



TCP in Presence of Bursty Losses

Eitan Altman, Konstantin Avrachenkov, Chadi Barakat

► To cite this version:

Eitan Altman, Konstantin Avrachenkov, Chadi Barakat. TCP in Presence of Bursty Losses. RR-3740, INRIA. 1999. inria-00072924

HAL Id: inria-00072924

<https://inria.hal.science/inria-00072924>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

TCP in presence of bursty losses

Eitan Altman — Kostya Avrachenkov — Chadi Barakat

N° 3740

July 1999

————— THÈME 1 —————



***rapport
de recherche***

TCP in presence of bursty losses

Eitan Altman , Kostya Avrachenkov , Chadi Barakat

Thème 1 — Réseaux et systèmes

Projet Mistral

Rapport de recherche n° 3740 — July 1999 — 30 pages

Abstract: We analyze in this paper the performance of TCP/IP-like congestion control in the presence of random losses. The input rate in the control scheme that we consider has a linear growth rate; whenever a loss occurs, the input rate is halved. This approximates the performance of several versions of TCP/IP that divide their congestion window by two whatever is the cause of the loss. We propose a mathematical model that allows to account for burstiness as well as for correlation in the loss process. Our aim is to study the impact of burstiness on the throughput of the connection. We compute the expected instantaneous input rate and its moments at some potential loss instants, and provide for a useful implicit expression for the Laplace Stieltjis Transform. This allows us to compute explicitly the average throughput and its moments. We show that the average throughput is indeed sensitive to the distribution of the loss process, and not just to the average loss rate: for a given average loss rate, we show that the throughput *increases* with the burstiness of the loss process. We finally examine the impact of burstiness of losses on the throughput variability.

Key-words: TCP, Congestion control, Random losses, Congestion losses, Markov chain, Performance evaluation, Simulations.

TCP en présence de pertes en rafale

Résumé : TCP utilise la perte de paquets pour détecter la congestion dans le réseau. Il divise sa fenêtre de congestion en deux lorsqu'une perte est détectée à la source. Ceci résulte en une dégradation de la performance du protocole lorsque les paquets TCP sont perdus sur le chemin pour d'autres raisons que la congestion (ex. erreurs de transmission). Plusieurs travaux ont étudié l'effet de ces pertes indésirables sur la performance. Ces travaux ont supposé que les paquets TCP se perdent avec la même probabilité et donc, la performance a été étudiée seulement en fonction du taux moyen des pertes. Dans ce travail, on développe un modèle mathématique qui nous permet d'étudier, non seulement l'effet du taux moyen des pertes, mais aussi l'effet de leur distribution. Particulièrement, on étudie l'effet de l'apparition en rafale des pertes sur la performance du protocole. Nous trouvons les expressions explicites du débit et de ses moments. Ces expressions nous montrent comment varie le débit en fonction des paramètres du processus de pertes. Notre résultat principal est que le débit d'une connexion TCP augmente lorsque les pertes tendent à apparaître en rafale même si le taux moyen des pertes par unité de temps reste le même. À la fin, nous validons nos résultats mathématiques par des simulations conduites avec le simulateur de réseau ns.

Mots-clés : TCP, Contrôle de congestion, Pertes aléatoires, Pertes de congestion, Chaîne de Markov, Évaluation de performance, Simulations.

1 Introduction

Communication channels (such as wireless and satellite channels) often suffer from transmission errors that result in corruption of information. Error detection techniques are frequently used to discover erroneous packets which enables the network or the receiver to discard corrupted packets. Corruption of information and packet losses due to noisy links occur even in the presence of a large amount of redundancy added for error correction (this is the case on satellite links). Characterization and experimental measures on lossy links can be found in [1, 14, 12]. In [1] the measured loss process of packets is characterized not only through its average rate but also through the average number of losses that occur in each "loss event", i.e. in each burst of losses. Measurements of the loss process with different types of transmission all show that losses occur in burst. For example, when the Intelsat IDR type digital transmission is used (this transmission contains a forward error correction of the ratio of 3/4 and additional scrambling), then the average number of errors per burst is around 10 for any average bit error rates between 10^{-4} to 10^{-11} . When the Intelsat IBS type digital transmission is used (it contains a forward error correction of the ratio of 1/2 and additional scrambling), then the average number of errors per burst is around 5.

Congestion control mechanisms in the Internet are sensitive to losses; these are interpreted as congestion signals and result in actions aimed at reducing congestion. When the rate of random losses is high then the throughput of these mechanism may decrease dramatically. Several previous works have analyzed the impact of a random loss process on a TCP/IP connection [2, 10, 11, 15]. But all these works did not consider the burstiness of the loss process. Either times between losses were replaced by their expectations, or independent losses were assumed.

We propose a mathematical model that allows to account for burstiness as well as for correlation in the loss process. Our aim is to study the impact of burstiness on the throughput of the connection. We compute the expected instantaneous input rate and its Laplace Stieltjes Transform at some potential loss instants, which allows us to compute the average throughput. We show that this average throughput is indeed sensitive to the distribution of the loss process, and not just to the average loss rate: for a given average loss rate, we show that the throughput *increases* with the burstiness of the loss process. We finally examine the impact of burstiness of losses on the throughput variability. We assume that input rate in the control scheme

that we consider has a linear growth rate; whenever a loss occurs the input rate is halved. This approximates the performance of several versions of TCP/IP [16]. Our mathematical model is then compared to simulations of the SACK version of TCP [5] performed with `ns` simulator [13], and good match is reported between the model and the simulations.

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 input rate in presence of losses. At the end of this section, we give the general expression of the average throughput. The average 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 performance. A particular case is considered for this study. 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

We consider a fluid model that might be used to approximate a flow control mechanism based on a discrete number of packets. The input rate is assumed to grow linearly at a rate of α . In the case of a window-based flow control as TCP [16], the input rate at any instant is equal to the window size divided by the Round Trip Time of the connection. Our fluid model thus uses a window size which may take non integer values. The growth continues till a loss occurs. As in many versions of the TCP flow control, we suppose that upon the detection of the loss the input transmission rate is halved [16].

We ignore in this work losses caused by a real congestion. Transmission errors are frequent so that they reduce the input rate before the occurrence of congestion in the network. At the end of this paper, we quantify this condition. Given a certain path with a certain bandwidth, we show where our analysis can be applied.

We wish the model for the lossy channel to allow to model both burstiness in the losses as well as correlations in the time between consecutive losses. The Gilbert model is often used in this context [4] in which the channel is assumed to have two

states: *Good* and *Bad*; losses are then assumed to occur in the bad state. In the Gilbert model however, 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 channel states. Our model is related to the MAP process [8].

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 occurring with a certain distribution. Let T_n denote the time at which the n th potential loss may occur. Let $D_n, n = 1, 2, \dots$ 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 Stieltjes Transform $D^*(s)$. Let X_n be the input rate just prior to the instant of the n th potential loss.

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 n th 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. The time between potential losses is therefore independent of the state of the channel.

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 steady-state probabilities:

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

The average loss rate over the channel is given by

$$R = \frac{(p_G \pi_G + p_B \pi_B)}{d},$$




Figure 1: The Markov chain associated to the channel

which is equal to the average number of times the source reduces its rate per unit of time. The unit of time is the one in which d is expressed. This model allows us to vary the average loss rate as well as the pattern of the loss process for a given average loss rate. The time between potential losses and the parameters of the Markov chain are used for this objective. In the following section, we present the analysis of the performance in presence of this 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 possible states of the channel. A value one of these variables means that the potential loss causes really a packet loss and thus a reduction in X_n . A value zero however means that the potential loss is not transformed into a real loss and therefore, X_n is not affected. We have then,

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

The evolution of the input rate is the following:

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

We begin by computing the first moments of X_n in steady-state. The calculation of these moments helps us in calculating the average input rate (or the average throughput of the connection) over a long time interval. We denote

$$x_G = \lim_{n \rightarrow \infty} E[X_n 1\{Y_n = G\}], \quad x_B = \lim_{n \rightarrow \infty} E[X_n 1\{Y_n = B\}].$$

In steady-state, the average value of the input rate prior to the occurrence of the potential loss is no other than

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

We show first that these moments converge to steady-state values. Then, we follow two approaches in calculating these moments. The first approach consists of:

- (a) using the relation between X_{n+1} and X_n given in equation (1).
- (b) moving n to infinity.

The second approach uses the Laplace Transforms of X_n in steady state.

Theorem 1 *The first moments of the input 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,

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

Proof: First we write,

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

$$\begin{aligned}
& E[X_{n+1}1\{Y_{n+1} = B\}] \\
&= \bar{g}E\left[1 - \frac{U_n}{2}\right] E[X_n1\{Y_n = G\}] \\
&+ bE\left[1 - \frac{V_n}{2}\right] E[X_n1\{Y_n = B\}] + \alpha dP(Y_{n+1} = B) \\
&= \bar{g}\gamma_G E[X_n1\{Y_n = G\}] + b\gamma_B E[X_n1\{Y_n = B\}] + \alpha dP(Y_{n+1} = B). \quad (3)
\end{aligned}$$

Then, we show that the sequences $E[X_n1\{Y_n = G\}]$, $n = 0, 1, \dots$ and $E[X_n1\{Y_n = B\}]$, $n = 0, 1, \dots$ 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,

$$\mathcal{X}_n = \begin{bmatrix} E[X_n1\{Y_n = G\}] \\ E[X_n1\{Y_n = B\}] \\ P(Y_n = G) \\ P(Y_n = B) \end{bmatrix}.$$

The convergence properties of the augmented sequence $\mathcal{X}_n$, $n = 0, 1, \dots$ 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

$$\begin{aligned} 0 < b < 1, \quad 0 < g < 1, \\ 1/2 \leq \gamma_G \leq 1, \quad 1/2 \leq \gamma_B < 1. \end{aligned}$$

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

$$\lambda_2 < 1,$$

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

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

Since $g < 1$ and $b < 1$, the left 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, then 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 $E[X_n 1\{Y_n = G\}]$, $n = 0, 1, \dots$ and $E[X_n 1\{Y_n = B\}]$, $n = 0, 1, \dots$ by letting $n \rightarrow \infty$ in the original equations (2) and (3). By moving n to infinity, we get the following system of equations:

$$\begin{aligned} x_G &= \gamma_G g x_G + \gamma_B \bar{b} x_B + \alpha d \pi_G, \\ x_B &= \gamma_B \bar{g} x_G + \gamma_B b x_B + \alpha d \pi_B. \end{aligned}$$

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

Remark 1 *We show in the Appendix that in fact X_n converges to a unique stationary regime.*

Laplace Stieltjis Transforms and the moments of X_n

Define the following Laplace Stieltjis Transforms:

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

where we assume that X_n is in steady-state.

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

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

Proof: We write,

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

Using the fact that

$$E \left[e^{-s(1-U_n/2)X_n} 1\{Y_n = G\} \right]$$

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

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 steady state) 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 relations:

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

More general, the order k moments can be obtained in a similar way from the implicit relations for the Laplace Transforms $Z(s, G)$ and $Z(s, B)$ using

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

The average throughput

Theorem 3 *The time average input denoted $\bar{x}$, can be expressed as*

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

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

Proof: We have almost surely

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

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 can be established from the convergence result in the Appendix. This concludes the proof. $\diamond$

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 obtained in the bursty case is compared to the throughput achieved when the channel is subjected to a non-bursty loss process having the same average loss rate. We denote this throughput $\bar{x}_r$ and we use it as a reference to evaluate the effect of burstiness. A non-bursty loss process is obtained when the channel has 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, the non-bursty channel must have a p equal to:

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

Lemma 1 *On a non-bursty channel, the source achieves an average input rate of:*

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

Proof: The given expression for $\bar{x}_r$ can be easily obtained when substituting in the general expression of the average input rate $\bar{x}$ (theorem (3)), γ_G and γ_B by their values as a function of p , the loss probability in the two states. On a non-bursty channel, we have,

$$\gamma = \gamma_G = \gamma_B = 1 - p/2.$$

As one must predict, by writing

$$\bar{x}_r = \gamma(x_G + x_B) + \frac{1}{2} \alpha \frac{d^{(2)}}{d},$$

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. $\diamond$

Comparison with previous works

Consider a particular case where the non-bursty channel has $p = 1$. In this case, all potential losses cause a reduction in the throughput. This forms a loss process similar (even more general) to the one used in many previous works [6, 11, 15]. These works suppose that in steady state, the window (or the input rate) varies linearly between two values w and $2w$. They calculate the average window size over a long time interval as being equal to $3w/2$. The average time (or the average number of packets) between two losses is used to calculate w . It is taken equal to the increase in the window size during this average time. Given the average window size, they divide it by the Round Trip Time of the connection in order to get the average throughput as a function of the average time between losses or the average loss rate.

Now, our model shows well that in the presence of a non-bursty channel with $p = 1$, the average input rate just prior to a loss is equal to $2\alpha d$. The average window

size just after a real loss is simply αd . Thus, αd in our model corresponds to their w . However, our model doesn't give the average window size $3w/2$ they find. In addition to their result, 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 loss distribution of value d . Although they use a Bernoulli loss model, which corresponds to an exponential loss distribution in the continuous time space, these works transform the loss distribution into a deterministic one which results in the disappearance of the term $d^{(2)}$ in their analysis. Our model however, in addition to the consideration of the burstiness, considers the exact expression of the average throughput over a long time interval. In addition to the average time between losses, the second moment of this time must be considered otherwise the throughput will be underestimated. In the case of an exponential loss distribution, $d^{(2)}$ is equal to $2d^2$ and the average window size (or the average input rate) is equal simply to the average window size just prior to losses, hence to $2w$.

4 A case study

In the sequel we consider the special case where

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

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 its burstiness.

Substituting in the expressions of x_G and x_B (theorem 1), we get:

$$x_B = 2\alpha d, \quad x_G = \alpha d(\bar{b} + \pi_G)/\bar{g}. \quad (5)$$

The average input rate is given by

$$\bar{x} = x_G + \frac{1}{2}x_B + \frac{1}{2}\alpha \frac{d^{(2)}}{d}. \quad (6)$$

Remark 2 *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*

$$\frac{1}{R} = \frac{d}{\pi_B},$$

so the mean increase in the X between two consecutive losses is $\alpha d/\pi_B$. Since we assume that we are in steady-state, 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,

$$E[X_n | Y_n = B] = \frac{2\alpha d}{\pi_B}.$$

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 channel that has the same average loss rate is obtained when taking a loss probability p equal to π_B . Recall that p is the probability that a potential loss causes a real loss regardless of the channel state. Choosing the same loss probability in the two states results in a non-bursty loss process. The value of p however is chosen so that we get the same average loss rate as a bursty channel. 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 certain distribution of potential losses and 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

$$\bar{x} = \bar{x}_r + \alpha d \pi_G \left[\frac{1}{\bar{g}} - \frac{1}{\pi_B} \right]. \quad (7)$$

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. On a non-bursty channel, these two probabilities must be equal. At the same time, they must be equal to the average probability that a potential loss occurs while the channel is in the bad state π_B and this independently of the current state.

4.2 Effect of the loss model parameters on the performance

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 average throughput increases when the burstiness increases. To facilitate the analysis, we suppose that the time between potential losses is exponentially distributed which means that $d^{(2)}$ is equal to $2d^2$. This is similar to the geometrical distribution used in the literature [11, 15] to model the number of packets correctly received between real losses.

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

$$\bar{x} = \alpha d \left[2 + \frac{\bar{b}}{\bar{g}} + \frac{\bar{b}}{\bar{g}(\bar{b} + \bar{g})} \right]. \quad (8)$$

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.

Now, we suppose that d is fixed as well as π_B and π_G . We increase b and g in order to increase the burstiness. Note that the change in the Markov chain parameters doesn't affect the reference throughput, which is equal to $2\alpha/R$ in the exponential case. $\bar{x}_r$ is only a function of the average loss rate. However, it is clear from equation (7) that the average input rate improves when losses appear in bursts.

Theoretically, equation (7) shows that the throughput increases indefinitely when $\bar{g}$ moves to zero. This is because we didn't limit the input rate when the channel

is in the good state. In reality, the throughput cannot move to infinity and it is limited by the available resources in the network. Even if the channel is always in the good state, congestion losses will occur forcing the source to reduce its rate. We can consider as if there is always a minimum loss probability prohibiting the input rate from reaching higher values than the available bandwidth.

4.3 Computation of second moments

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

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

It is clear that $x^{(2)} = x_B^{(2)} + x_G^{(2)}$. The variance of X_n in steady state is no other than

$$\text{Var}(X_n) = x^{(2)} - x^2 = x_B^{(2)} + x_G^{(2)} - (x_B + x_G)^2. \quad (9)$$

By using, either the relation between the expectations of X_{n+1}^2 and X_n^2 as in equations (2) and (3), or the Laplace Transform approach, we can prove the following theorem

Theorem 4 *In steady state,*

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

As in the case of the average throughput over a long time interval (theorem 3), 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 input rate over a long time interval is equal to:*

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

The effect of the variance of X_n on the average throughput can be showed if we write $\bar{x}$ as a function of the first and second moments of X_n . It is easy to show that in our particular case, the average throughput can be written as

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

By substituting $x_B^{(2)}$ by its value in this new expression of $\bar{x}$, we get the previous expression (equation (6)). As we know, x_B is independent of the parameters of the Markov chain (equation (5)). 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 input rate upon real losses. Indeed, when losses become clustered, the input rate suffers from an important reduction in its value when a burst of losses occurs. Then, the channel enters a long good state where the source has enough time to increase its rate again to an important value. Thus, the variance of the input rate increases causing an improvement in performance. The expression for $x_B^{(2)}$ we gave shows this increase as a function of burstiness.

4.4 Congestion versus random losses

In the absence of the random losses, the input rate increases until reaching the maximum available bandwidth in the network. Beyond this point, queues start to build up in network buffers. Some time later, these buffers overflow leading to loss of packets. Here, the network is said to be congested. When detecting these losses, the source reduces its input rate by one half to alleviate the congestion in the network.

In our work, we study the impact of undesirable random losses on the performance. We assumed that random losses are so frequent so that the input rate is reduced

before the occurrence of congestion in the network. It is clear that if random losses are not so frequent, the input rate will increase until a congestion occurs. As a consequence, the source reduces its performance which is not predicted by our model.

In this section, we explain how much transmission errors must be frequent so as to be able to neglect the effect of congestion losses and, then to apply our results. For simplicity, our particular case is considered in the section.

Suppose that the maximum possible input rate is X_{max} . A congestion collapse occurs when the input rate exceeds this maximum rate. In the case of a window-based protocol, X_{max} denotes the maximum window size, which is equal to the maximum number of packets that can be in flight between the source and the destination (the pipe size). Congestion losses are avoided if the input rate $X(t)$ is reduced due to a transmission error before it reaches X_{max} . This requires a certain minimum transmission error rate. Also, given a certain average loss rate, this requires a certain maximum burstiness. As we will see later, the increase in burstiness increases the length of the good state and therefore lets the input rate reach higher values during this state. It is not enough to consider the average loss rate in the calculation of the required condition to neglect congestion losses. The burstiness of the loss process must also be considered.

The maximum value of $X(t)$ is seen just before the first potential loss in a bad state. This is the first reduction in the input rate after getting out of a good state. On average, the expectation of the input rate at this point must be much smaller than the maximum input rate. In steady state, this condition can be written as

$$\begin{aligned} E[X_n | Y_n = B, Y_{n-1} = G] &<< X_{max}, \\ \text{i.e.} \quad x_G + x_B + \alpha d &<< X_{max}. \end{aligned} \tag{10}$$

The larger the average loss rate and the lower the burstiness are, the more likely it is that the above condition holds. Even if the average loss rate remains the same (via the same d , π_B and π_G), the increase in burstiness stretches the duration in which the channel is in the good state and increases then the input rate X_n that appears in (10). This behavior cannot be seen if we consider all the points at which the input rate is reduced. The expectation in this later case is equal to

$$E[X_n | Y_n = B] = \frac{x_B}{\pi_B},$$

and it accounts for the average loss rate not for the burstiness.

The farther $E[X_n|Y_n = B, Y_{n-1} = G]$ increases beyond the maximum input rate, the more important congestion losses become. The evolution of the input rate is more and more controlled by the available bandwidth in the network and the buffer sizes rather than by random losses. This is the assumption a congestion control protocol as TCP/IP [9] makes about the network. The best performance is obtained when random losses are sufficiently rare. The condition for random losses to be negligible is that the average input rate during the bad state of the channel given by our model (i.e. when ignoring congestion losses) is much larger than the maximum input rate in the network (in absence of random losses). This can be written as

$$E[X_n|Y_n = B, Y_{n-1} = B] \gg W_{max},$$

$$\text{i.e.} \quad (x_G + x_B)/2 + \alpha d \gg W_{max}.$$

For a random loss process satisfying the later condition, congestion losses are dominant, regardless of the channel state. In contrast to the first region where random losses are frequent, the performance in this region is mainly a function of network parameters (bandwidth and buffer sizes) rather than the loss process parameters.

4.5 Simulations

To validate our analysis, we simulate a flow control protocol that reduces its rate by half upon loss detection, and that increases its rate linearly as a function of time. The protocol used is TCP and the simulator is `ns`, the Network Simulator [13]. TCP uses a sliding window to control the flow of packets in the network. The window is increased by one packet for every Acknowledgment (ACK) during slow start and by one packet for every window's worth of ACKs during congestion avoidance. The window growth during congestion avoidance can be considered as linear if queuing time in network buffers is negligible with respect to the propagation delay. To run the protocol in this linear phase, we consider frequent losses so that the window is reduced in general before the building of queues in network buffers. Also, we use the SACK version of TCP that manages most of the time to recover from losses without resorting to Timeout and slow start [5]. Note that when a Timeout occurs, the slow start threshold is set to half the current window, the window is reset to the value of 1 packet instead of halving it and a slow start phase is called. The increase




Figure 2: The variation of $X(t)$ as a function of time

in the window during that phase is exponential instead of linear. Once the slow start threshold is reached, slow start ends and the linear increase in the window is resumed.

We suppose that the receiver acknowledges every data packet. This results in a window growth of one packet every Round Trip Time (RTT). In practice there are implementations of TCP in which only every second RTT the window grows by one packet; this happens when only each second received packet generates an ACK. Such ACKs are known as “Delayed ACKs” [16].

When plotting our results, we show the window size in packets instead of TCP throughput. The throughput of a TCP connection at any instant is simply the window size divided by the RTT.

The simulation scenario consists of a TCP connection crossing a 2Mbps lossy link. TCP packets are of total size 1000 Bytes. We add our loss model to the simulator. Time between potential losses is considered to be exponential. Figure 2 shows a typical variation of the congestion window of the TCP connection when operating over a lossy channel. 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






Figure 3: The variation of x_B as a function of d

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 channel 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 the one hour expressed in terms of Packets/s times the RTT.

First, we fix the parameters of the Markov chain of the channel 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. In both cases, we choose the parameters for the lossy channel so that congestion losses can be ignored.

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 (5) and (8)). Figures 3, 4 and 5 show well the match between simulation and analytical results.

Figure 4: The variation of x_G as a function of d Figure 5: The variation of $\bar{x}$ as a function of 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 analysis. The simulated model consists of individual packets that are sent in bursts on the channel rather than a fluid, as in our approximating model. In the simulated scenario, the lossy channel 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 channel 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 channel 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 .

Next we consider the impact of burstiness on the average throughput. We thus keep the average rate of losses constant. We fix the average time between potential losses to 5 seconds and we change the transition probabilities b and g but keeping $b = g$. This results in $\pi_G = \pi_B = 0.5$. Our analysis shows that x_B must not change (equation (5)). x_G and $\bar{x}$ however must increase as a result of the increase in burstiness (equations (5) and (8)). The simulation validates our analytic results, as can be seen in figures 6, 7 and 8. In particular, it is seen 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 more and more clustered.

Finally, we plot in Figure 9 the variance of the input rate upon potential loss occurrence (the window size in case of TCP). This variance is given in equation (9). As predicted by our analysis, the simulations show the increase in the variation of X_n when burstiness increases. On a bursty channel, the source input rate varies between

Figure 6: The variation of x_B as a function of b Figure 7: The variation of x_G as a function of b




Figure 8: The variation of $\bar{x}$ as a function of b

important values when the channel is in the good state and small values when the bad state appears.

5 Conclusions

In this paper, we have studied the performance of a TCP-like congestion control protocol over a lossy channel. In addition to the average loss rate considered in previous works, we evaluate in our work the impact on the performance of burstiness in the loss process. We define a model for losses using potential losses and a two-states Markov chain to account for the change in the channel state. This model lets us change the average loss rate as well as the correlation between losses. We calculate then the average as well as the variance of the throughput over a long time interval. Two approaches have been considered, a direct one and a Laplace Transform approach. The average throughput is compared to the throughput achieved by the protocol when it operates over a non-bursty lossy channel having the same average loss rate.




Figure 9: The variation of $Var(X_n)$ as a function of b

First, our mathematical analysis shows well a deterioration in the performance when the average loss rate increases. Second, at constant average loss rate, we show that the performance improves when losses tend to appear in bursts. This increase in performance with burstiness is caused by an increase in the second moments of the input rate. We conduct then 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 average throughput has to be calculated. Often in the literature, the source input rate (or the source window in case of TCP) has been supposed to vary in steady state between two fixed values w and $2w$. The expectation of the time between two losses is used to characterize this stationary behavior (calculate W and $2w$). The average throughput is taken equal $3w/2$. We show in our analysis that, in the general case, the second moment of the time between losses has to be considered. In the particular case of an exponentially distributed inter-losses time, the average input has to be taken equal to $2w$ rather than $3w/2$.

6 Appendix: Convergence of X_n to a stationary regime

We rewrite the dynamics (1) of X_n as

$$X_{n+1} = A_n X_n + B_n, \quad (11)$$

where

$$A_n = (1 - U_n/2)1\{Y_n = G\} + (1 - V_n/2)1\{Y_n = B\},$$

and

$$B_n = \alpha D_n.$$

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 (11), i.e. to a process X_n^* satisfying

$$X_{n+1}^* = A_n^* X_n^* + B_n^*,$$

for all $n \geq 0$. For any initial condition X_0 , we obtain by iterating (11)

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

for all $n \geq 0$.

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,*

$$X_n^* = \sum_{j=0}^{n-1} \left(\prod_{i=n-j}^{n-1} A_i \right) B_{n-j-1} \quad (12)$$

is the only stationary solution of (11). The sum on the right hand side of (12) converges absolutely almost surely. Furthermore, $|X_n - X_n^| \rightarrow 0$ 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_0^* as $n \rightarrow \infty$.*

Proof: We use Theorem 2A in [7] (based on [3, 17]). 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). $\diamond$

Remark 3 *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).*

References

- [1] Recommendation S.1062-1 of the l'ITU-R, "Allowable Error Performance for a Hypothetical Reference Digital Path Operating at or above the Primary Rate".
- [2] C. Barakat and E. Altman, "The tradeoff between TCP and link-level FEC", submitted to Infocom'2000.
- [3] A. Brandt, "The stochastic equation $Y_{n+1} = A_n Y_n + B_n$ with stationary coefficients", Adv. Appl. Prob., vol. 18, pp. 211-220, 1986.
- [4] E.N. Gilbert, "Capacity of a burst-noise channel", Bell Systems Technical Journal, Vol. 39, pp. 1253-1265, September 1960.
- [5] K. Fall and S. Floyd, "Simulation-based Comparisons of Tahoe, Reno, and SACK TCP", Computer Communications Review, July 1996.

- [6] S. Floyd, "Connections with Multiple Congested Gateways in Packet-Switched Networks Part 1: One-way Traffic", *Computer Communications Review*, Vol. 21, No. 5, October 1991.
- [7] P. Glasserman and D. D. Yao, "Stochastic vector difference equations with stationary coefficients", *Journal of Applied Probability*, vol. 32, pp. 851-866, 1995.
- [8] G. Koole and S. Asmussen, "Marked point processes as limits of Markovian arrival streams", *Journal of Applied Probability*, 1993.
- [9] V. Jacobson, "Congestion avoidance and control", in *Proc. ACM Sigcomm'88*, Stanford, CA, USA, August 1988.
- [10] A. Kumar, "Comparative Performance Analysis of Versions of TCP in a Local Network with a Lossy Link", *IEEE/ACM Transactions on Networking*, August 1998.
- [11] 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*, pp. 336-350, June 1997.
- [12] J. Lunsford, S. Narayanaswamy, D. Chitre and M. Neibert, "Link enhancement for ATM over satellite links", *ICDSC-10*, pp. 129-136, 1995.
- [13] S. McCanne and S. Floyd, NS (Network Simulator), 1995, available at URLs <http://www-nrg.ee.lbl.gov/ns>.
- [14] L. Mola, P. J. Lizcano, P. Meuret, H. Wännman and S. Agnelli, "An initiative for conformance ATM-networks relying on satellite-links", *COST 226 final symposium*, Vol. 1, pp. 518-523, 1995.
- [15] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP Throughput: a Simple Model and its Empirical Validation", *UMASS CMPSCI Tech Report TR98-008*, February 1998.
- [16] W. Stevens, "TCP Slow-Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms", *RFC 2001*, January 1997.
- [17] W. Vervaat, "On a stochastic difference equation and a representation of non-negative infinitely divisible random variables", *Adv. Appl. Prob.*, vol. 11, 750-783, 1979.



Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - B.P. 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Lorraine : Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - B.P. 101 - 54602 Villers lès Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot St Martin (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - B.P. 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, B.P. 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399