

Some Observations on Equation-Based Rate Control*

Milan Vojnović and Jean-Yves Le Boudec

Swiss Federal Institute of Technology

CH-1015 Lausanne, Switzerland

e-mail: {milan.vojnovic,jean-yves.leboudec}@epfl.ch

We consider one aspect of the general problem of unicast equation-based rate control in the Internet, which we formulate as follows. When a so called “loss-event” occurs, a data source updates its sending rate by setting it to $f(\hat{p}_n)$, where \hat{p}_n is an estimate of \bar{p} , the loss-event ratio. Function f (the target loss-throughput function) defines the objective of the control method: we would like that the throughput \bar{x} , attained by the source, satisfies the equation $\bar{x} \leq f(\bar{p})$. If so, we say that the control is conservative. In the Internet, function f is obtained by analyzing the dependency of throughput versus the loss-event ratio for a real TCP source. A non-TCP source which implements a control system as we describe is said to be TCP-friendly if the control is conservative. Our main finding is a set of two conditions whose conjunction is sufficient for the control to be conservative. Suppose that $1/\hat{p}_n$ is an unbiased estimator of $1/\bar{p}$. Then, it is sufficient that (1) $f(p)$ is a concave function with $1/p$, and (2) the expected time between two consecutive loss-events, given the current rate x , is non-increasing with x . To verify our finding, we formulate a model of the control and study numerical and simulation results for some special cases. We show there exists statistics of the loss-event inter-arrival times such that the control is non-conservative, even if $f(p)$ is a concave function of $1/p$, which should necessarily imply that our second sufficient condition does not hold in those cases. As a by-product, our theory explains why the control is overly conservative when the loss-event ratio is high. Another aspect of unicast equation-based rate control in the Internet is the influence of the variability of round-trip times, which is not analyzed in this paper.

1. Introduction

Our work is motivated by the concept of TCP friendliness. We quote [5] for the definition of TCP compatibility: “a TCP-compatible flow is responsive to congestion notification, and in steady-state it uses no more bandwidth than a conformant TCP running under comparable conditions”. TCP-friendliness [6] is a weaker property corresponding to the second part of the latter definition.

Consider a TCP sender adjusting its sending rate $x(t)$ in response to a sequence of loss-events, modeled as a stationary ergodic system. Then, the average throughput satisfies the “loss-throughput equation” $\mathbb{E}[x(t)] = f(\bar{p})$, where f captures the performance of TCP in typical network conditions; it is a function of the long-run average of the loss-event ratio

*ITC-17, 17th International Teletraffic Congress, Salvador da Bahia, Brazil, September 24–28, 2001.

\bar{p} , and also of other parameters such as the autocovariance of the loss-event inter-arrival times and some statistics of the round-trip time [10,11,2]. Consider also an adaptive non-TCP source with rate $x(t)$ which experiences a loss-event ratio \bar{p} ; the requirement to be TCP-friendly can be expressed as

$$\mathbb{E}[x(t)] \leq f(\bar{p}). \quad (1)$$

There is a large number of proposals on how to obtain this; see [1] and the references therein. We focus on one type of solutions which is exemplified in TFRC [7,9], and which we refer to as “equation-based rate control”. It can be formulated as follows. At the time of occurrence of the n -th loss-event, the source computes

$$X_n = f(\hat{p}_n), \quad (2)$$

where \hat{p}_n is an estimator of \bar{p} , the loss-event ratio. Function f (the target loss-throughput function) is obtained by TCP modeling as discussed above. The sending rate $x(t)$ is then set to X_n until the next loss-event:

$$x(t) = X_n, \text{ for } T_n \leq t < T_{n+1}, \quad (3)$$

where T_n is a time instant of the n -th loss-event.

Then we expect that such a source would be TCP-friendly. However, as we see later, this is not necessarily true.

In this paper, we focus on the simple control system described by Equation (2) and Equation (3) and ask whether Equation (1) holds. If so, we say that the control is conservative. This models one aspect of TCP-friendliness. Other aspects such as the influence of variability of the round-trip times are not analyzed in this paper.

We assume that our system, and in particular the sequence of the loss-event inter-arrival times, is stationary and ergodic. Indeed, this is a minimal assumption that we may impose. For the validity of this assumption in the real-world Internet we refer to the experimental study by Zhang, Paxson, and Shenker [13].

We need to distinguish between the expected value of X_n , namely the average rate seen at update instants, from the expected value of $x(t)$ (they are assumed to be independent of n or t , by the stationarity assumption). The expectation of X_n is the expected rate at special instants, and is traditionally called a Palm expectation [3]. We denote it with $\mathbb{E}_T^0[x(t)] = \mathbb{E}[X_0]$ and call it the “rate Palm expectation”. We call $\mathbb{E}[x(t)]$ the “rate time average”, or simply the “throughput”. In general, Palm expectations differ from time averages, even if the special instants form a Poisson process.

Our first observation, in Section 2, is that, if $1/\hat{p}_n$ is an unbiased estimator of $1/\bar{p}$, and if $f(p)$ is concave with $1/p$, then

$$\mathbb{E}[X_0] \leq f(\bar{p}). \quad (4)$$

The condition that $1/\hat{p}_n$ is an unbiased estimator of $1/\bar{p}$ is true for smoothed averages such as in TFRC [7,9]. The condition that $f(p)$ is concave with $1/p$ is true for the simple TCP loss-throughput formulas in [10] (the “square-root” formula), but is no longer true for exact (28) and approximate (30) formulas in [11] with non-zero retransmission timeout parameter (they are convex with respect to small values of $1/p$).

Our main result (Theorem 1 in Section 2) gives us sufficient conditions under which the control is conservative, that is, our control objective (1) is met. The result tells us that if $1/\hat{p}_n$ is an unbiased estimator of $1/\bar{p}$, $f(p)$ is concave with $1/p$, and $\mathbb{E}[S_n|X_n = x]$ is non-increasing with x , then the control (2) is conservative.

Under these conditions, we show that the throughput is less than or equal to the rate Palm expectation. Then, due to inequality (4) the control is conservative. Thus, we identify two elements that contribute to the conservative nature of the control under the given conditions. The first one is the fact of updating the rate at loss-event occurrences, which is the origin of the difference between the throughput and the rate Palm expectation. The second one is the concavity of the function $f(p)$ with respect to $1/p$. For the design of a unicast equation-based rate control, it tells us that, under our conditions, the rate control should indeed be TCP-friendly. In fact, if $f(p)$ is strictly concave with $1/p$, we should expect that such a rate control would give less rate than TCP would.

We show numerically that if $f(p)$ is not concave with $1/p$, the control may be non-conservative. The case where $\mathbb{E}[S_n|X_n = x]$ is not non-increasing with x is not addressed explicitly. However, in Section 4, we formulate a model and show numerically and by simulation that even for $f(p)$ concave with $1/p$ we can construct statistics of the loss-event inter-arrival times such that the control is non-conservative. This should indeed imply that $\mathbb{E}[S_n|X_n = x]$ is not non-increasing with x in those cases. An explicit verification of this is work in progress.

Note that in our modeling we do not aim at providing the best, or even a good, model of the interaction between a rate controlled source and the network. In contrast, we consider models that are reasonable approximations of some limiting cases, and use them as a framework in which we can give some answers to our original question.

We model the loss-event inter-arrival times as an exogeneous stationary random process. For tractability, we analyze the case where the loss event process is driven by a hidden Markov chain. A system which can be reasonably modeled by this is an audio source which adjusts its rate not by sending fewer packets, but by adapting the packet size, assuming that the network drops packets independent of their size. Thus for such a source, the dropping events depend on the rate of the source only through the state of the hidden Markov chain. We show that there exists statistics of the loss-event inter-arrival times such that the control is non-conservative, even if $f(p)$ is concave with $1/p$. This suggests that if the packet dropping probability in a RED gateway is independent of packet size, then such an audio application may exceed its TCP-friendly rate. The exact solving for the throughput of our model seems to be out of our reach, even with our hidden Markov chain assumption. Our approach to circumvent that difficulty is as follows. First, we solve the model numerically in the special case where the estimator \hat{p}_n is not smoothed. This gives us one exact solution and is used as a benchmark. Second, for the general case of this model, we use a classical approximation technique which consists in linearizing the system around its equilibrium. We confirm our approximations by a discrete-event simulation of the control system (not of TFRC, since we want to avoid the effect of other factors). Our computations show that for a two-state hidden Markov chain, the autocorrelation structure of loss-events plays a key role.

Lastly, in Section 5, as a by-product, we give an explanation to the empirically observed phenomena that TFRC throughput is overly conservative for high loss-event ratios.

2. Preliminary results

We use the following notation. Let $\{S_n\}_{n=-\infty}^{\infty}$ be a stationary ergodic stochastic process representing the loss-event inter-arrival times. Then, $\{T_n\}_{n=-\infty}^{\infty}$ is the associated point process with intensity $\lambda = 1/\mathbb{E}_T^0[S_0]$. Stationarity implies that the law of $S_n, S_{n+1}, \dots, S_{n+m}$ is independent of n .

Let θ_n be the number of the packets sent in $[T_n, T_{n+1})$; following [7] we call this the loss-event interval. We consider a positive-valued non-increasing function f with respect to the loss-event ratio \hat{p}_n ; defined on $[0, 1]$. In the Internet, f is also a function of the round-trip time, and the retransmission timeout, but in this paper we consider those parameters to be fixed, equal to τ and ρ , respectively.

For notation convenience, we define $f_{\theta}(\hat{\theta}_n) = f(1/\hat{\theta}_n)$, where $\hat{\theta}_n$ is an estimator of the expected loss-event interval; thus the rate control formula in (2) is equivalent to

$$X_{n+1} = f_{\theta}(\hat{\theta}_n). \quad (5)$$

We suppose the control (5) admits stability conditions [8,4], under which there exists a unique finite stationary rate process.

Proposition 1 *Suppose $\hat{\theta}_n$ is an unbiased estimator of $1/\bar{p}$. If $f(p)$ is concave with $1/p$, then*

$$\mathbb{E}[X_0] \leq f_{\theta}(\mathbb{E}[\theta_0]) = f(\bar{p}).$$

Proof. For $\hat{\theta}_n$ to be an unbiased estimator of $1/\bar{p}$ means that $\mathbb{E}[\hat{\theta}_n] = \mathbb{E}[\theta_n] = 1/\bar{p}$. The proof is then a direct application of Jensen's inequality to (5). \square

Remark 1 *Note that we request $\hat{\theta}_n$ is an unbiased estimator of $1/\bar{p}$, which is inline with [7,9]. Under this assumption, \hat{p}_n is a biased estimator of the loss-event ratio, that is, $\mathbb{E}[\hat{p}_n] \geq \bar{p}$, with equality iff \hat{p}_n is fixed (Jensen's inequality again).*

It follows from this proposition that the control is conservative as soon as $\mathbb{E}[x(t)] \leq \mathbb{E}[X_0]$. Whether this holds depends on the properties of the point process; in general, we should expect that time averages differ from Palm averages. In the section to follow we identify sufficient conditions under which it indeed holds that $\mathbb{E}[x(t)] \leq \mathbb{E}[X_0]$. Before getting there, we derive general relations which will be used in the sequel.

The rate Palm average and the throughput are related by Palm's inversion formula [3]:

$$\mathbb{E}[x(t)] = \lambda \mathbb{E}_T^0 \left[\int_0^{T_1} x(t) dt \right] = \lambda \mathbb{E}_T^0 [X_0 S_0]. \quad (6)$$

Also observe that the long-run average of the loss-event ratio \bar{p} is equal to $1/\mathbb{E}[\theta_0]$. Indeed,

$$\bar{p} = \lim_{t \rightarrow \infty} \frac{N(t)}{\int_0^t x(t) dt} = \frac{1}{\lim_{t \rightarrow \infty} \left\{ \frac{1}{N(t)} \sum_{n=0}^{N(t)-1} \theta_n + \frac{1}{N(t)} \theta(t - T_{N(t)}, t) \right\}} = \frac{1}{\mathbb{E}[\theta_0]}.$$

Combining the latter equation with (6) it follows: $\mathbb{E}[x(t)] = \lambda/\bar{p}$. Thus the control is conservative if $\mathbb{E}[x(t)] \leq f(\lambda \mathbb{E}[x(t)]^{-1})$, or equivalently: $\mathbb{E}[x(t)] \leq f_{\theta}(\lambda^{-1} \mathbb{E}[x(t)])$.

In some cases (for example the square-root formula in [10]), the equation $x = f_{\theta}(\lambda^{-1} x)$ has a unique positive solution x^* and then the control is conservative if $\mathbb{E}[x(t)] \leq x^*$. In such cases, we call x^* the maximum TCP-friendly throughput.

Table 1

If $f(p)$ is not concave with $1/p$ then the control may be non-conservative – an example with $f(p) = e^{a/p} + b$, $a = 0.2$, $b = 9$, around $\bar{p} = 1/6$. The results are normalized with $f(\bar{p})$; the normalized throughput larger than 1 indicate a non-conservative control. We use $\hat{\theta}_n = \sum_{l=1}^L w_l \theta_{n-l+1}$ (with TFRC weights).

L	$(\mathbb{E}[X_0] \pm CI_{\mathbb{E}[X_0]}^{95\%})/f(\bar{p})$	$(\mathbb{E}[x(t)] \pm CI_{\mathbb{E}[x(t)]}^{95\%})/f(\bar{p})$
1	5.0268 ± 2.2195	1.0325 ± 0.0048
2	1.3932 ± 0.2011	1.0258 ± 0.0051
4	1.0852 ± 0.0119	1.0189 ± 0.0060
8	1.0360 ± 0.0054	1.0109 ± 0.0045
16	1.0166 ± 0.0060	1.0069 ± 0.0061

3. Main result

Theorem 1 *Suppose $\hat{\theta}_n$ is an unbiased estimator of $1/\bar{p}$. If (C1) $f(p)$ is concave with $1/p$, and (C2) $\sigma(x) := \mathbb{E}[S_0|X_0 = x]$ is non-increasing with x , then*

$$\mathbb{E}[x(t)] \leq f(\bar{p}), \quad (7)$$

in other words, the control is conservative.

Proof. Note that

$$\mathbb{E}[x(t)] = \frac{\mathbb{E}[X_0 S_0]}{\mathbb{E}[S_0]} = \frac{\mathbb{E}[X_0 \sigma(X_0)]}{\mathbb{E}[\sigma(X_0)]}. \quad (8)$$

We use Harris' inequality (e.g., see [3], p. 225) which states the following. For g and h , both non-decreasing functions, and X a random variable, $\mathbb{E}[g(X)h(X)] \geq \mathbb{E}[g(X)]\mathbb{E}[h(X)]$. From (C2), $-\sigma$ is indeed a non-decreasing function, thus Harris' inequality gives us

$$\mathbb{E}[X_0 \sigma(X_0)] \leq \mathbb{E}[X_0] \mathbb{E}[\sigma(X_0)]. \quad (9)$$

Combining (8) with (9), we obtain $\mathbb{E}[x(t)] \leq \mathbb{E}[X_0]$, thus the rate time average is less than equal to the rate Palm expectation. Finally, given (C1) and $\hat{\theta}_n$ is an unbiased estimator of $1/\bar{p}$, Proposition 1 completes the proof of the theorem. \square

3.1. Discussion of the result

We show by an example that if (C1) does not hold and (C2) hold, then the control may be non-conservative. To that end, take $f(p) = e^{a/p} + b$, for positive-valued constants a and b . Suppose $\mathbb{P}(S_0 > s|X_0 = x) = e^{-\lambda x s}$, then $\mathbb{E}[S_0|X_0 = x] = \frac{1}{\lambda x}$. In Table 1, we show estimates of the rate Palm expectation and throughput computed out of 5 random experiments. The results show us that the control is non-conservative. Another case to consider would be that (C1) hold, but not (C2). We do not show that explicitly in the present paper (for our models in Section 4 we were not able to explicitly compute $\sigma(x)$; we defer this to further study). However, in Section 4.4 we show numerical results for which (C1) holds, but the control is non-conservative, which should necessarily imply that (C2) does not hold in those cases.

Condition (C2), $\sigma(x)$ is non-increasing with x , is indeed verified for memoryless packet dropping where the source has a negligible impact on the network state. For example, consider a simple model (e.g. [11]) where packets are dropped with fixed probability q . Then, it is readily shown that $\sigma(x) = \frac{1}{qx}$, and the condition (C2) is verified. Even more generally, for packet losses that are governed by some hidden discrete-time Markov chain with transitions at the packet transmission times, (C2) would hold. However, if we model dynamics of the probability to drop a packet with a hidden continuous-time Markov chain than it may be shown there exist cases such that $\sigma(x)$ is not non-increasing with x . Nevertheless, it remains to be seen how the condition (C2) is verified in the Internet.

4. Validation by modeling and simulation

In this section we first formulate a general model that accommodate any function f and any statistics of the loss-event inter-arrival times. For tractability, we then study two special cases: the square-root formula with not smoothed $\hat{\theta}_n$ and then linearization of our originally non-linear model with smoothed $\hat{\theta}_n$. In the former case, we give a procedure to compute the throughput numerically, while in the latter, we obtain a closed-form solution; both for a hidden Markov chain that we introduce shortly.

4.1. General model

We suppose a sliding-window estimator of the loss-event interval (proposed in [7,9])

$$\hat{\theta}_n = \sum_{l=1}^L w_l \theta_{n-l+1}, \quad (10)$$

where $\{w_1, w_2, \dots, w_L\}$ are positive-valued weights such that $\sum_{l=1}^L w_l = 1$. Given $x(t)$ is piece-wise constant, $\theta_n = X_n S_n$. Thus

$$X_{n+1} = f_\theta\left(\sum_{l=1}^L w_l S_{n-l+1} X_{n-l+1}\right). \quad (11)$$

We observe that the model corresponds to a non-linear autoregressive process with stationary random coefficients (in general, for non-independent $\{S_n\}$, (11) is a special case of stochastic recursive sequence; resp., for independent $\{S_n\}$, (11) defines a L -dimensional Markov chain [4]). We note that non-linearity of f_θ renders the analytical analysis difficult.

We setup a hidden Markov chain as follows. Let $\{Z_n\}_{n=-\infty}^{\infty}$ be an irreducible positive recurrent discrete-time Markov chain with a finite state space $E = \{1, 2, \dots, N\}$, where $N = |E|$. Let $\mathbf{P} = \{p_{ij}\}_{i,j \in E}$ be the transition matrix, and $\boldsymbol{\pi} = [\pi_1, \pi_2, \dots, \pi_N]^T$ be the stationary distribution. Suppose that

$$\mathbf{P}(S_n \leq s | Z_{n+1-k}, S_{n-k}, k = 1, 2, \dots) = \mathbf{P}(S_n \leq s | Z_n). \quad (12)$$

Note that, in fact, $\{S_n\}$ is a Markov renewal process. We introduce additional notation $F_i(s) := \mathbf{P}(S_n \leq s | Z_n = i)$, $\gamma_i(u) := \mathbb{E}[S_n^u | Z_n = i]$, for all $i \in E$, $\boldsymbol{\gamma}(u) = [\gamma_1(u), \gamma_2(u), \dots, \gamma_N(u)]^T$, and $\boldsymbol{\Gamma}(u) = \text{diag}(\boldsymbol{\gamma}(u))$. In order to easy notation we use γ_i , $\boldsymbol{\gamma}$, $\boldsymbol{\Gamma}$ for $\gamma_i(1)$, $\boldsymbol{\gamma}(1)$, and $\boldsymbol{\Gamma}(1)$, respectively.

4.2. Special case I: the square-root formula with $L = 1$

In this section, we study a special case of (11), given as follows

$$X_{n+1} = K \sqrt{X_n S_n}, \quad (13)$$

where $K > 0$. We recover a linear recurrence with stationary random coefficients by taking the logarithm in (13) to obtain

$$\tilde{X}_{n+1} = \frac{1}{2}\tilde{X}_n + \frac{1}{2}\tilde{S}_n + \tilde{K}_n,$$

where $\tilde{X}_n := \ln X_n$, $\tilde{S}_n := \ln S_n$, for any $n \in \mathbb{Z}$, and $\tilde{K} := \ln K$.

Then, under the stability conditions [8,4], there exists a unique finite stationary \tilde{X}_n^* such that $|\tilde{X}_n - \tilde{X}_n^*| \rightarrow 0$, as $n \rightarrow \infty$, P-a.s.. From this, it follows

$$X_n^* = K^2 \prod_{k=1}^{\infty} S_{n-k}^{\frac{1}{2^k}}. \quad (14)$$

Computing the throughput by (6) with (14), we obtain

$$\mathbb{E}[x^*(t)] = K^2 \mathbb{E}\left[\prod_{k=0}^{\infty} S_{-k}^{\frac{1}{2^k}}\right]. \quad (15)$$

Note that the maximum TCP-friendly throughput solves $\mathbb{E}[x(t)] = f_{\theta}(\lambda^{-1}\mathbb{E}[x(t)])$, which for the square-root formula yields $\mathbb{E}[x(t)] = \frac{K^2}{\lambda}$. Thus, the rate control (13) is conservative if $\mathbb{E}[\prod_{k=0}^{\infty} S_{-k}^{\frac{1}{2^k}}] \leq \mathbb{E}[S_0]^2$.

Now we use the hidden Markov chain. Let $\Phi_i(m) = \mathbb{E}[\prod_{k=0}^m S_{-k}^{\frac{1}{2^k}} | Z_{-m} = i]$, and $\Phi(m) = [\Phi_1(m), \dots, \Phi_N(m)]^T$, $m \geq 0$. Then, it may be verified

$$\Phi(m) = \Gamma(2^{-m}) \mathbf{P} \Phi(m-1), \quad m > 0. \quad (16)$$

Finally, we obtain

$$\mathbb{E}[x^*(t)] = K^2 \lambda \pi^T \Phi(\infty).$$

The recursion (16) allows a numerical solving of the throughput. In the sequel, we give an alternative approach. Let $\tilde{\theta}_n = \ln \theta_n$, then $\mathbb{E}[\theta_0^*] = \mathbb{E}[e^{r\tilde{\theta}_0}] = \sum_{k=0}^{\infty} \frac{r^k}{k!} \mathbb{E}[\tilde{\theta}_0^k]$. Thus,

$$\mathbb{E}[x^*(t)] = \lambda \mathbb{E}[\theta_0] = \lambda \mathbb{E}[e^{\tilde{\theta}_0}] = \lambda \sum_{k=0}^{\infty} \frac{\mathbb{E}[\tilde{\theta}_0^k]}{k!}. \quad (17)$$

We need to compute the moments of $\tilde{\theta}_0$. To that end, define $\psi_i(u) = \mathbb{E}[e^{u\tilde{\theta}_0} | Z_n = i]$. Then, we identify the following functional:

$$\psi_i(u) = \sum_{j \in E} \frac{\pi_j}{\pi_i} p_{ji} K^u \mathbb{E}[e^{u\tilde{S}_{n+1}} | Z_{n+1} = i] \psi_j\left(\frac{u}{2}\right), \quad i \in E. \quad (18)$$

We omit derivation of (18) due to the space limitations. It is in general difficult to solve a functional in a closed-form, but at least a numerical solution can be retrieved. Instead, from (18), we derive the r -th conditional moment $\vartheta_i(r) := \mathbb{E}[\tilde{\theta}_0^r | Z_0 = i] = \frac{d^r \psi_i(u)}{du^r} \Big|_{u=0}$ and then $\mathbb{E}[\tilde{\theta}_0^r] = \sum_{i \in E} \vartheta_i(r) \pi_i$. Taking the r -th derivative of (18), at $u = 0$, we obtain the following system of linear equations:

$$\left(1 - p_{ii} \frac{1}{2^r}\right) \vartheta_i(r) - \sum_{j \in E, j \neq i} \frac{\pi_j}{\pi_i} p_{ji} \frac{1}{2^r} \vartheta_j(r) = \alpha_i(r), \quad i \in E, \quad (19)$$

where $\alpha_i(r) = \sum_{j \in E} \frac{\pi_j}{\pi_i} p_{ji} \alpha_{ij}(r)$ and

$$\alpha_{ij}(r) = \sum_{k=0}^{r-1} \frac{1}{2^k} \binom{r}{k} \left[\sum_{l=0}^{r-k} \binom{r-k}{l} (\ln K)^l \mathbb{E}[(\ln S_0)^{r-k-l} | Z_0 = i] \right] \vartheta_j(k).$$

Thus, it remains to solve the system (19), recursively, for $r \geq 0$, and then use this solution in (17). Similarly, we obtain a formulation for calculating the moments of the rate in respect to the Palm probability, but we omit that here due to the space limitations.

4.3. Special case II: the linearized system with $L \geq 1$

We now come back to the general case where $L \geq 1$ and any function f . We linearize the non-linear system (11) around $\mathbb{E}[\theta_0]$, and then obtain

$$X_{n+1} = a \sum_{l=1}^L w_l S_{n-1+l} X_{n-l+1} + b, \quad (20)$$

where for brevity $a := f'_\theta(\lambda^{-1} \mathbb{E}[x(t)])$, and $b := f_\theta(\lambda^{-1} \mathbb{E}[x(t)]) - f'_\theta(\lambda^{-1} \mathbb{E}[x(t)]) \lambda^{-1} \mathbb{E}[x(t)]$.

Note that we lowered the complexity of the system (5) to a linear autoregressive process with stationary random coefficients. By limited Taylor expansion, there exists $0 < \vartheta_n < 1$ such that

$$f_\theta(\hat{\theta}_n) = f_\theta(\mathbb{E}[\theta_0]) + (\hat{\theta}_n - \mathbb{E}[\theta_0]) f'_\theta(\mathbb{E}[\theta_0]) + \frac{1}{2} (\hat{\theta}_n - \mathbb{E}[\theta_0])^2 f''_\theta(\mathbb{E}[\theta_0]) + \vartheta_n (\hat{\theta}_n - \mathbb{E}[\theta_0]). \quad (21)$$

This shows a rather intuitive observation that for concave f_θ , $f''_\theta \leq 0$, the rate Palm expectation of the linearized system is an upper-bound to that of the original non-linear system (11) (resp. a lower-bound for convex f_θ , $f''_\theta \geq 0$).

Now, (20) can be cast into a form of L -dimensional linear recurrence with stationary random coefficients, and then, under stability conditions, it follows that $|X_n - X_n^*| \rightarrow 0$, as $n \rightarrow \infty$, with \mathbf{P} -a.s., where

$$\begin{aligned} X_n^* = & b \left\{ 1 + a \sum_{l=1}^L w_l S_{n-l} + a^2 \sum_{l=1}^L \sum_{k=1}^L w_l w_k S_{n-l-k} S_{n-l} + \right. \\ & \left. + a^3 \sum_{l=1}^L \sum_{k=1}^L \sum_{m=1}^L w_l w_k w_m S_{n-l-k-m} S_{n-l-k} S_{n-l} + \dots \right\}. \end{aligned} \quad (22)$$

Then, by plugging the latter into (6) we obtain

$$\begin{aligned} \mathbb{E}[x^*(t)] = & b \left\{ 1 + \lambda a \sum_{l=1}^L w_l R(l) + \lambda a^2 \sum_{l=1}^L \sum_{k=1}^L w_l w_k m_3(k, l) + \right. \\ & \left. + \lambda a^3 \sum_{l=1}^L \sum_{k=1}^L \sum_{m=1}^L w_l w_k w_m m_4(m, k, l) + \dots \right\}, \end{aligned} \quad (23)$$

where by definition $R(l) = \mathbb{E}[S_0 S_{-l}]$, $m_3(n, m) = \mathbb{E}[S_0 S_{-m} S_{-m-n}]$, $l, n, m \in \mathbb{N}$, and analogously we define $m_p(\dots)$ for $p > 3$.

We observe that the throughput (23) does not admit a simple form in terms of the autocovariance of the loss-event inter-arrival times, solely, but it depends on the higher-order autocorrelation terms. It is worth mentioning that the throughput does not depend on the variance of the loss-event inter-arrival times, which is not the case with TCP [2]. It may be shown that the underlying reason is that in our model we suppose the rate is kept constant in-between two consecutive loss-events, while TCP increases the rate, and this is where the variance term comes out.

Now we come back to the hidden Markov chain and use the notation introduced at the beginning of Section 4. It can be easily shown that

$$R(l) = \mathbb{E}[S_0 S_l] = \begin{cases} \sum_{i \in E} \mathbb{E}[S_0^2 | Z_n = i] \pi_i, & l = 0 \\ \boldsymbol{\pi}^T \mathbf{\Gamma} \mathbf{P}^l \boldsymbol{\gamma}, & l > 0 \end{cases}, \quad (24)$$

and $m_3(l, k) = \mathbb{E}[S_0 S_k S_{k+l}] = \boldsymbol{\pi}^T \mathbf{\Gamma} \mathbf{P}^k \mathbf{\Gamma} \mathbf{P}^l \boldsymbol{\gamma}$ with analogous definitions of $m_p(\dots)$, for $p > 3$. Let $\mathbf{G} = \mathbf{\Gamma} \sum_{l=1}^L w_l \mathbf{P}^l$. Then, it may be verified

$$\mathbb{E}[x^*(t)] = b\lambda \boldsymbol{\pi}^T (\mathbf{I} - a\mathbf{G})^{-1} \boldsymbol{\gamma}. \quad (25)$$

In the remainder of this section we obtain a closed-form solution to the throughput in (25), for two-state Markov modulated events. To that end, let \mathbf{P} be defined such that $p_{12} = p$ and $p_{21} = q$, for $p, q \in [0, 1]$. Then, from (25) we obtain the throughput is a solution to:

$$\begin{aligned} & \mathbb{E}[x^*(t)] - a\mathbb{E}[x^*(t)](g_{11} + g_{22} - \frac{1}{\lambda}) - a^2\mathbb{E}[x^*(t)](g_{12}g_{21} - g_{11}g_{22} - B) - \\ & - f_\theta(\lambda^{-1}\mathbb{E}[x^*(t)]) - \lambda B a f_\theta(\lambda^{-1}\mathbb{E}[x^*(t)]) = 0, \end{aligned} \quad (26)$$

where, for brevity, $B := \gamma_1(g_{21}\pi_2 - g_{22}\pi_1) + \gamma_2(g_{12}\pi_1 - g_{11}\pi_2)$.

Thus, we identified the implicit function (26) whose root is the throughput $\mathbb{E}[x^*(t)]$ for a given f . In the sequel, we suppose the square-root formula, $f_\theta(\theta) = \frac{K}{\tau} \sqrt{\theta}$, for which a closed-form solution is easily retrieved. Indeed, from (26) the throughput is a root of the quadratic equation:

$$\left(\sqrt{\mathbb{E}[x^*(t)]} \right)^2 - \frac{K\sqrt{\lambda}}{2\tau} \left(g_{11} + g_{22} + \frac{1}{\lambda} \right) \sqrt{\mathbb{E}[x^*(t)]} - \frac{K^2\lambda}{4\tau^2} (g_{12}g_{21} - g_{11}g_{22} + B) = 0. \quad (27)$$

And, finally, the solution is

$$\mathbb{E}[x^*(t)] = \frac{K^2}{\tau^2\lambda} e, \quad (28)$$

where we define

$$e = \frac{\lambda^2}{16} \left\{ g_{11} + g_{22} + \frac{1}{\lambda} \pm \sqrt{(g_{11} + g_{22} + \frac{1}{\lambda})^2 + 4(g_{12}g_{21} - g_{11}g_{22} + B)} \right\}^2. \quad (29)$$

Recall that the maximum TCP-friendly throughput is equal to $\mathbb{E}[x(t)] = \frac{K^2}{\tau^2\lambda}$. In this respect, e is a bias factor. Ideally, when $e = 1$ there is no bias, otherwise, for $e > 1$ the bias is non-conservative, and for $e < 1$ the bias is conservative.

4.4. Numerical Results

We ran random experiments over 10000 samples of the loss-event inter-arrival times, and estimate the rate Palm expectation $\mathbb{E}[X_0]$, and the throughput $\mathbb{E}[x(t)]$, as the average over 5 such runs, with associated 0.95-confidence intervals.

In Fig. 1, we confront our analytical results for two-state Markov modulated loss-events (with parameters p and q) to the simulation results. In Fig. 1 (Left) we show the results for the square-root formula with $L = 1$, and in Fig. 1 (Right) the results for the linearized system with $L = 8$ (TFRC weights). We use a simple two-state Markov modulated loss-events, such that the residence time in the state 1 is fixed to $\gamma_1 = 0.2$, and γ_2 is set such that $\mathbb{E}[S_0] = 1$ sec. Given these values, we may interpret the state 1 as a congestion state, and the state 2 as a non-congestion state. The results in Fig. 1 are shown normalized with $f(\bar{p})$; thus, the values smaller than 1 correspond to conservative throughput (resp. the values larger than 1 to non-conservative throughput). We note that in all the cases indeed it holds $\mathbb{E}[X_0] \leq f(\bar{p})$, which verifies Proposition 1. However, for small p we observe that the throughput is non-conservative. Thus, there exist statistics of the loss-event inter-arrival times such that the throughput is non-conservative. We expect that in those cases our condition (C2) in Theorem 1 necessarily does not hold. It remains to verify this explicitly. For the square-root formula with $L = 1$ (Fig. 1, Left) there is a perfect match between the analytical and simulation results. For the square-root formula with $L = 8$ (Fig. 1, Right), the analytical results that we derive upon the linearized system deviate from the simulations for larger q to p ratio. It is readily verified that the variance of the loss-event inter-arrival times is equal to $\frac{16q}{15p}$, which explains the observed deviation. We observe that non-conservative throughput happens to be for positive autocorrelation of the loss-event inter-arrival times (not shown here due to the space limitations).

5. Discussion and conclusion

It was observed through empirical studies that TFRC is overly conservative for high loss-event ratios (e.g., see [12]), however, no explanation is provided. Our observations made in the present paper suggest an answer to this. It is due to randomness of the loss-event ratio estimator \hat{p}_n , and non-linearity involved due to both updating the rate at the loss events and non-linearity of the function f . We support this with a simple example. Suppose Bernoulli (q) packet loss model, thus $\mathbb{P}(S_n > \frac{n}{x} | X_n = x) = (1 - q)^n$, $n = \{1, 2, \dots\}$. In this case it holds $\mathbb{E}[x(t)] = 1/\mathbb{E}[1/X_0]$. Take the function f proposed in [9], $f(p) = 1/(\tau a p^{1/2} + \rho b p^{3/2} + \rho c p^{5/2})$ pcks/sec, where τ is the round-trip time, ρ TCP retransmission timeout, and a, b, c positive-valued constants. Then, we obtain

$$\mathbb{E}[x(t)] = \frac{1}{\tau a \mathbb{E}[\hat{p}_0^{1/2}] + \rho b \mathbb{E}[\hat{p}_0^{3/2}] + \rho c \mathbb{E}[\hat{p}_0^{5/2}]} \leq \frac{1}{\tau a \bar{p}^{1/2} + \rho b \bar{p}^{3/2} + \rho c \bar{p}^{5/2}},$$

where the inequality follows from our assumption that $1/\hat{p}_n$ is an unbiased estimator of $1/\bar{p}$, and then by Jensen's inequality $\mathbb{E}[\hat{p}_0^{1/2}] \geq \bar{p}^{1/2}$, $\mathbb{E}[\hat{p}_0^{3/2}] \geq \bar{p}^{3/2}$, and $\mathbb{E}[\hat{p}_0^{5/2}] \geq \bar{p}^{5/2}$. The last two terms in the denominator of the function f come into play for small $1/\hat{p}_n$; the terms are very steep in this region and their convexity gear the throughput into conservative direction. We did verify this through simulation, which we do not show here due to the space limitation. An obvious fix to lessen this effect is to increase smoothness of \hat{p}_n , but that would diminish responsiveness of the control.

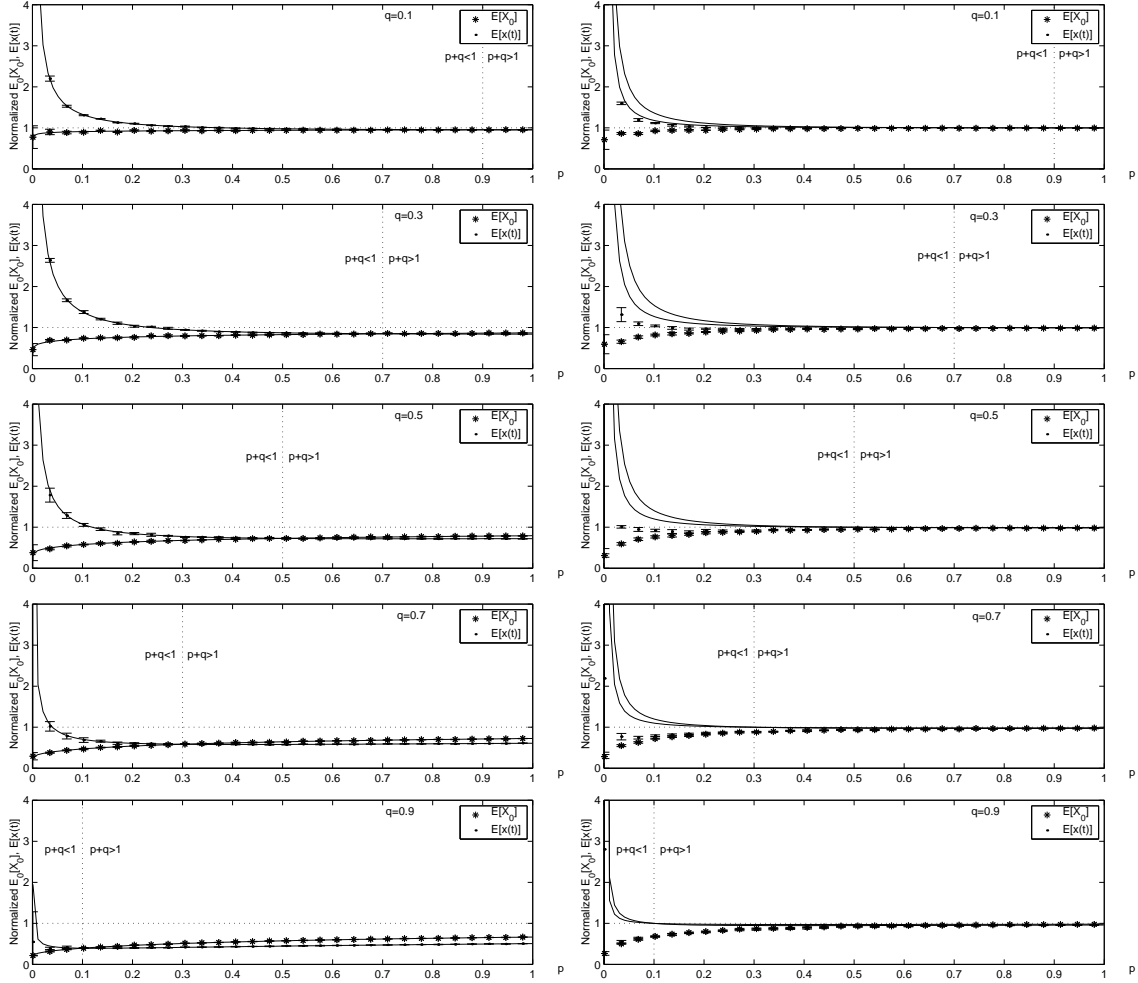


Figure 1. Numerical and simulation results for the square-root formula with $L = 1$ (Left), and the square-root formula with $L = 8$ and TFRC weights (Right). Simulation estimates of the throughput $\mathbb{E}[x(t)]$ and rate Palm expectation $\mathbb{E}[X_0]$ are denoted with dots and stars, respectively. Solid lines are analytical results ($\mathbb{E}[X_0]$ top line, $\mathbb{E}[x(t)]$ bottom line). The hidden Markov chain is with parameters p and q ; x -axis is p , and q ranges from 0.1 to 0.9, top to bottom, respectively. The results are normalized with maximum TCP-friendly throughput $\frac{K^2}{\tau^2\lambda}$; values at y -axis larger than 1 indicate a non-conservative bias (above the dotted horizontal line). Other parameters are set as follows: $1/\lambda = 1$ sec, $\gamma_1 = 0.2$ sec, $\gamma_2 = 1 + \frac{4}{5}p$ sec, $K = \sqrt{3/2}$, and $\tau = 0.1$ sec.

We briefly comment on one of the differences of the control we study and TFRC. TFRC uses an additional mechanism that amounts to update the rate not only at the loss-event occurrences. Call X the number of packets sent since the latest loss-event. If the estimator of the loss-event interval increases with X included as a sample, then update the loss-event interval estimator. Intuitively, we may expect that this mechanism would gear the throughput into a non-conservative direction. In our linearized system we were able to identify an additive positive term that adds to the throughput expression (not shown here). The term is increasing with the variance of the loss-event inter-arrival times.

There are many avenues to continue with future work. For instance: effects of randomness of the round-trip times; impact of additional control mechanisms; stability conditions.

REFERENCES

1. The TCP-friendly website. http://www.psc.edu/networking/tcp_friendly.html.
2. Eitan Altman, Konstantin Avrachenkov, and Chadi Barakat. A stochastic model of TCP/IP with stationary random losses. In *Proc. of the Sigcomm'00*, pages 231–242, 2000.
3. Francois Baccelli and Pierre Brémaud. *Elements of Queueing Theory*, volume 26. Applications of Mathematics, Springer-Verlag, 1991.
4. A. A. Borovkov. *Ergodicity and Stability of Stochastic Processes*. John Wiley & Sons, Wiley Series in Probability and Statistics, 1998.
5. B. Braden, D. Clark, J. Crowcroft, B. Davie, S. Deering, D. Estrin, S. Floyd, V. Jacobson, G. Minshall, C. Patridge, L. Peterson, K. Ramakrishnan, S. Shenker, J. Wroclawski, and L. Zhang. Recommendations on queue management and congestion avoidance in the Internet. *IETF RFC 2309*, April 1998.
6. Sally Floyd and Kevin Fall. Promoting the use of end-to-end congestion control in the Internet. *IEEE/ACM Trans. on Networking*, 7(4):458–472, August 1999.
7. Sally Floyd, Mark Handley, Jitendra Padhye, and Jörg Widmer. Equation-based congestion control for unicast applications. In *Proc. of the Sigcomm'00*, pages 43–56, 2000.
8. P. Glasserman and D.D. Yao. Stochastic vector difference equations with stationary coefficients. *Journal on Applied Probability*, 32:851–866, 1995.
9. Mark Handley, Jitendra Padhye, Sally Floyd, and Jörg Widmer. TCP friendly rate control (TFRC) protocol specification, ietf internet-draft, May 2001. <http://search.ietf.org/internet-drafts/draft-ietf-tsvwg-tfrc-02.txt>.
10. Matthew Mathis, Jeffrey Semke, Jamshid Mahdavi, and Teunis Ott. The macroscopic behavior of the TCP congestion avoidance algorithm. *Computer Communication Review*, 27(3), July 1997.
11. Jitendra Padhye, Victor Firoiu, Don Towsley, and Jim Kurose. Modeling TCP Reno performance: A simple model and its empirical validation. *IEEE/ACM Trans. on Networking*, 8(2):133–145, 2000.
12. Y. Richard Yang, Min Sik Kim, and Simon S. Lam. Transient behaviors of TCP-friendly congestion control protocols. In *Proc. of IEEE Infocom'2001*, March 2001.
13. Y. Zhang, V. Paxson, and S. Shenker. The stationarity of Internet path properties: Routing, loss, and throughput. Technical report, ACIRI, May 2000.