TCP in Presence of Bursty Losses

Eitan Altman, Konstantin Avrachenkov, Chadi Barakat
INRIA, 2004 route des Lucioles, 06902 Sophia Antipolis, France
Email:{altman,kavratch,cbarakat}@sophia.inria.fr

ABSTRACT

In this paper we analyze the performance of a TCP-like flow control in a lossy environment. The transmission rate in the control scheme that we consider has a linear growth rate; whenever a loss occurs, the transmission rate is halved. This approximates the performance of several versions of TCP that divide their congestion window by two when a loss is detected. We propose a mathematical model that allows to account for burstiness in the loss process. We compute the expected transmission rate and its moments at some potential loss instants, and provide a useful implicit expression for the Laplace Stieltjis Transform. This allows us to compute explicitly the time average of the transmission rate as well as its moments. We show that the time average of the transmission rate is indeed sensitive to the distribution of losses, and not just to the average loss rate: for a given average loss rate, we show that the time average of the transmission rate *increases* with the burstiness of losses. We finally examine the impact of burstiness of losses on the transmission rate variability.

1. INTRODUCTION

Flow control mechanisms in the Internet, particularly those of the Transmission Control Protocol [12], use the loss of packets as an indication of network congestion. In general, the transmission rate of the controlled flow is linearly increased until a loss occurs. The network is supposed here to be congested and the transmission rate is multiplicatively decreased in order to alleviate this congestion. In TCP as an example, the transmission rate is controlled via a congestion window which is increased in absence of losses and decreased upon loss detection [17]. Another set of flow control mechanisms recently introduced into multimedia applications consist in measuring the loss rate of packets and in controlling the transmission rate in a way to be friendly with TCP transfers [9]. Explicit expressions for TCP throughput for a given loss rate (e.g.[16]) are used for this purpose.

A good understanding of the impact of a loss process on the performance of a flow control mechanism is required for a good network and protocol tuning. Several previous works have addressed the problem of TCP performance as a function of data packet losses. The focus on TCP is due to the dominance of TCP traffic in today's Internet. Some of these works [8, 13, 14, 16] have studied the impact of the intensity of losses (or the average loss rate) on the performance. TCP packets are assumed to be lost independently with the same probability. Explicit expressions for the throughput of the TCP connection are derived by simply dividing the average number of packets transmitted between losses to the average time between losses. No other parameter than the packet loss probability is used to characterize the distribution of loss instants over time. As we will see in this paper, this will cause a wrong estimation of the throughput when a certain burstiness of losses exists. Other works [1, 4, 5] have addressed the problem of burstiness of packet losses but in the wireless environment context. It is known that due to multiple phenomena such as multipath fading [6], wireless links as those found in terrestrial wireless networks or satellite networks present a certain degree of transmission error burstiness. The impact of consecutive packet losses on the different versions of TCP is studied in these works [1, 4, 5]. They model losses with a two-state Markov chain where small bursts of losses appear in an independent and uniform manner. They study then TCP performance as a function of the average rate of bursts as well as the average burst size. But, the new versions of TCP (New Reno, SACK) [7] are able to resist to consecutive packet losses and to reduce their window once for all losses in the same Round Trip Time. This will result in these models becoming similar to the previous ones since they study the impact on TCP performance of only the average rate at which the window is reduced.

In this paper, we propose a completely different model for a TCP-like flow control protocol that, in addition to the average rate of window reduction events, it accounts for the burstiness of these events. Rather than looking at the packet level and considering the probability that a packet is lost, we look at the transmission rate level and look at instants when the transmission rate is reduced. We associate then a new loss process to the moments at which the transmission rate is reduced. A loss event is equivalent to a transmission rate reduction event. This loss can be the result of a single packet loss or multiple consecutive packet losses during the same RTT. Our aim is to study the impact of burstiness of this new loss process on the throughput of the connection. We compute the expectation of the transmission rate as well as an implicit expression for its Laplace Stieltjis Transform at some potential loss instants. This allows us to compute the time average of the transmission rate which we call the throughput of the transfer. We show that this throughput is indeed sensitive to the distribution of losses, and not just to the average loss rate: for a given average loss rate, we show that the throughput *increases* with the burstiness. We finally examine the impact of burstiness on the transmission rate variability. Our results are compared to simulations done with the `ns` simulator developed at LBNL [15] and a good match is reported.

The structure of the paper is as follows. In the next section we present our model for losses and for the controlled rate. Section 3 contains our analysis of the performance of the transmission rate in presence of losses. At the end of this section, we give the general expression of the throughput. The throughput in the case of an independent loss process having the same average loss rate as a bursty loss process is defined. This second throughput is then used as a reference to show the effect of burstiness. In section 4, we study the impact of the parameters of the loss process on the performance. The analytical results are compared to simulation ones. The paper is concluded in section 5. In the Appendix, we prove the convergence of the dynamics to a unique stationary regime.

2. THE MODEL

Consider a flow control mechanism where the transmission rate grows linearly at a rate α per unit of time. The growth continues until a loss occurs. The transmission rate is halved and the linear growth is then resumed. This model approximates the performance of several versions of TCP/IP where the transmission rate at any instant is equal to the window size divided by the RTT and where the window increases linearly by 1 packet every RTT [12]. If the delay ACKs mechanism is implemented at the destination, the increase in TCP window is by one packet every two RTTs. This linear window increase corresponds to the congestion avoidance mode of TCP. The slow start mode is neglected in this paper due its fast exponential window increase. The model can also approximate any additive-increase multiplicative-decrease flow control mechanism.

Let us propose a model which accounts for burstiness of losses. The Gilbert model is often used in this context [6]. The path between the source and the destination called channel in the wireless terminology is assumed to have two states: *Good* and *Bad*; losses are assumed to occur in the Bad state. The time during which the channel is in a Good or in a Bad state is taken to be geometrically (or exponentially) distributed. We propose an extension of this model in order to handle generally distributed periods of Good and Bad states. Our model is related to the MAP (Markovian Arrival Process) process [11]. We allow losses to occur both in the Good state as well as in the Bad state; the occurrence of losses in each of these states is different. To that end, we define a series of potential losses. Let t_n denote the time at which the nth potential loss may occur. Let $D_n, n = 1, 2, \ldots$ be the sequence of times between potential losses: $D_n = t_{n+1} - t_n$. D_n are assumed to be i.i.d. with expectation d, second moment $d^{(2)}$ and Laplace Stieltjis Transform $D^*(s)$. Let X_n be the transmission rate just prior to the instant of the nth potential loss.

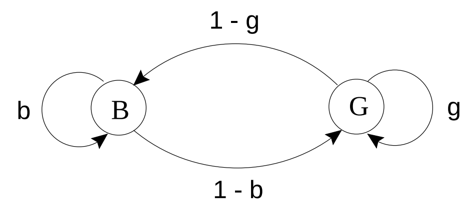

Figure 1: The Markov chain associated to the path

Potential losses are transformed to real losses with a certain probability. This is similar to MAP processes in which at each state transition an arrival can occur with a probability that depends on the state. Let Y_n be the state of the channel at the nth potential loss instant. We consider the states B (for Bad) in which a potential loss is transformed to a real loss with probability p_B, and G, (for Good) in which it is transformed with a smaller probability p_G. We shall assume throughout that $p_G \leq p_B$ and that $p_B > 0$. We assume further that the sequences $\{Y_n\}$ and $\{D_n\}$ are independent.

Y_n is assumed to be a Markov chain with the following transition probabilities (Figure 1),

$$\begin{aligned} P(Y_{n+1} = G | Y_n = G) &= g \\ P(Y_{n+1} = B | Y_n = G) &= \bar{g} = 1 - g \\ P(Y_{n+1} = B | Y_n = B) &= b \\ P(Y_{n+1} = G | Y_n = B) &= \bar{b} = 1 - b \end{aligned}$$

We shall assume throughout that $g, b \in (0, 1)$. $\{Y_n\}_{n=1}^{+\infty}$ is then ergodic with stationary probabilities,

$$\pi_G = \frac{1-b}{2-b-g} = \frac{\bar{b}}{\bar{b}+\bar{g}}, \qquad \pi_B = \frac{1-g}{2-b-g} = \frac{\bar{g}}{\bar{b}+\bar{g}}$$

The average loss rate is given by,

$$R = \frac{p_G \pi_G + p_B \pi_B}{d}.$$

This is equal to the average number of times the source reduces its rate per unit of time. With this loss process, we are able to vary the average loss rate as well as the the burstiness of losses. Let us now calculate the throughput under the above described loss process.

3. PERFORMANCE ANALYSIS

Define the two random variables U_n and V_n describing the behavior of X_n when a potential loss occurs. They correspond to the two states of the channel. A value one of these variables means that the potential loss causes really a reduction in X_n. A value zero however means that X_n is not affected (i.e. a real loss didn't occur). We have,

$$\begin{aligned} P(U_n = 1) = p_G, \quad P(U_n = 0) = 1 - p_G \\ P(V_n = 1) = p_B, \quad P(V_n = 0) = 1 - p_B \end{aligned}$$

The evolution of the transmission rate is the following,

$$\begin{aligned} X_{n+1} &= (1-U_n)X_n 1\{Y_n = G\} + U_n \frac{X_n}{2} 1\{Y_n = G\} \\ &+ (1-V_n)X_n 1\{Y_n = B\} + V_n \frac{X_n}{2} 1\{Y_n = B\} \\ &+ \alpha D_n \end{aligned}$$

$$= \left(1 - \frac{U_n}{2}\right) X_n 1\{Y_n = G\}$$
$$+ \left(1 - \frac{V_n}{2}\right) X_n 1\{Y_n = B\} + \alpha D_n \quad (1)$$

1{A} is the indicator function which is equal to 1 if expression A is true and to 0 otherwise.

We begin by computing the first moments of X_n in the stationary regime. The calculation of these moments allows us to calculate the time average of the transmission rate, called also the throughput. We denote,

$$x_G = \lim_{n \to +\infty} E[X_n 1\{Y_n = G\}]$$

$$x_B = \lim_{n \to +\infty} E[X_n 1\{Y_n = B\}]$$

$$x = \lim_{n \to +\infty} E[X_n] = x_G + x_B$$

We show first that these moments exist, then we follow two approaches to calculate them. The first approach consists in (a) using the relationship between X_{n+1} and X_n given in equation (1) and (b) moving n to infinity. The second approach uses the Laplace Stieltjis Transforms of X_n.

Theorem 1: *The first moments of the transmission rate just prior to the potential loss occurrence converge and they are equal to,*

$$x_G = \alpha d \frac{\gamma_B(\pi_B - b) + \pi_G}{1 - \gamma_B b - \gamma_G g + \gamma_B \gamma_G (g + b - 1)}$$
$$x_B = \alpha d \frac{\gamma_G(\pi_G - g) + \pi_B}{1 - \gamma_B b - \gamma_G g + \gamma_B \gamma_G (g + b - 1)}$$

with,

$$\gamma_G = 1 - \frac{p_G}{2}, \qquad \gamma_B = 1 - \frac{p_B}{2}$$

Proof: First we write,

$$E[X_{n+1} 1\{Y_{n+1} = G\}]$$
$$= gE\left[1 - \frac{U_n}{2}\right] E[X_n 1\{Y_n = G\}]$$
$$+ \bar{b}E\left[1 - \frac{V_n}{2}\right] E[X_n 1\{Y_n = B\}] + \alpha d P(Y_{n+1} = G)$$
$$= g\gamma_G E[X_n 1\{Y_n = G\}]$$
$$+ \bar{b}\gamma_B E[X_n 1\{Y_n = B\}] + \alpha d P(Y_{n+1} = G) \quad (2)$$

$$E[X_{n+1} 1\{Y_{n+1} = B\}]$$
$$= \bar{g}E\left[1 - \frac{U_n}{2}\right] E[X_n 1\{Y_n = G\}]$$
$$+ bE\left[1 - \frac{V_n}{2}\right] E[X_n 1\{Y_n = B\}] + \alpha d P(Y_{n+1} = B)$$
$$= \bar{g}\gamma_G E[X_n 1\{Y_n = G\}]$$
$$+ b\gamma_B E[X_n 1\{Y_n = B\}] + \alpha d P(Y_{n+1} = B) \quad (3)$$

Then we show that the sequences $E[X_n 1\{Y_n = G\}]$ and $E[X_n 1\{Y_n = B\}]$, $n = 0, 1, ...$ converge. We rewrite the recursive equations (2) and (3) in the following matrix form,

$$\mathcal{X}_{n+1} = \Pi \mathcal{X}_n = \begin{bmatrix} \gamma_G g & \gamma_B \bar{b} & \alpha d g & \alpha d \bar{b} \\ \gamma_G \bar{g} & \gamma_B b & \alpha d \bar{g} & \alpha d b \\ 0 & 0 & g & \bar{b} \\ 0 & 0 & \bar{g} & b \end{bmatrix} \mathcal{X}_n$$

with,

$$\mathcal{X}_n = \begin{bmatrix} E[X_n 1\{Y_n = G\}] \\ E[X_n 1\{Y_n = B\}] \\ P(Y_n = G) \\ P(Y_n = B) \end{bmatrix}$$

The convergence properties of the augmented sequence $\mathcal{X}_n$, $n = 0, 1, ...$ depend on the spectrum of the transition matrix Π. In turn, the spectrum of Π is formed by the spectra of the matrices,

$$P = \begin{bmatrix} g & \bar{b} \\ \bar{g} & b \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} \gamma_G g & \gamma_B \bar{b} \\ \gamma_G \bar{g} & \gamma_B b \end{bmatrix}$$

Since P is the transition matrix of the Markov chain $\{Y_n\}$, it always has an eigenvalue which is equal to one. Further, because $\{Y_n\}$ is ergodic, P has only one unit eigenvalue. Let us show that all the eigenvalues of matrix A are less than 1 if g and b are less than 1. Actually, the eigenvalues of A can be given in the closed analytic form,

$$\lambda_{1,2} = \frac{(g\gamma_G + b\gamma_B)}{2} \pm \frac{\sqrt{(g\gamma_G + b\gamma_B)^2 - 4(gb\gamma_G\gamma_B - \bar{g}\bar{b}\gamma_G\gamma_B)}}{2}$$

The term under the root square is always positive. Indeed we write,

$$(g\gamma_G + b\gamma_B)^2 - 4(gb\gamma_G\gamma_B - \bar{g}\bar{b}\gamma_G\gamma_B)$$

$$= (g\gamma_G - b\gamma_B)^2 + 4\bar{g}\bar{b}\gamma_G\gamma_B$$

which is always positive since,

$$0 < b < 1, \qquad 0 < g < 1$$
$$1/2 \le \gamma_G \le 1, \qquad 1/2 \le \gamma_B < 1$$

The smallest eigenvalue (λ_1) is smaller than one given that $(g\gamma_G + b\gamma_B)/2$ is. For the other eigenvalue (λ_2) we write,

$$\lambda_2 < 1$$

i.e. $(g\gamma_G + b\gamma_B)^2 - 4(gb\gamma_G\gamma_B - \bar{g}\bar{b}\gamma_G\gamma_B) < (2 - (g\gamma_G + b\gamma_B))^2$

$$\text{i.e. } \gamma_G\gamma_B + g\gamma_G(1 - \gamma_B) + b\gamma_B(1 - \gamma_G) < 1$$

Since $g < 1$ and $b < 1$, the left-hand term of the latter equation is smaller than,

$$\gamma_G\gamma_B + \gamma_G(1 - \gamma_B) + \gamma_B(1 - \gamma_G)$$

which can be written as

$$\gamma_B(1 - \gamma_G) + \gamma_G.$$

Since we assumed that γ_B is strictly smaller than 1, this latter equation is also strictly smaller than 1. Thus, $p_G > 0$, $g < 1$ and $b < 1$ imply that $\lambda_{1,2} < 1$. We conclude that the transition matrix Π of the augmented sequence has only simple unit eigenvalue (it is of multiplicity one). This implies that the powers of matrix Π converge to its eigenprojection corresponding to the unit eigenvalue. However, it is more simple to compute the limits of sequences $E[X_n 1\{Y_n = G\}]$ and $E[X_n 1\{Y_n = B\}], n = 0, 1, ...$ by moving n to $+\infty$ in the original equations (2) and (3). We get,

$$x_G = \gamma_G g x_G + \gamma_B \bar{b} x_B + \alpha d \pi_G$$
$$x_B = \gamma_B \bar{g} x_G + \gamma_B b x_B + \alpha d \pi_B$$

The solution of this system in x_G and x_B concludes the proof. $\diamond$

Remark: *The existence of x_G and x_B means the existence of x. In the Appendix we show that in fact X_n converges to a unique stationary regime.*

3.1 Laplace Transforms and moments of X_n

Define the following Laplace Stieltjis Transforms,

$$\begin{aligned} Z(s,G) &= E\left[e^{-sX_n}1\{Y_n=G\}\right] \\ Z(s,B) &= E\left[e^{-sX_n}1\{Y_n=B\}\right] \end{aligned}$$

where we assume that X_n is in the stationary regime.

Theorem 2: *The Laplace Stieltjis Transforms are the solutions of the following implicit equations,*

$$\begin{aligned} Z(s,G) &= D^*(\alpha s)\left[g(1-p_G)Z(s,G)+gp_G Z(s/2,G)\right] \\ &+ D^*(\alpha s)\left[\bar{b}(1-p_B)Z(s,B)+\bar{b}p_B Z(s/2,B)\right] \\ Z(s,B) &= D^*(\alpha s)\left[\bar{g}(1-p_G)Z(s,G)+\bar{g}p_G Z(s/2,G)\right] \\ &+ D^*(\alpha s)\left[b(1-p_B)Z(s,B)+bp_B Z(s/2,B)\right] \end{aligned}$$

Proof: We write,

$$\begin{aligned} & E\left[e^{-sX_{n+1}}1\{Y_{n+1}=G\}\right] \\ &= gE\left[e^{-s((1-\frac{U_n}{2})X_n+\alpha D_n)}1\{Y_n=G\}\right] \\ &+ \bar{b}E\left[e^{-s((1-\frac{V_n}{2})X_n+\alpha D_n)}1\{Y_n=B\}\right] \\ & E\left[e^{-sX_{n+1}}1\{Y_{n+1}=B\}\right] \\ &= \bar{g}E\left[e^{-s((1-\frac{U_n}{2})X_n+\alpha D_n)}1\{Y_n=G\}\right] \\ &+ bE\left[e^{-s((1-\frac{V_n}{2})X_n+\alpha D_n)}1\{Y_n=B\}\right] \end{aligned}$$

Using the fact that,

$$\begin{aligned} & E\left[e^{-s(1-\frac{U_n}{2})X_n}1\{Y_n=G\}\right] \\ &= (1-p_G)E\left[e^{-sX_n}1\{Y_n=G\}\right]+p_G E\left[e^{-s\frac{X_n}{2}}1\{Y_n=G\}\right], \\ & E\left[e^{-s(1-\frac{V_n}{2})X_n}1\{Y_n=B\}\right] \\ &= (1-p_B)E\left[e^{-sX_n}1\{Y_n=B\}\right]+p_B E\left[e^{-s\frac{X_n}{2}}1\{Y_n=B\}\right] \end{aligned}$$

and by taking the limit as n goes to infinity, we get the required relations. $\diamond$

Although the Laplace Stieltjis Transforms in Theorem 2 are only given as solutions of implicit equations, all moments of $X_n 1\{Y_n=G\}$ and $X_n 1\{Y_n=B\}$ (in the stationary regime) can be obtained explicitly. In particular, the first moments are no other than the opposite of the derivatives of $Z(s,G)$ and $Z(s,B)$ at $s=0$. By differentiating the implicit equations in Theorem 2 and substituting $s=0$, one can obtain a system of two linear equations with two unknowns, whose solution coincides with what we already obtained in Theorem 1. The calculation requires the following equations,

$$\begin{aligned} & Z(0,G)=\pi_G, && Z(0,B)=\pi_B \\ & D^*(0)=1, && \left.\tfrac{dD^*(\alpha s)}{ds}\right|_{s=0}=-\alpha d \end{aligned}$$

More general, the order k moments can be obtained in a similar way using,

$$\begin{aligned} E[X_n^k 1\{Y_n=G\}] &= (-1)^k \left.\frac{d^k Z(s,G)}{ds^k}\right|_{s=0} \\ E[X_n^k 1\{Y_n=B\}] &= (-1)^k \left.\frac{d^k Z(s,B)}{ds^k}\right|_{s=0} \end{aligned}$$

3.2 The average throughput

Theorem 3: *The throughput or the time average of the transmission rate we denote by $\bar{x}$ can be expressed as,*

$$\bar{x}=\lim_{t\to+\infty}\frac{1}{t}\int_0^t X(t)dt=\gamma_G x_G+\gamma_B x_B+\frac{1}{2}\alpha\frac{d^{(2)}}{d}$$

where x_G and x_B are given in Theorem 1, and where the limit holds almost surely.

Proof: We have almost surely,

$$\begin{aligned} \bar{x} &= \lim_{n\to+\infty}\frac{\sum_{i=0}^{i=n-1}\int_{t_i}^{t_{i+1}}X(t)dt}{\sum_{i=0}^{i=n-1}D_i} \\ &= \lim_{n\to+\infty}\frac{\frac{1}{n}\sum_{i=0}^{i=n-1}D_i(1-U_i/2)X_i 1\{Y_i=G\}}{\frac{1}{n}\sum_{i=0}^{i=n-1}D_i} \\ &+ \lim_{n\to+\infty}\frac{\frac{1}{n}\sum_{i=0}^{i=n-1}D_i(1-V_i/2)X_i 1\{Y_i=B\}}{\frac{1}{n}\sum_{i=0}^{i=n-1}D_i} \\ &+ \lim_{n\to+\infty}\frac{\frac{1}{n}\sum_{i=0}^{i=n-1}\frac{1}{2}\alpha D_i^2}{\frac{1}{n}\sum_{i=0}^{i=n-1}D_i} \qquad (4) \\ &= \lim_{n\to+\infty}\frac{E[1-U_n/2]E[X_n 1\{Y_n=G\}]E[D_n]}{E[D_n]} \\ &+ \lim_{n\to+\infty}\frac{E[1-V_n/2]E[X_n 1\{Y_n=B\}]E[D_n]}{E[D_n]} \\ &+ \lim_{n\to+\infty}\frac{\frac{1}{2}\alpha E[D_n^2]}{E[D_n]} \\ &= \gamma_G x_G+\gamma_B x_B+\frac{1}{2}\alpha\frac{d^{(2)}}{d} \end{aligned}$$

The equality that appears just after (4) can easily be shown to follow from the convergence of X_n to a stationary ergodic regime, which follows from the Appendix. This concludes the proof. $\diamond$

3.3 The reference throughput

To study the effect of burstiness, we change in what follows the parameters of the Markov chain (b and g) while keeping the average loss rate unchanged. The throughput in the bursty case is then compared to the throughput when the channel is subjected to a non-bursty loss process having the same average loss rate. We denote this latter throughput $\bar{x}_r$ and we use it as a reference to evaluate the impact of burstiness. A non-bursty loss process is obtained when we have the same loss probabilities in the two states. We call this probability p. To get the same average loss rate as in the bursty case, p must be equal to,

$$p=dR=p_G\pi_G+p_B\pi_B.$$

Lemma 1: *On a non-bursty path, the source achieves a throughput of*

$$\bar{x}_r=\frac{2-p}{p}\alpha d+\frac{1}{2}\alpha\frac{d^{(2)}}{d} \qquad (5)$$

Proof: This expression of $\bar{x}_r$ can be easily obtained by substituting in the expression of $\bar{x}$ (Theorem 3), γ_G and γ_B by their values as a function of p, the loss probability in the two states. We have,

$$\gamma = \gamma_G = \gamma_B = 1 - \frac{p}{2}.$$

The parameters of the Markov chain disappear and we get an expression of the reference throughput as a function of p and the distribution of potential losses. $\diamond$

3.4 Comparison with previous works

Consider a particular case where $p = 1$. In this case, all potential losses cause a reduction in the transmission rate. This forms a loss process similar (even more general) to the one used in many previous works [8, 14, 16]. These works suppose that in the stationary regime, TCP window (or the transmission rate) varies in a cyclic manner between two values X and $2X$. They found that the time average transmission rate is about $3X/2$. Our model shows well that in the presence of a non-bursty channel with $p = 1$, the expectation of the transmission rate just prior to a loss is equal to $2\alpha d$. The expectation of the transmission rate just after a real loss is simply αd. Thus, αd in our model corresponds to their X. However, our model doesn't give the same throughput they found. In our expression for the throughput, we see the appearance of the second moment of the time between losses ($d^{(2)}$). To get their result, the second moment of the time between losses must be equal to the square of its average. This is only the case for a deterministic inter-loss distribution of value d. Although they are using a probabilistic loss model, these works transform the loss process into a deterministic one which results in the disappearance of the term $d^{(2)}$ in their analysis. The second moments of X_n and D_n are taken equal to the square of their average rates. The difference in these works is that the packet loss probability is used for the calculation of these quantities. This deterministic evolution of the window can be seen as a normal result of the mutual independence that they assume between processes $\{X_n\}$ and $\{D_n\}$. Our model however, in addition to the consideration of the burstiness, considers the exact expression of the throughput. It shows that the average time between losses as well as the second moment of this time must be considered otherwise the throughput will be underestimated. As an example, in the case of an exponential loss distribution, $d^{(2)}$ is equal to $2d^2$ and the throughput is simply equal to the average transmission rate just prior to losses.

4. A CASE STUDY

In the sequel we consider the special case where,

$$p_G = 0, \gamma_G = 1, \qquad p_B = 1, \gamma_B = 1/2.$$

In other words we suppose that if the channel is in the Bad state, each potential loss is transformed into a real loss, and if it is in the Good state no real losses occur. This model is sufficiently general to allow to vary both the average loss rate as well as the burstiness. Substituting in the expressions of x_G and x_B (Theorem 1), we get

$$x_B = 2\alpha d, \qquad x_G = \alpha d \frac{\bar{b} + \pi_G}{\bar{g}} \tag{6}$$

The throughput is given by,

$$\bar{x} = x_G + \frac{1}{2} x_B + \frac{1}{2} \alpha \frac{d^{(2)}}{d} \tag{7}$$

Remark : *It may seem remarkable that x_B does not depend on the transition probabilities of the Markov chain. This can easily be explained using the following argument. The mean time between losses is clearly $1/R = d/\pi_B$, so the mean increase in the X between two consecutive losses is $\alpha d/\pi_B$. Since we assume that we are in the stationary regime, the mean decrease in X between losses should thus equal to the mean increase. But the mean decrease in X is half its mean value at loss. Thus,*

$$E[X_n | Y_n = B] = 2\alpha d/\pi_B.$$

We conclude that indeed,

$$x_B = E[X_n 1\{Y_n = B\}] = E[X_n | Y_n = B] P(Y_n = B) = 2\alpha d.$$

4.1 The deviation of $\bar{x}$ from $\bar{x}_r$

The non-bursty path that has the same average loss rate is obtained when taking a loss probability p equal to π_B in the two states. The reference throughput in the non-bursty case is then,

$$\bar{x}_r = \frac{2\alpha}{R} - \alpha d + \frac{1}{2} \alpha \frac{d^{(2)}}{d}.$$

Given a a certain average loss rate, we increase the burstiness by increasing b and g in such a way that their ratio remains unchanged. This guarantees that π_B and π_G, and therefore the average loss rate R, remain the same. To study the deviation of the throughput from the non-bursty case, we express $\bar{x}$ as a function of $\bar{x}_r$ and the parameters of the Markov chain. We get,

$$\bar{x} = \bar{x}_r + \alpha d \pi_G \left[\frac{1}{\bar{g}} - \frac{1}{\pi_B} . \right] \tag{8}$$

It is clear from this expression of $\bar{x}$, that the non-bursty case is obtained when $\bar{g} = b = \pi_B$. In our particular case, $\bar{g}$ is the probability that the next potential loss causes a real loss given that we are in the Good state. b is that the probability that it causes a loss given that we are in the Bad state. In the non-bursty case, these two probabilities must be equal. At the same time, they must be equal to π_B, the probability that the next potential loss causes a real loss independently of the current state.

4.2 Effect of the loss model parameters

In this section, we study how the throughput varies as a function of R and the burstiness (via d, b and g). We shall show in particular that for a fixed loss rate R, the throughput increases when the burstiness increases. To facilitate the analysis, we suppose that the time between potential losses is exponentially distributed.

First, we study the effect of an increase in R on the performance. An increase in R can be caused by an increase in the number of potential losses per unit of time ($1/d$) or by an increase in π_B. To study these two cases, we write $\bar{x}$ as,

$$\bar{x} = \alpha d \left[2 + \frac{\bar{b}}{\bar{g}} + \frac{\bar{b}}{\bar{g}(\bar{b} + \bar{g})} . \right] \tag{9}$$

It is clear that when d decreases, the throughput deteriorates. The increase in π_B can be caused by an increase in $\bar{g}$ or a decrease in $\bar{b}$. The two cases result also in throughput deterioration.

Suppose now that d is fixed as well as π_B and π_G. We increase b and g in order to increase the burstiness. The reference throughput remains constant given that it is only a function of the average loss rate. Equation (8) shows well that the average transmission rate improves when losses start to appear in bursts.

4.3 Computation of second moments

In this section, we briefly mention our calculation of the second moments of X_n in the stationary regime. These moments will be shown to have an impact on the average throughput. We define,

$$
\begin{aligned}
x_B^{(2)} &= \lim_{n\to+\infty} E[X_n^2 1\{Y_n = B\}] \\
x_G^{(2)} &= \lim_{n\to+\infty} E[X_n^2 1\{Y_n = G\}] \\
x^{(2)} &= \lim_{n\to+\infty} E[X_n^2] = x_B^{(2)} + x_G^{(2)}
\end{aligned}
$$

The variance of X_n in the stationary regime is no other than,

$$\lim_{n\to+\infty} Var(X_n) = x^{(2)} - x^2 = x_B^{(2)} + x_G^{(2)} - (x_B + x_G)^2 \quad (10)$$

By using, either the relation between the expectations of X_{n+1}^2 and X_n^2 or the Laplace Transform approach, we can prove the following theorem.

Theorem 4: *In the stationary regime,*

$$
\begin{aligned}
x_B^{(2)} &= \frac{4}{3}\left[2\alpha d x_G + \alpha d x_B + \alpha^2 d^{(2)}\right] \\
x_G^{(2)} &= \frac{1}{\bar{g}}\left[\alpha^2 d^{(2)}\left(\frac{1}{3}\bar{b} + \pi_G\right) + 2\left(\frac{1}{3}\bar{b} + g\right)\alpha d x_G + \frac{4}{3}\bar{b}\alpha d x_B\right]
\end{aligned}
$$

We can also prove that,

Theorem 5: *Let $d^{(3)}$ be the third moment of the time between potential losses. The second moment of the transmission rate over a long time interval is equal to,*

$$
\begin{aligned}
\bar{x}^{(2)} &= \lim_{t\to+\infty} \frac{1}{t}\int_0^t X^2(t)dt = x_G^2 + \frac{1}{4}x_B^2 \\
&+ \alpha\frac{d^{(2)}}{d}x_G + \frac{1}{2}\alpha\frac{d^{(2)}}{d}x_B + \frac{1}{3}\alpha^2\frac{d^{(3)}}{d}
\end{aligned}
$$

The effect of the second moments on the throughput can be showed by writing $\bar{x}$ in the following form,

$$\bar{x} = \frac{3}{4}\frac{x_B^{(2)}}{x_B}.$$

As we know, x_B is independent of the parameters of the Markov chain (equation (6)). Thus, the increase in throughput caused by an increase in burstiness can be only the result of an increase in the second moment of the transmission rate upon real losses. Indeed, when losses become clustered, the transmission rate suffers from an important reduction in its value when a burst of losses occurs. The channel enters then in a long Good state where the source has enough time to increase again its rate to an important value. Thus, the variance of the transmission rate increases causing an improvement in performance.

4.4 Impact of transmission rate limitation

Consider the case of TCP flow control. In the absence of losses on the link, the transmission rate increases until reaching a maximum value given by the window advertised by the receiver [17]. Once this window is reached, the transmission rate remains constant until the next loss occurs. Our model does not account for this limitation. It works well when losses are frequent so that the maximum window is rarely reached. We write first some conditions on the loss process to define the region where our previous model works properly. Then, we present a simple approximate calculation to account for this window limitation.

Suppose that the transmission rate is bounded by X_{max}. The point where the transmission rate is most likely to reach the maximal value corresponds to the Good state and is just before the first potential loss in a Bad state. This is the first reduction in the transmission rate after getting out of a Good state. For our previous model to be correct, the expectation of the transmission rate at this point must be much smaller than the upper bound. This condition can be written as,

$$E[X_n | Y_n = B, Y_{n-1} = G] << X_{max}.$$

Taking into account that

$$
\begin{aligned}
&E[X_n | Y_n = B, Y_{n-1} = G] = E[X_n | Y_{n-1} = G] = \\
&E[X_{n-1} | Y_{n-1} = G] + \alpha d = x_G/\pi_G + \alpha d
\end{aligned}
$$

we get the following condition,

$$x_G/\pi_G + \alpha d << X_{max}$$

The larger the average loss rate and the lower the burstiness are, the more likely is that the the above condition holds. For a given loss rate, the increase in burstiness stretches the duration of the Good state and makes it more likely that the transmission rate reaches the upper bound. This impact of burstiness on the correctness of the model cannot be seen if we consider all the points at which the transmission rate is reduced. The expectation in this latter case is equal to, $E[X_n|Y_n = B] = x_B/\pi_B$ and it accounts only for the average loss rate not for the burstiness.

The closer $E[X_n|Y_n = B, Y_{n-1} = G]$ is to X_{max}, the more important is the impact of the receiver window. At the beginning, the receiver window starts to impact the transmission rate evolution only during the Good state of the channel. We can assume that during the Bad state, the transmission rate still have the same evolution as that predicted by our previous model. The receiver window starts to impact the two states once the expectation of the transmission rate just prior to losses in the Bad state becomes larger than X_{max}. This latter condition can be written as,

$$
\begin{aligned}
&E[X_n | Y_n = B, Y_{n-1} = B] >> X_{max} \\
\text{i.e.} \quad &\frac{1}{2}\frac{x_B}{\pi_B} + \alpha d >> X_{max}
\end{aligned}
$$

Once the upper bound starts to impact a state, we make the assumption that during that state the transmission rate always reaches its maximal value. This is the kind of assumption made in [16]. Using the above two conditions, we separate first the space into three regions. In the first region, the transmission rate is not affected by X_{max}. In the second region, the Good state is affected. In the third region, both states are affected. We use then the above assumption to calculate the throughput of the transfer during each state of the channel. Let

$$\bar{x}_G = E[X(t)|Y(t) = G], \qquad \bar{x}_B = E[X(t)|Y(t) = B]$$

where the expectation is w.r.t. the stationary probability. Thus, the throughput is simply equal to,

$$\bar{x} = \pi_G \bar{x}_G + \pi_B \bar{x}_B. \tag{11}$$

Let us define the three regions and calculate $\bar{x}_G$ and $\bar{x}_B$ for each of them:

$\mathbf{E[X_n|Y_n = B, Y_{n-1} = G] < X_{max}}$**:** The transmission rate limitation in this case has no influence and the throughput given by equation (7) can be considered.

$\mathbf{E[X_n|Y_n = B, Y_{n-1} = G] > X_{max}}$ **but** $\mathbf{E[X_n|Y_n = B, Y_{n-1} = B] < X_{max}}$**:** During the Bad state, the transmission rate limitation has no impact and $\bar{x}_B$ can be simply approximated by taking $p = 1$ in equation (5). This is the throughput obtained when the transmission rate is reduced at every potential loss, which is the case for the Bad state. Thus,

$$\bar{x}_B = \alpha d + \alpha \frac{d^{(2)}}{d}.$$

During the Good state, however, another throughput is to be considered. In average, the transmission rate at the beginning of the Good state is equal to,

$$x_0 = E[X_n|Y_n = G, Y_{n-1} = B] = \frac{1}{2}\frac{x_B}{\pi_B} + \alpha d.$$

The average duration of the Good state is $d/\bar{g}$. Using our assumption that the transmission rate during the Good state always reaches X_{max}, we consider that the transmission rate grows first from x_0 to X_{max}, then stays at X_{max} until the beginning of the Bad state. We can find the following expression for $\bar{x}_G$,

$$\begin{aligned}
\bar{x}_G &= \frac{\bar{g}}{d}\left(\int_0^{(X_{max}-x_0)/\alpha} (x_0 + \alpha t)dt + \int_{(X_{max}-x_0)/\alpha}^{d/\bar{g}} X_{max} dt\right) \\
&= \frac{\bar{g}}{d}\left(\frac{X_{max}^2 - X_0^2}{2\alpha} + X_{max}\left(\frac{d}{\bar{g}} - \frac{X_{max} - X_0}{\alpha}\right)\right).
\end{aligned}$$

Given $\bar{x}_G$ and $\bar{x}_B$, the throughput can be calculated using equation (11).

$\mathbf{E[X_n|Y_n = B, Y_{n-1} = B] > X_{max}}$**:** In this case we assume that the transmission rate always reaches X_{max}. The transmission rate just before the occurrence of a real loss can be taken equal to X_{max}, that is, we can now take $x_0 = X_{max}$. Hence,

$$\begin{aligned}
\bar{x}_G &= X_{max} \\
\bar{x}_B &= \frac{1}{d}\left(\int_0^{(X_{max}/2\alpha)} (X_{max/2} + \alpha t)dt + \int_{(X_{max})/2\alpha}^{d} X_{max} dt\right) \\
&= X_{max} - \frac{X_{max}^2}{8\alpha d}.
\end{aligned}$$

And the total throughput is equal to

$$\bar{x} = \pi_G \bar{x}_G + \pi_B \bar{x}_B = X_{max} - \frac{X_{max}^2 \pi_B}{8\alpha d}.$$

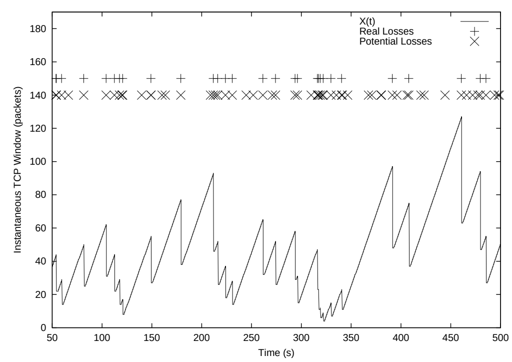

Figure 2: The variation of $X(t)$ vs. time

The difference between our calculation here and the calculation in [16] is that we benefit from the use of a Markov chain, so we can introduce two refined conditions instead of one as in [16]. It was assumed in [16] that only when $E[X_n]$ exceeds X_{max}, the transmission rate limitation starts to impact the throughput.

4.5 Validation of the model

4.5.1 The simulation scenrario

We validate our model using the TCP implementation in the `ns` simulator [15]. Recall that TCP is a window-based flow control protocol that increases its window exponentially during slow start and linearly during congestion avoidance [12]. We consider long TCP transfers to eliminate the impact of the transient behavior at the beginning of the connection. Also, we use the SACK version [7] of TCP since it is able to recover from losses quickly and with a low probability of Timeout and slow start. We suppose that the receiver acknowledges every data packet. This results in a window growth during congestion avoidance of approximately one packet every RTT [17]. We suppose also that the receiver window is very large so that it does not affect the transmission rate. Later, we will show that the estimations derived in Section 4.4 well agree with simulation results obtained in the case when the window evolution is limited by the receiver window. We consider the TCP window size in packets as the transmission rate in our mathematical model since this window varies linearly as a function of time and is divided by two upon loss detection. The different rates in our model are then expressed in terms of packets and need to be divided by the RTT in order to get the real rates.

The simulation scenario consists of a TCP connection crossing a 2Mbps link. The RTT of the connection is taken equal to 560 ms. TCP packets are of total size 1000 Bytes. We add our loss model to the simulator and we associate it to the 2Mbps link. We account only for losses on the link and we study their impact on the throughput. We chose the parameters of the simulation in a way to not get losses in the other parts of the network. This clearly requires that losses are frequent so that the buffers in network routers do not overflow. The first condition of Section 4.4 is always satisfied with X_{max} equal to the network capacity. The purpose of the present paper is to create and to validate the model which takes into account the burstiness of losses. In a subsequent work, we will show how the parameters of our model can be inferred from a real TCP trace.

The time between potential losses is taken to be exponentially distributed. As we will explain later, we simulate a potential loss that it is to be transformed to a real loss by dropping all packets leaving the lossy link during a small time interval (a 100 ms is chosen in the following simulations). Loss events in our model correspond then to bursts of packet losses on the lossy link. The reason for dropping packets in bursts is that TCP traffic is not actually fluid but rather bursty. Thus, it is very likely that at the instant of loss, there is no packets leaving the link. By dropping packets in bursts, we guarantee that some packets are at least lost when our model tells us that the throughput of the connection has to be reduced.

Figure 2 shows a typical variation of the congestion window of the TCP connection. We see well how potential losses are transformed into real losses and how real losses cause a reduction of the window by a factor of two. In what follows, we run the connection for one hour and then we calculate the values of x_G, x_B and $\bar{x}$. These simulation results are then compared to those given by our analysis. When simulating, x_G (resp. x_B) is calculated by summing the window sizes when a potential loss occurs and the link in the Good state (resp. in the Bad state), then by dividing this sum by the total number of potential losses. $\bar{x}$ is calculated as the throughput of the connection over one hour expressed in terms of Packets/s times the RTT. This gives the time average of the congestion window. First, we fix the parameters of the Markov chain of the link and we vary the time between potential losses. This allows to check the impact of the average loss rate on the throughput for a given burstiness. Afterwards, we shall vary the burstiness while fixing the average rate of losses.

4.5.2 *Impact of the average loss rate*

In our first set of simulations, d is varied between 1 and 10 seconds. b and g are however taken equal to 0.6. Our analysis predicts a linear variation of the three quantities x_G, x_B and $\bar{x}$ (equations (6) and (9)). Figures 3, 4 and 5 show well the match between simulation and analytical results.

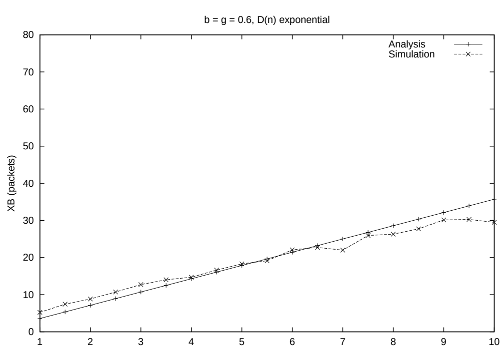

Figure 3: The variation of x_B vs. d

We shall give some more details about the way the losses are generated and then explain the small deviations from the analytical results. We see that the slope of the line given by simulation is slightly smaller than the one given by

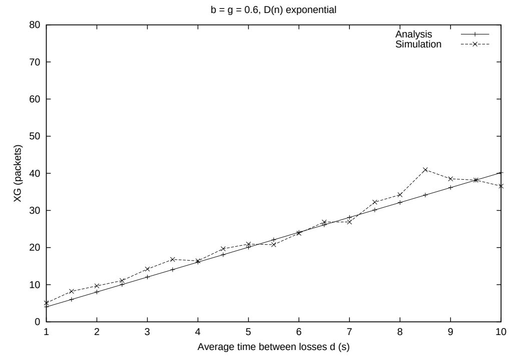

Figure 4: The variation of x_G vs. d

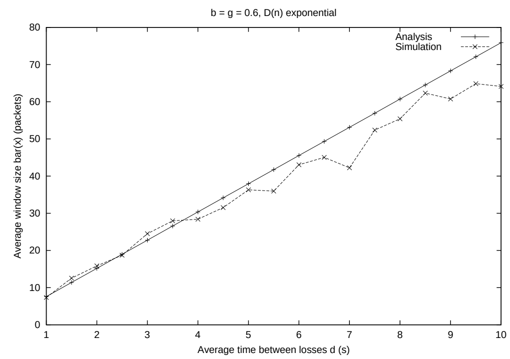

Figure 5: The variation of $\bar{x}$ vs. d

analysis. The simulated model consists of individual packets that are sent in bursts on the link. The lossy link may not be carrying TCP packets when a potential loss has to be transformed into a real loss. At small d, losses are frequent and the window is most of the time of small size. When the window is very small and an event of loss is simulated, there might not be an actual packet on the link to which this loss corresponds. This results in many real losses considered by the analytical model but not considered by the simulation. Now, when d increases, the window becomes larger and the probability that the link is not carrying TCP packets when a potential loss occurs becomes smaller. Thus, the simulation line becomes closer to the analytical line.

To overcome the above problem, we simulate a loss as an event that causes the loss of all the packets that cross the lossy link during a certain time interval. By taking a large time interval to represent potential losses, we solve the problem of small windows. However, large windows see a large number of lost packets which causes sometimes a Timeout and a slow start. For this reason, we see that the simulation results fall below the analytical ones at large d.

4.5.3 *The impact of burstiness*

We fix here the average time between potential losses to 5 seconds and we change the transition probabilities b and g while keeping $b = g$. This results in $\pi_G = \pi_B = 0.5$ which guarantees that the average loss rate remains constant. Our analysis shows that x_B must not change (equation (6)). x_G and $\bar{x}$ must however increase as a result of the increase in burstiness (equations (6) and (9)). Figures 6, 7 and 8 val-

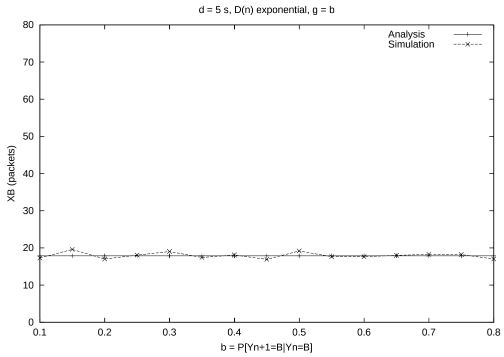

Figure 6: The variation of x_B vs. b

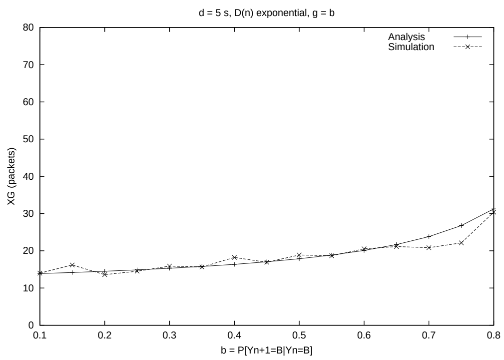

Figure 7: The variation of x_G vs. b

idate our analytical results. In particular, it is clear from Figure 8 that by increasing b from 0.1 to 0.8, the average throughput increases by around 60% even though the average loss rate remains unchanged. This confirms our result concerning the improvement in performance when losses become clustered.

We plot finally in Figure 9 the variance of the transmission rate upon potential loss occurrence (the window size in case of TCP). This variance is given in equation (10). As predicted by our analysis, the simulations show the increase in the variation of X_n when burstiness increases. On a bursty path, the source transmission rate varies between important values when the path is in the Good state and small values when it is in the Bad state.

4.5.4 Case of a limitation on the transmission rate

We now consider a case where the receiver window is set to a finite value so that it limits the evolution of the congestion window. We set b and g to 0.6 and we take an exponential time between potential losses of average 5 s. We reduce the RTT of the connection to 250 ms and we set the receiver window to the bandwidth delay product. We change d from 1 to 10. By simple calculation, we see that in this setting, we cross the three regions we defined in Section 4.4 while introducing the limitation on the transmission rate for our model. Figure 10 shows how our approximation correctly estimates the real throughput. The model without limitation on the transmission rate leadsto a clear overestimation of the real throughput in this scenario.

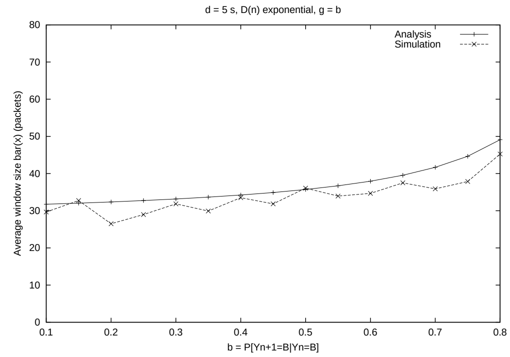

Figure 8: The variation of $\bar{x}$ vs. b

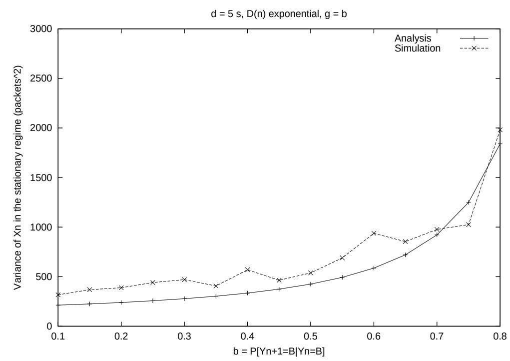

Figure 9: The variation of $Var(X_n)$ vs. b

5. CONCLUSIONS

In this paper, we studied the performance of a TCP-like flow control protocol as a function of losses. In addition to the average loss rate considered in the previous works, we evaluate the impact on the performance of burstiness in the loss process. We define a model for losses using potential losses and a two-state Markov chain to account for burstiness. We then calculate the throughput and the moments of the transmission rate at some potential loss instants. The throughput is compared to the one achieved when operating over a non-bursty path having the same average loss rate. Our main result is that for a given loss rate, the performance improves when losses tend to appear in bursts. This increase in performance with burstiness is caused by an increase in the second moment of the transmission rate. We conduct

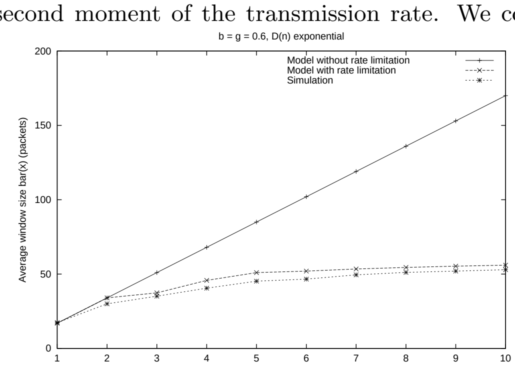

Figure 10: $\bar{x}$ vs. d with a limitation on the rate

a set of simulations with `ns` to validate the analytical results. A good match between simulation and analysis has been noticed.

Another result of our analysis is a better understanding of how the throughput has to be calculated. In the literature, the source transmission rate (or the congestion window in case of TCP) has been shown to vary in the stationary regime between two fixed values X and $2X$ with a time average transmission rate equal to $3X/2$. We showed in our analysis that this is correct only when the time between losses is constant which is not the case given the random nature of the loss process. This transformation of the random problem to a a deterministic one is the result of the assumptions made to simplify the analysis. In general, the second moment of the inter-loss time has to be considered for an accurate calculation of the throughput. In the particular case of an exponentially distributed inter-loss time, we found that the average transmission has to be taken equal to $2X$ rather than $3X/2$.

6. APPENDIX: CONVERGENCE OF X_N TO A STATIONARY REGIME

We rewrite the dynamics (1) of X_n as,

$$X_{n+1} = A_n X_n + B_n \tag{12}$$

where,

$$\begin{aligned} A_n &= (1 - U_n/2)1\{Y_n = G\} + (1 - V_n/2)1\{Y_n = B\} \\ B_n &= \alpha D_n \end{aligned}$$

For any initial condition X_0, we obtain by iterating (12),

$$X_n = \sum_{j=0}^{n-1} \left(\prod_{i=n-j}^{n-1} A_i \right) B_{n-j-1} + \left(\prod_{i=0}^{n-1} A_i \right) X_0$$

for all $n \geq 0$. If we assume that Y_n is initially in steady state then (A_n, B_n) are jointly stationary. We denote by (A_n^*, B_n^*) this stationary process. We show in this section that the process X_n converges to a stationary solution of (12), i.e. to a process X_n^* satisfying $X_{n+1}^* = A_n^* X_n^* + B_n^*$ for all $n \geq 0$.

Theorem 6: *Assume that Y_n contains a single recurrent class and is initially in steady state. Consider an arbitrary initial state X_0. Then,*

$$X_n^* = \sum_{j=0}^{+\infty} \left(\prod_{i=n-j}^{n-1} A_i \right) B_{n-j-1} \tag{13}$$

is the only stationary solution of (12) and is ergodic. The sum on the right hand side of (13) converges absolutely almost surely. Furthermore, $|X_n - X_n^*| \to 0$ a.s. for all X_0 on the same probability space as $\{(A_n, B_n)\}$. In particular, the distribution of X_n converges to that of X_n^* as $n \to +\infty$.

Proof: We use Theorem 2A in [10] (based on [3, 18]). The assertion follows directly if we establish the following conditions of the Theorem (i) $-\infty \leq E[\log|A_0|] < 0$ (ii) $E[\log|B_0|^+] < \infty$. We show that these conditions indeed hold.

The only possible values of A_0 are $1/2$ and 1. Thus the only possible values of $\log|A_0|$ are $\log 0.5$ or 0. Under the assumptions of our model, the value $\log 0.5 < 0$ has positive probability. This implies conditions (i). By Jensen's inequality we have $e^{E[\log|B_0|]} \leq E|B_0|$ which is finite. This implies condition (ii). Ergodicty follows from [2], p. 14. $\diamond$

Remark: *The conclusions of the above theorem can be extended to the case that the Markov chain Y_n is initially not in its steady state distribution. This is due to the fact that coupling of Y_n to a stationary regime occurs in a time which is a.s. finite (since the Markov chain Y_n contains a single ergodic class).*

7. REFERENCES

[1] F. Anjum and L. Tassiulas, "On the Behavior of Different TCP Algorithms over a Wireless Channel with Correlated Packet Losses", *ACM SIGMETRICS*, Mar 1999.

[2] A.A. Borovkov, "Stochastic Processes in Queuing Theory", Springer-Verlag.

[3] A. Brandt, "The stochastic equation $Y_{n+1} = A_n Y_n + B_n$ with stationary coefficients", *Advances in Applied Probability*, Vol. 18, 1986.

[4] H. Chaskar, T. V. Lakshman, and U. Madhow, "On the design of interfaces for TCP/IP over wireless", *IEEE MILCOM*, 1996.

[5] A. Chockalingam, M. Zorzi, and R.R. Rao, "Performance of TCP on Wireless Fading Links with Memory", *IEEE ICC*, Jun 1998.

[6] E.N. Gilbert, "Capacity of a burst-noise channel", *Bell Systems Technical Journal*, Sep. 1960.

[7] K. Fall and S. Floyd, "Simulation-based Comparisons of Tahoe, Reno, and SACK TCP", *ACM Computer Communication Review*, Jul 1996.

[8] S. Floyd, "Connections with Multiple Congested Gateways in Packet-Switched Networks Part 1: One-way Traffic", *ACM Computer Communication Review*, Oct 1991.

[9] S. Floyd and K. Fall, "Promoting the Use of End-To-End Congestion Control in the Internet", *IEEE/ACM Transactions in Networking*, Aug 1999.

[10] P. Glasserman and D. D. Yao, "Stochastic vector difference equations with stationary coefficients", *Journal of Applied Probability*, Vol. 32, 1995.

[11] G. Koole and S. Asmussen, "Marked point processes as limits of Markovian arrival streams", *Journal of Applied Probability*, 1993.

[12] V. Jacobson, "Congestion avoidance and control", *ACM SIGCOMM*, Aug 1988.

[13] A. Kumar, "Comparative Performance Analysis of Versions of TCP in a Local Network with a Lossy Link", *IEEE/ACM Transactions on Networking*, Aug 1998.

[14] T.V. Lakshman and U. Madhow, "The performance of TCP/IP for networks with high bandwidth-delay products and random loss", *IEEE/ACM Transactions on Networking*, Jun 1997.

[15] The LBNL Network Simulator, *ns*, http://www-nrg.ee.lbl.gov/ns.

[16] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP Throughput: a Simple Model and its Empirical Validation", *ACM SIGCOMM*, Sep 1998.

[17] W. Stevens, "TCP Slow-Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms", *RFC 2001*, Jan 1997.

[18] W. Vervaat, "On a stochastic difference equation and a representation of non-negative infinitely divisible random variables", *Advances in Applied Probability*, Vol. 11, 1979.