainly often correct, we believe that this result has been interpreted too pessimistically, leading to a common belief that network traffic cannot be measured using any ordinary statistical methods.

However, in the present paper, we show that important parameters of network traffic, such as the end-to-end path available bandwidth, can indeed be measured, without suffering from the analytical complexity of tail-heavy distributions. The reason for this is that whatever the statistical properties of the traffic demand, the communication channels have finite capacity, which will by necessity lead to the involved traffic parameters having finite variance. This is important, since it allows us to apply statistical analysis methods such as Kalman filters [2] for estimating the parameters.

In this paper, we shall describe an approach where we have used a Kalman filter in order to estimate the available bandwidth over an end-to-end path. We include the results of some experiments, using both Poisson/uniform distributions and tail-heavy Pareto distributions for generated cross-traffic patterns.

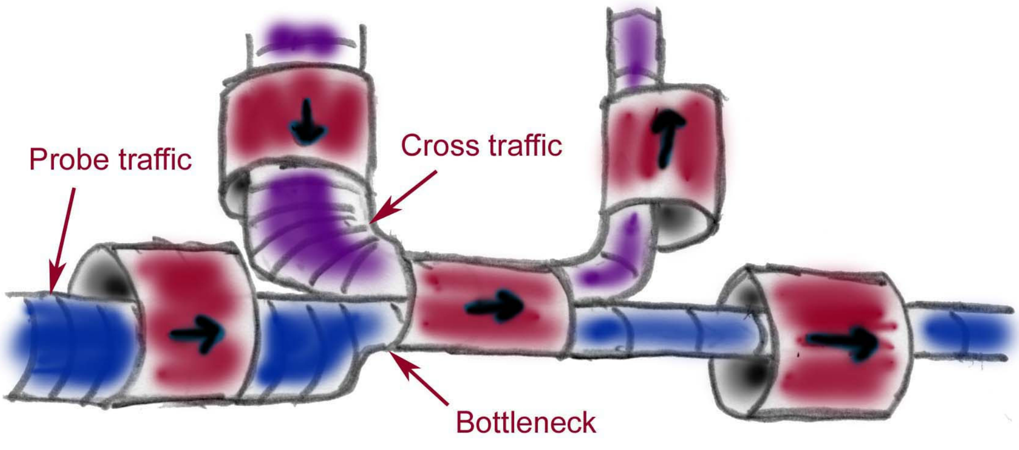

Fig. 1: The available bandwidth can be measured by sending probe packets.

2. MEASURING AVAILABLE BANDWIDTH

2.1 The variance of a stochastic variable

A *bounded* stochastic variable *must* have finite variance. Suppose that we have a stochastic variable X taking on values in the interval $[a,b]$. Its variance is then bounded, since

$$\begin{aligned} V(X) &= E((X-E(X))^2) = \\ &= E(X^2)-(E(X))^2 \qquad (1) \\ &\le \max(a^2,b^2) \end{aligned}$$

If we introduce $Y = X-(a+b)/2$, using (1), we have

$$\begin{aligned} & V(X)=V(Y)\le \\ & \max\left(\frac{(a-b)^2}{4},\frac{(b-a)^2}{4}\right)=\frac{(a-b)^2}{4} \qquad (2) \end{aligned}$$

Statistical methods for estimating the probability distribution of a parameter are usually designed for stochastic variables with finite variance. If the variance is infinite, measurement data will jump so wildly that it is very hard to extract any meaningful or reliable data.

2.2 Available bandwidth

The *available bandwidth* (ABW) of an end-to-end path is an important parameter for network traffic. Compared to measuring the theoretical maximum *capacity* [bits/s] of the path, ABW also takes into account the reduction of the capacity by cross traffic. More specifically, ABW is the smallest value of the capacity minus the cross traffic over all links in the path. Clearly, ABW is a bounded stochastic variable, because the available bandwidth is bounded below by zero, and above by the capacity of the narrowest link, which is finite. *This implies that ABW cannot have a tail-heavy distribution.*

2.3 Measuring ABW

Several methods based on *active probing* [fig. 1] have been proposed. This means that probe packets are sent over the path, and when the packets arrive, the *strain*, i.e. the relative increase in time separation between consecutive packets, is measured. From this value the available bandwidth can be estimated.

Fig. 2: Strain as a function of probe intensity.

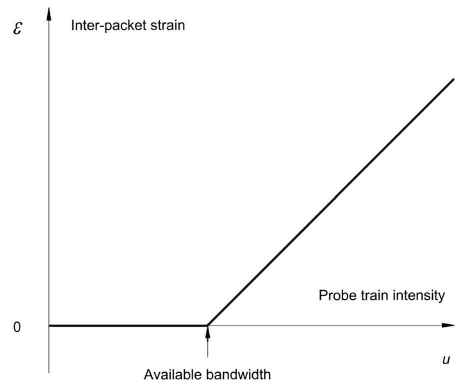

The strain itself is a stochastic variable. This variable will be bounded if the routers on the network path use a first-come-first-served (FCFS) policy. This condition is satisfied by an overwhelming majority of routers today.

Several methods based on active probing have been proposed, and a good overview, including the TOPP method, can be found in [3,4]. The TOPP method models the system by expressing the strain as a function of probe-packet rate, leading to a piece-wise linear relationship [fig. 2]. The available bandwidth corresponds to the first corner in the bandwidth-strain graph. The TOPP inventors proposed linear regression for analyzing the graph. Ekelin and Nilsson observed that Kalman filters could be applied to the model, and proposed the BART (*Bandwidth Available in Real-Time*) estimation algorithm [5]. Although the model has a strong non-linearity, it can be circumvented in order to apply a Kalman filter, allowing fast measurements.

3. INTRODUCTION TO BART

Monitoring of available bandwidth for an end-to-end network path would theoretically be possible without active probing, by having access to management data from all the network nodes in the path. However, network equipment owners do not usually make such data available. Measurement is only feasible by actively probing the network path, in order to determine at which probe rate the path shows signs of being congested.

In BART, care is taken to minimize negative effects on the real network traffic, and the packets experiencing congestion are temporarily stored in the router buffers. Typically, BART should causes negligible packet loss, and uses on the average only a small fraction of the bottleneck link capacity for probe traffic.

BART uses the inter-packet separation strain as a convenient, dimensionless measure of the interaction between the cross traffic and the probe traffic. The expectation value for the strain is zero for an uncongested state, but rises linearly with the overload rate when congestion occurs [fig. 2]. BART applies a Kalman filter, in order to maintain and update the estimate of the available bandwidth for each new measurement point. Although the system is non-linear over the whole range of probe rates, it is linear in the overload range; BART uses separate Kalman filters for the linear subsections.

Besides simulations, we have also carried out some initial tests of the method in a laboratory network, using carefully controlled cross traffic in order to be able to verify that the measurement method produces reasonable results. The results from BART so far show good agreement with the "true" available bandwidth, and also quick step response to a sudden change in cross traffic rate. Please refer to [5] for further details.

4. RESULTS

For the purposes of this paper, we have tested the behavior of the BART method for tail-heavy cross traffic using a simulator. We have tested the performance when cross traffic arrivals are Poisson-distributed, and the instantaneous load is either a) uniformly distributed or b) Pareto-distributed (tail-heavy) with $\alpha = 1.6$. In b), the traffic was shaped through a finite-capacity link before arriving at the router. The results are shown qualitatively in diagrams 1 and 2, respectively.

5. DISCUSSION

The optimality of Kalman filters requires that the system is linear, and that process and measurement noise is zero-mean gaussian. Experience has shown that when these conditions are relaxed, Kalman filters can still produce good results, although they aren't optimal. However, for tail-heavy noise distributions, Kalman filters will not work.

In the case of measuring available bandwidth, we can make the system model "nearly" linear, so the critical question is what the variance is. Clearly, both ABW and strain have finite variance, and consequently there is no parameter in the estimation process, which has a tail-heavy distribution.

A second question is that of being *infinite* vs. being *large*: Can the variance be finite, but very large, and in fact so large that it can be considered infinite for all practical purposes? This is relevant in situations where the average, or the "normal range" of the parameter is much less than the bound. However, this is not the case for ABW, where the bound is the capacity of the narrowest link, which is by no means an exceptional value for the ABW, average- or magnitude-wise. A similar argument holds for the strain.

Aggregation, or the situation where many small, independent traffic flows arrive in the same router, will tend to improve the performance of Kalman filtering, since the distribution of cross traffic will tend to become more Gaussian.

6. CONCLUSIONS

Although *many* parameters related to network traffic have tail-heavy distributions, this is by no means true for *all* parameters. We have shown in this paper that end-to-end available bandwidth is one such well-behaved parameter. It is an important issue for future network research to carefully identify precisely which parameters can be effectively measured and which cannot.

7. ACKNOWLEDGMENTS

The authors are grateful to Andreas Johnsson, Bob Melander and Mats Björkman from the TOPP team, and the other members of our Vinnova-financed EvaluNet project group for valuable discussions. This work is also in part sponsored by the European Community's FP6 Information Society Technologies programme under contract IST-001935, Evergrow, with the URL www.evergrow.org.

REFERENCES

[1] V. Paxson and S. Floyd, "*Wide-area traffic: The failure of Poisson modeling*," IEEE/ACM Trans. on Networking, vol. 3, pp. 226-244, 1995.

[2] G. Bishop, and G. Welch, "An Introduction to the Kalman Filter," TR 95-041, Department of Computer Science, University of North Carolina at Chapel Hill, NC, USA. 2004.

[3] B. Melander, M. Björkman, and P. Gunningberg, "A new end-to-end probing and analysis method for estimating bandwidth bottleecks", in *Proceedings of IEEE Globecomm '00*, San Francisco, USA, November 2000.

[4] A. Johnsson, B. Melander and M. Björkman, "On the analysis of packet train probing schemes", *International Conference on Communications in Computing (CIC) 2004*, Las Vegas, USA, June 2004.

[5] S. Ekelin and M. Nilsson, "Continuous Monitoring of Available Bandwidth over a Network Path", in *Proc. 2nd Swedish National Computer Networking Workshop (SNCNW'04)*, pp. 31-36, Karlstad, Sweden, October 2004.

Diagram 1: Estimated ABW for uniformly distributed load

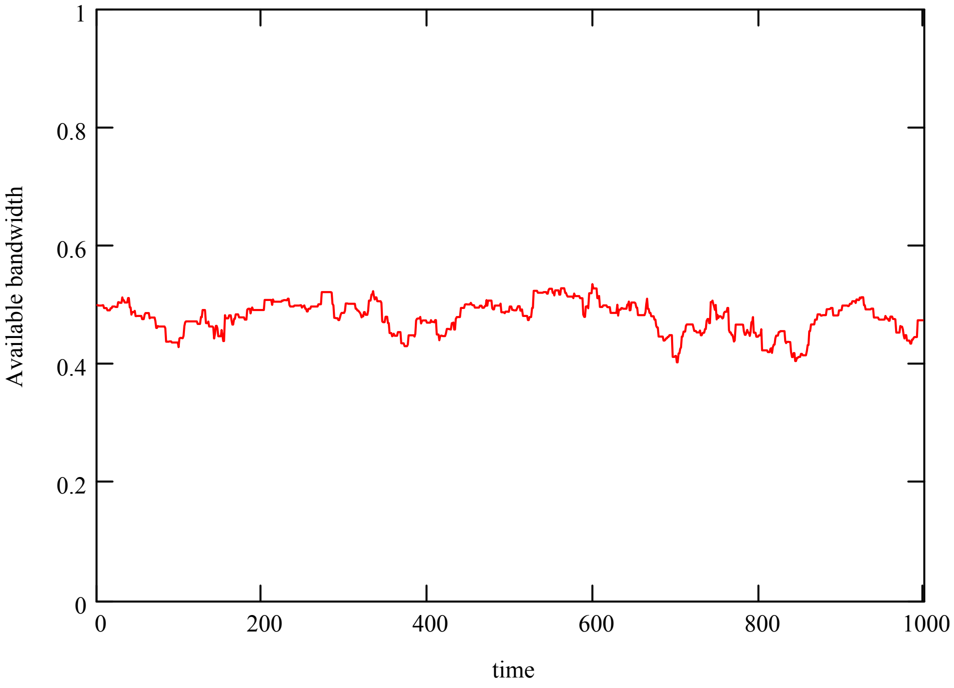

Diagram 2: Estimated ABW for Pareto-distributed load