ON THE CORRELATION OF TCP TRAFFIC IN BACKBONE NETWORKS

H. Nguyen [*], *P.Thiran*

EPFL: Swiss Federal Institute of Technology
LCA-ISC-I&C (Institute of Communication Systems)
CH-1015 Lausanne, Switzerland

C. Barakat [†]

INRIA - Planete group
06902 Sophia Antipolis
France

ABSTRACT

In this paper, we study the second order statistics of traffic in an Internet backbone. We model the traffic at the flow level by a Poisson shot noise process. This model is quite parasimonious, and is driven only by variables that can be easily obtained from measurements, namely flow sizes, durations and arrival rate. We consider the auto-correlation of TCP traffic where the loss process of each TCP connection is assumed to be Poisson. Using a stochastic differential equation, we are able to provide an upper bound on the auto-covariance function of the aggregated TCP traffic whose tightness is shown by simulations with the network simulator-*ns*.

1. INTRODUCTION

Second order statistics of network traffic, namely the auto-correlation, plays an important role in network performance evaluation (e.g., [1]) and traffic modelling (e.g., [2]). There is a huge amount of work in the literature which investigates the second order characteristics of network traffic (see [3] and the references therein). A recent trend in studying network traffic is to model traffic as fluid flows where each flow is a stream of packets which have common source and/or destination [4, 5]. Modelling traffic at the packet level is very difficult, since traffic on a link is the result of a high level of multiplexing of numerous flows whose behavior is strongly influenced by the transport protocol and by the application. From a simplicity stand point, it is much easier to monitor flows than to monitor packets in routers. The instantaneous rate of a flow depends on the network and transport protocol dynamics. Transmission Control Protocol (TCP) flows, which account for about 95% of the Internet traffic, can have a highly varying instantaneous rate due to the complexities in the mechanisms that TCP employs. However, the majority of work on auto-correlation of network traffic in the literature concentrates on either studying the correlation of a single TCP connection [3], or of the aggregated traffic in which flows have simple shapes such as rectangle.

The objective of this work is to study the auto-correlation of traffic in an Internet backbone where links are usually not congested. To this end, we use a Poisson shot noise process where a flow is a generic notion which can be a TCP flow or a UDP stream. We will concentrate only on the auto-correlation of traffic where all flows are TCP connections. Instead of using Markovian models for studying auto-correlation of TCP traffic (e.g., [3]), our approach uses a stochastic differential equation to obtain a tight bound on the auto-covariance function of TCP flows. We find an upper bound on the auto-correlation of the aggregated TCP traffic. This upper bound is a function of only a few flow parameters such as flow arrival rate, flow sizes and flow durations. These parameters can be easily computed by a router (e.g., using a tool such as NetFlow, which provides flow information in Cisco routers). Such a bound on the auto-covariance function of backbone traffic is useful for backbone operators who currently have only very basic information about the traffic.

The rest of the paper is organized as follows. Section 2 gives a summary of the shot noise model that was developed in previous work. Section 3 presents the auto-correlation of TCP traffic. Section 4 contains the simulation results and Section 5 concludes the paper.

2. SHOT NOISE MODEL

Our traffic model in this paper is the shot-noise model which was developed by Barakat et.al. in [6, 4]. In [6], one definition of flow is a stream of packets having the same source and destination IP addresses, same source and destination port numbers, and the same protocol number. Alternatively, a flow can also be defined as a stream of packets having the same /24 destination address prefix (i.e., only the 24 most significant bits).

Let T_n, $n \in \mathbf{Z}$, denote the arrival time of the n-th flow to the backbone link under consideration. Let S_n and D_n denote the size and duration of the n-th flow. The size of a

[*] H.Nguyen is supported by grant DICS 1830 of the Hasler Foundation, Bern, Switzerland.

[†] C. Barakat performed the work while at EPFL

flow is the volume of data it transports during its lifetime. Let $X_n(t-T_n)$ denote the rate of the n-th flow at time t (e.g., in bits/s), with $X_n(t-T_n)$ equals to zero for $t < T_n$ and for $t > T_n + D_n$. We call $X_n(.)$ the flow rate function or shot. The total rate of data through the backbone link, which we denote by $R(t)$, is the result of multiplexing of all shots: $R(t) = \sum_{n\in\mathbf{Z}} X_n(t-T_n)$. Flow rate functions $\{X_n(.)\}$ are assumed to be independent and identically distributed. These assumptions hold relatively well on a backbone link [6]. The total data rate $R(t)$ can be seen as a shot-noise process [7], where the term "shot" is synonymous of the "flow rate function".

Assume that traffic flows arrive at the backbone link as a homogeneous Poisson process of finite rate. Using elements from theory of Poisson shot-noise process, the authors in [6] found that the auto-covariance function of the aggregated traffic can be expressed as

$$\begin{aligned} C_R(\tau) &= \mathbf{E}[R(t-\tau)R(t)] - \mathbf{E}[R^2(t)] \\ &= \lambda \mathbf{E}\left[\mathbf{1}_{\{D_n>|\tau|\}} \int_0^{D_n-|\tau|} X_n(u)X_n(u+|\tau|)du\right], \end{aligned} \quad (1)$$

Where $\mathbf{1}_A$ is the indicator function of event A.

3. CORRELATION STRUCTURE OF TCP TRAFFIC

The auto-correlation of the traffic strongly depends on the shot shape we consider, and thus on the dynamics of the flow rate, which in turn depends on many factors such as the definition of flows, the transport mechanism, the application nature, etc. In some important cases, we can make use of the protocol information to derive the shot function $X_n(.)$. The most typical example is TCP, whose dynamics shape the flows and can be captured by analytical models. TCP is a window-based flow control protocol that provides reliable end-to-end communication in data networks. It is designed to adapt to the various traffic conditions of the network: a TCP connection progressively increases its transmission rate until it receives some indication that the capacity along its path is almost fully utilized. On the other hand, when the network cannot accommodate the traffic, the data rate of the connections is reduced. More specifically, the transmission rate of a TCP connection is governed by the additive-increase multiplicative-decrease (AIMD) mechanism which is as follows. Between congestion events (we also call them loss events, since they are usually the times at which a packet loss is detected at the sender), the rate of a TCP connection increases linearly with a slope α, which is inversely proportional to the square of the round-trip time (RTT) of the connection. At the congestion events, the rate of a TCP connection is reduced by half. Precisely, α is related to the RTT by: $\alpha = 1/bRTT^2$ where b is the acknowledgement factor which indicates how many packets are included in one acknowledgement [8] (typically, $b = 2$).

In this work, we assume that congestion events have a Poisson distribution with intensity λ_ℓ. This assumption may not hold in practice, but is needed for the following analytical derivation.

Assume that all traffic flows are long-lived TCP flows, i.e. the congestion avoidance phase is dominant over the life time of the flows. Rewrite equation (1), we get

$$\begin{aligned} C_R(\tau) &= \lambda \mathbf{P}(D_n > |\tau|) \\ & \mathbf{E}\left[\left. \int_0^{D_n-|\tau|} X_n(u)X_n(u+|\tau|)\mathrm{d}u \right| D_n > |\tau|\right]. \end{aligned} \quad (2)$$

Now if we condition on $D_n = d$ with $d > |\tau|$ then

$$\begin{aligned} &\mathbf{E}\left[\left. \int_0^{D_n-|\tau|} X_n(u)X_n(u+|\tau|)\mathrm{d}u \right| D_n = d, d > |\tau|\right] \\ &= \int_0^{d-|\tau|} \mathbf{E}_d\left[X_n(u)X_n(u+|\tau|)\,|\, d>|\tau|\right]\mathrm{d}u \\ &= (d-|\tau|)\mathbf{E}_d\left[X_n(t)X_n(t+|\tau|)\,|\, d>|\tau|\right]. \end{aligned}$$

Notation $\mathbf{E}_d$ is used to indicate the expected value is calculated under the condition $D_n = d$. The last equality is derived under the assumption that the TCP rate is in its stationary regime at time 0 so that $\mathbf{E}_d\left[X_n(t)X_n(t+|\tau|)\right]$ does not depend on time t.

From here on, we do the calculation for $\tau > 0$. For $\tau < 0$, the calculation is similar. Let $r_{d|\tau}(\tau) = \mathbf{E}_{d|\tau}[X_n(t)X_n(t+\tau)]$, where the subscript $d \mid \tau$ indicates the expected value is calculated under the conditions $D_n = d$, and $d > \tau$. Let

$$\begin{aligned} \mathrm{d}r_{d|\tau}(\tau) &= r_{d|\tau}(\tau+\mathrm{d}\tau) - r_{d|\tau}(\tau) \\ &= \mathbf{E}_{d|\tau}[X_n(t)(X_n(t+\tau+\mathrm{d}\tau) - X_n(t+\tau))]. \end{aligned}$$

We now consider the condition of occurrence of loss in $[t+\tau, t+\tau+\mathrm{d}\tau]$. When a loss appears in $[t+\tau, t+\tau+\mathrm{d}\tau]$, the rate is divided by 2, so

$$X_n(t+\tau+\mathrm{d}\tau) - X_n(t+\tau) = -\frac{X_n(t+\tau)}{2} + \alpha\mathrm{d}\tau,$$

Whereas when there is no loss in $[t+\tau, t+\tau+\mathrm{d}\tau]$, we have

$$X_n(t+\tau+\mathrm{d}\tau) - X_n(t+\tau) = \alpha\mathrm{d}\tau.$$

Since the loss process is Poisson with rate λ_ℓ, the probability that a loss appears in $[t+\tau, t+\tau+\mathrm{d}\tau]$ is independent of $X_n(.)$, and is equal to $\lambda_\ell\mathrm{d}\tau$. If all the flows have the same RTT, and thus the same α, it follows that

$$\begin{aligned} \mathrm{d}r_{d|\tau}(\tau) &= -\frac{\mathbf{E}_{d|\tau}[X_n(t)X_n(t+\tau)]}{2}\lambda_\ell\mathrm{d}\tau \\ & +\mathbf{E}_{d|\tau}[X_n(t)]\alpha\mathrm{d}\tau. \end{aligned} \quad (3)$$

From [8], we have for a Poisson loss process: $\mathbf{E}_{d|\tau}[X_n(t)] = 2\alpha/\lambda_\ell$.

Inserting this value in (3), we obtain the following ordinary differential equation

$$\frac{\mathrm{d}r_{d|\tau}(\tau)}{\mathrm{d}\tau} + \frac{\lambda_\ell}{2} r_{d|\tau}(\tau) = \frac{2\alpha^2}{\lambda_\ell}.$$

From [4], we have the initial condition for Poisson losses $r_{d|\tau}(0) = \mathbf{E}_{d|\tau}[X_n^2(t)] = 4\mathbf{E}_{d|\tau}[S_n]^2/3d^2$.

Solving the above equation and simplifying the result using $\mathbf{E}_{d|\tau}[X_n(t)] = \mathbf{E}_{d|\tau}[S_n]/d$ ([4]), we obtain:

$$r_{d|\tau}(\tau) = \frac{1}{3}\frac{\mathbf{E}_{d|\tau}[S_n]^2}{d^2} e^{-\alpha\tau\frac{d}{\mathbf{E}_{d|\tau}[S_n]}} + \frac{\mathbf{E}_{d|\tau}[S_n]^2}{d^2}.$$

Let us define $f_{d|\tau}$ as the probability density function (pdf) of the random variables D_n whose probability space is $[\tau, \infty]$. Substituting the expression of $r_{d|\tau}$ in (2), we get

$$\begin{aligned} C_R(\tau) &= \lambda\mathbf{P}(D_n > \tau)\int_\tau^\infty (d-\tau) \\ &\left(\frac{1}{3}\frac{\mathbf{E}_{d|\tau}[S_n]^2}{d^2} e^{-\alpha\tau\frac{d}{\mathbf{E}_{d|\tau}[S_n]}} + \frac{\mathbf{E}_{d|\tau}[S_n]^2}{d^2}\right) f_{d|\tau}\mathrm{d}d. \end{aligned}$$

Since the function $f(x) = x^2 e^{-\alpha\tau d/x}/3d^2$ is convex for all $x > 0$, we have:

$$\frac{1}{3}\frac{\mathbf{E}_{d|\tau}[S_n^2]}{d^2} e^{-\alpha\tau\frac{d}{\mathbf{E}_{d|\tau}[S_n]}} \leq \mathbf{E}_{d|\tau}[\frac{1}{3}\frac{S_n^2}{d^2} e^{-\alpha\tau\frac{d}{S_n}}].$$

Thus,

$$\begin{aligned} C_R(\tau) &\leq \lambda\mathbf{P}(D_n > \tau) \\ &\int_\tau^\infty (d-\tau)\left(\mathbf{E}_{d|\tau}[\frac{1}{3}\frac{S_n^2}{d^2} e^{-\alpha\tau\frac{d}{S_n}}] + \frac{\mathbf{E}_d[S_n^2]}{d^2}\right) f_{d|\tau}\mathrm{d}d \\ &\leq \lambda\mathbf{P}(D_n > \tau)\Bigg(\frac{1}{3}\mathbf{E}\left[\frac{S_n^2(D_n-\tau)}{D_n^2} e^{-\alpha\tau\frac{D_n}{S_n}}\middle| D_n > \tau\right] \\ &\quad + \mathbf{E}\left[\frac{S_n^2(D_n-\tau)}{D_n^2}\middle| D_n > \tau\right]\Bigg). \end{aligned}$$

As $C_R(\tau)$ is an even function, we have the general result:

$$\begin{aligned} C_R(\tau) &\leq \lambda\mathbf{P}(D_n > |\tau|) \\ &\quad \left(\frac{1}{3}\mathbf{E}\left[\frac{S_n^2(D_n-|\tau|)}{D_n^2} e^{-\alpha|\tau|\frac{D_n}{S_n}}\middle| D_n > |\tau|\right]\right. \\ &\quad + \left.\mathbf{E}\left[\frac{S_n^2(D_n-|\tau|)}{D_n^2}\middle| D_n > |\tau|\right]\right). \end{aligned} \tag{4}$$

There are two separate terms in the bound. The first term is an exponential function of τ, which vanishes quickly for large values of τ. The second term is a linear function of τ, which decreases slowly as τ increases. As a result, the bound has two distinct behaviors: one for small values of τ, when the exponential term dominates and one for large values of τ, when the linear term dominates. Furthermore, the loss rate λ_l does not figure in the bound and the round-trip time only appears in the exponential term via α.

4. SIMULATION RESULTS

We present a validation of the bound on the auto-covariance function of TCP traffic by simulation. We use the *ns* simulator to study two different scenarios. In the first scenarios, all flows have the same size but different durations. In the second scenario, both sizes and durations of flows change during each simulation.

4.1. Simulation scenario

In our simulations, a set of TCP Newreno flows transmit files over a 10 Mbps link which corresponds to the backbone link. Each flow transmits one file and all flows cross the backbone link in the same direction. The duration of each simulation is equal to 1000 seconds. Delayed acknowledgement option of TCP is enabled and each packet has a size of 500 bytes. Before arriving on the backbone link, all flows experience some packet losses with a probability of 3% (to introduce randomness in the durations of flows). TCP flows are generated according to a Poisson process. The rate of the Poisson process and the file sizes are chosen such that the 10Mbps backbone link always remains under-utilized. The round trip time of all TCP flows is set to 80ms. We compute the rate with which data cross the 10Mpbs link and store the variation of this rate as a function of time. This rate is used to calculate the auto-covariance of TCP traffic. We also measure the size and duration of each flow, which produces samples for the processes $\{D_n\}$ and $\{S_n\}$. The instantaneous rate $R(t)$ is measured by averaging the number of bytes that cross the 10Mpbs link over the interval of 100ms. In each simulation, we plot the auto-covariance of the real simulation traffic, and the upper bound of the auto-covariance for TCP traffic (4). In our simulations, the RTT was set to 80ms. α will then be: $\alpha = 1/(bRTT^2) = 500 * 8/(2 * 0.08^2)(bits/sec^2)$.

4.2. Constant-size flows

We set the arrival rate of TCP flows to 2 flows per second and we give file sizes constant values equal to 25Kbytes, 50Kbytes, 100Kbytes, 250Kbytes and 500Kbytes. We run a set of 10 simulations for each value of the file size. In each simulation, all files have the same size. The average of the 10 values obtained from simulations are plotted. The 95% confidence intervals for the actual traffic are also plotted, while the 95% confidence interval for the approximations are omitted because they are too narrow. The results are plotted in Figures 1(a) to (e).

4.3. Variable-size flows

We repeat the previous simulations with variable file sizes. To generate variable sizes, for each flow we pick a real num-

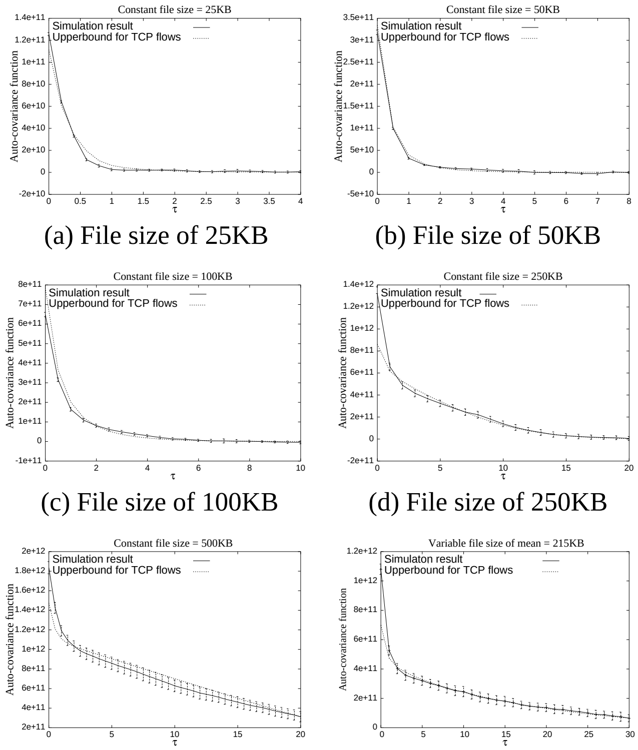

(a) File size of 25KB (b) File size of 50KB

(c) File size of 100KB (d) File size of 250KB

(e) File size of 500KB (f) Variable file size

Fig. 1. Simulation results showing the auto-covariance function $C_R(\tau)$ versus time lag τ.

ber randomly between 1 and 3 with a uniform distribution. The size of the flow in Kbytes is 10 to the power of the selected number. This way we get an average file size of 215KB. The results are plotted in Figure 1(f).

4.4. Observations

From the simulation results, we can observe that for all values of τ, the gap between the upper bound in (4) and the real auto-covariance of the traffic is very small. For large values of τ, both the upper bound and the auto-covariance of the simulated traffic decrease slowly as a function of τ with a rate close to linear. This is consistent with our finding in (4), where for large value of τ the exponential term vanishes and only the linear term contributes to the bound. For small values of τ, both the auto-covariance of the simulated traffic and the upper bound decrease exponentially fast as a function of τ. This can be explained by the domination of the exponential term in the bound. Note here that the bound (4) is obtained under some strong assumptions (i.e., losses are Poisson, and other approximations detailed in [6]), which may not hold in practice. This explains the instances in our simulation results where the bound (4) is not respected.

5. CONCLUSION

In this paper, we use a Poisson shot noise model to study the auto-correlation of Internet traffic in non-congested backbone links. We provide an upper bound on the correlation of aggregated TCP traffic where all flows are long-lived TCP flows. The upper bound is a function of only three flow parameters (λ: arrival rate, D_n flow duration, and S_n flow size) which can be obtained from passive measurements quite easily [4]. Such a bound can be used in network dimensioning and management to study the impact of flow arrival, flow sizes and durations on the auto-correlation of the traffic and hence on dimensioning the backbone.

6. REFERENCES

[1] F.Park, G.Kim, and M.Crovella, "On the effect of traffic self-similarity on network performance," *Proc. of SPIE International Conference on the Performance and Control of Netwokk systems*, Nov 1997.

[2] M.Krunz and A.Makowski, "Modelling video traffic using M/G/Infinity input processes: A compromise between Markovian and LRD models," *IEEE JSAC*, vol. 16, pp. 733–748, Jun 1998.

[3] D. R. Figueiredo, B. Liu, V. Misra, and D. Towsley, "On the autocorrelation structure of TCP traffic," *'Computer Networks Journal' Special Issue on 'Advances in Modeling and Engineering of Long-Range Dependent Traffic'*, 2002.

[4] C. Barakat, P. Thiran, G. Iannaccone, C. Diot, and P. Owezarski, "Modeling Internet backbone traffic at the flow level," *IEEE Transactions on Signal processing*, vol. 51, no. 8, Aug 2003.

[5] N. Hohn, D. Veitch, and P. Abry, "Cluster processes: a natural language for network traffic," *IEEE Transactions on Signal processing*, vol. 51, no. 8, Aug 2003.

[6] C. Barakat, P. Thiran, G. Iannaccone, C. Diot, and P. Owezarski, "A flow-based model for Internet backbone traffic," *ACM SIGCOMM Internet Measurement Workshop*, Nov 2002.

[7] D.Daley and D.Vere-Jones, *An introduction to the theory of point processes*, New York: Springer-Verlag, 1988.

[8] E. Altman, K. Avrachenkov, and C. Barakat, "A stochastic model for TCP/IP with stationary random losses," *ACM SIGCOMM*, Sep. 2000.