On the Long-Run Behavior of Equation-Based Rate Control

Milan Vojnović and Jean-Yves Le Boudec

EPFL

CH-1015, Lausanne, Switzerland

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

ABSTRACT

We consider unicast equation-based rate control, where a source estimates the loss event ratio p , and, primarily at loss events, adjusts its send rate to $f(p)$. Function f is assumed to represent the loss-throughput relation that TCP would experience. When no loss occurs, the rate may also be increased according to some additional mechanism. We assume that the loss event interval estimator is non-biased. If the loss process is deterministic, the control is TCP-friendly in the long-run, i.e., the average throughput does not exceed that of TCP. If, in contrast, losses are random, it is a priori not clear whether this holds, due to the non-linearity of f , and a phenomenon similar to Feller's paradox. Our goal is to identify the key factors that drive whether, and how far, the control is TCP friendly (in the long run). As TCP and our source may experience different loss event intervals, we distinguish between TCP-friendliness and conservativeness (throughput does not exceed $f(p)$). We give a representation of the long term throughput, and derive that conservativeness is primarily influenced by various convexity properties of f , the variability of loss events, and the correlation structure of the loss process. In many cases, these factors lead to conservativeness, but we show reasonable experiments where the control is clearly non-conservative. However, our analysis also suggests that our source should experience a higher loss event ratio than TCP, which would make non-TCP friendliness less likely. Our findings provide guidelines that help understand when an equation base control is indeed TCP-friendly in the long-run, and in some cases, excessively so. The effects of round trip time and its variations are not included in this study.

Categories and Subject Descriptors

C.2 [Computer-Communication Networks]: Network Protocols;

C.4 [Performance of Systems]: Modeling techniques

General Terms

Algorithms, Performance, Theory, Verification

Keywords

Internet, Congestion control, Equation-based rate control, TCP-friendly, Point processes, Palm calculus

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'02, August 19-23, 2002, Pittsburgh, Pennsylvania, USA.

Copyright 2002 ACM 1-58113-570-X/02/0008 ...\$5.00.

1. INTRODUCTION

We consider an adaptive source that employs unicast equation-based rate control: the source estimates the loss event ratio p , and, primarily at loss events, adjusts its send rate to $f(p)$. Function f is assumed to represent the loss-throughput relation that TCP would experience. When no loss occurs, the rate may also be increased according to some additional mechanism, as we define below. An example of such an equation-based rate control is TFRC [10], which we use most of the time in this paper as a recurring example. Because f is assumed to represent TCP's loss throughput equation, it is expected that such a rate control is TCP friendly, i.e., our adaptive source shares the network fairly with competing TCP sources [11]. More precisely, this is required to happen on two time scales: in the short-run, response to congestion should be commensurate to that of TCP; in the long-run, average throughput should not exceed that of TCP. In this paper, we focus on the latter (we call it "long-run TCP-friendly"); we refer to [11, 7, 8] for some definitions regarding the concept of TCP-friendliness. Our goal is to identify the key factors that drive whether, and how far, the equation-based rate control is long-run TCP friendly.

We first point out that TCP and our source may experience different loss event ratios, and thus distinguish between TCP-friendly and conservative control. We say that the equation-based rate control is *conservative*, or *undershoots*, when $\bar{x}_A \leq f(\bar{p}_A)$, where $\bar{x}_A$ is the long-run time average of the send rate (= throughput) of our adaptive source and $\bar{p}_A$ its average loss event ratio. Note that we require a careful definition of loss rate, to avoid bias. Following [10] we assume that our source employs an unbiased estimator of the loss event interval, which is the amount of data sent in the interval between two successive loss events (see assumption (E) in Section 2). If there is convergence, then at the limit we would have $\bar{x}_A = f(\bar{p}_A)$. In practice, though, the control is required to be responsive, and thus the non-linearity of function f leaves little hope that $\bar{x}_A = f(\bar{p}_A)$. In fact, we do show in Section 3.4 that there are cases where the control undershoots, and others where it overshoots. Undershooting may appear to be a non-problem, however, we show in this paper that there are structural elements that lead to systematic, and sometimes excessive undershoot. In contrast, we find that overshooting may occur, but that there are fundamental reasons that limit its impact. Further, our analysis also shows that it is not only the non-linearity of f that plays a role, but also another phenomenon related to Feller's paradox¹.

Our findings are applied to the following choices of functions f (see also Section 2.4): the square root formula [12] (we call it SQRT); the loss-throughput formula (we call it PFTK-standard),

¹Feller's paradox [6] is that the average interval seen by a random observer is larger than the average interval sampled by an observer at interval boundaries.

and a slight variant (we call it PFTK-simplified). Other formulae exist that differ by a constant [3, 18]; our analysis applies directly to those as well. Yet other formulae are for short-lived TCP flows [5]; we did not include them in our analysis, since we focus on long-run behavior. Note however that most of our findings are independent of the specific function f being used.

We further distinguish between basic control, which consists in updating the send rate at every loss event, and comprehensive control, which adds a rate increase mechanism during intervals where no loss occurs. The comprehensive control reflects what is implemented in TFRC. We perform a detailed, exact analysis of the basic control. We find an approximate representation of comprehensive control and use it together with simulations. We find, numerically and by simulation, that the comprehensive control adds only a modest increase in long-run throughput to the basic control (remember that we focus here on long-run analysis; clearly, the comprehensive control may have a larger impact in the short-run).

For the basic control, we find fairly exhaustive results. First, conservativeness is strongly influenced by the convex or concave nature of two functionals of f , and the joint statistics of the loss event interval θ_n . If (C1) the statistics are such that θ_n and the loss event interval estimator $\hat{\theta}_n$ are lightly or negatively correlated (there are indications [20] that loss event intervals may be lightly correlated), then the control is conservative. Further, the higher the variability of the loss event estimator, the more conservative the control is; similarly, the higher the loss event ratio, the more conservative the control. Both of these effects are more pronounced with PFTK than with SQRT; with PFTK, this causes the control to be excessively conservative in regions of heavy loss. SQRT does not have this problem (but is also a less accurate representation of TCP's behavior). This conservative nature of TFRC control has been empirically observed in [9, 19, 2]. Second, if the correlation condition (C1) does not hold, then results may be radically different, and strongly depend on the nature of function f . We identified one useful case where we can conclude about non-conservativeness. If (C2c) the correlation of the duration between loss events and the send rate is non-negative, then for PFTK and heavy loss, the control systematically overshoots. For PFTK with low to medium loss, or for SQRT in all cases, this does not occur. This is due to convexity properties of some functionals of f , which holds differently in these cases. An example of protocol to which these assumptions may apply is an audio source with a constant packet rate, which adapts its data rate by varying the packet size [4]. These findings are exact for the basic control; for the comprehensive control, we pose them as claims and verify them by numerical and *ns* experiments.

To complete the analysis, we address two additional points. First, we compare, analytically and by simulation, the loss event ratio experienced by our adaptive source ($\bar{p}_A$) and by TCP ($\bar{p}_T$). We find that $\bar{p}_T \leq \bar{p}_A$ should be expected in most cases; this drives our source to have a smaller long-run throughput than TCP, beyond and above the effects mentioned above. Second, it remains to be seen whether the TCP equation used in a system is an accurate representation of TCP; we do not have an in-depth evaluation of this point; instead, we point to [3] for a discussion of this issue. However, we do find indications that there may be significant differences. This shows that designers of TCP-friendly protocols should clearly separate, in their evaluation, the three elements of conservativeness, TCP loss event ratio versus this protocol's loss ratio, and TCP's obedience to its formula. Failing to do so blurs the setting of parameters and may lead to undesired corrections. In contrast, knowing the expected deviations given in this paper leads to a safer understanding and tuning. Our results are based on an-

alytical findings (backed up by simulations) and should thus apply to a wide class of situations, including dynamic environments.

Limitations. We have identified some fundamental factors of equation-based rate control, by a mathematical analysis and experiments. Once the factors are identified, it remains to be seen what their values are in the real Internet. While existing measurements already provide some indication [20], this remains for further work. Also, throughout the paper, we assume that the round trip time (RTT) experienced by our source is constant, and known. It is outside the scope of this paper to address the impact on the control of variations of the RTT (the same method can be used). Last, we focus on the relationship between loss event ratio and throughput, taking both as observed quantities; stability and convergence are not addressed in this paper.

The paper is organized as follows. Section 2 describes our assumptions and notations. Section 3 gives our analytical findings. They are derived for the basic control, then verified numerically and by simulation of the control. We summarize our main findings in the form of two claims, given in Section 3.3. In Section 4 we validate our findings with *ns* simulations. In Section 5 we address the two additional points of TCP loss event ratio versus this protocol's loss event ratio, and TCP's obedience to its formula. All proofs are given in the appendix.

2. ADDITIONAL ASSUMPTIONS AND NOTATIONS

2.1 Notation

We consider an adaptive source with the send rate at time t equal to $X(t)$. We assume that $X(t)$ can be described by an ergodic process, and thus equate the long-run average with the expected value: $\bar{x} = E[X(0)]$. Index n refers to the n -th loss event. We use the following additional notation. T_n is the time at which the n -th loss event is detected by the source. $S_n = T_{n+1} - T_n$ is the elapsed time between two successive loss events. $X_n = X(T_n)$ is the rate set at the n -th loss event. $(\theta_n)_n$ is the number of packets sent between two successive loss events (θ_n packets in $[T_n, T_{n+1})$). Following TFRC, we call θ_n the loss event interval.

With E_T^0 we denote expectation with respect to the Palm probability P_T^0 , which is, roughly speaking, the probability given there is a loss event at time 0 (see for instance [1] for a formal definition). Let

$$\text{cov}_T^0[X(0), Y(0)] := E_T^0[X(0)Y(0)] - E_T^0[X(0)]E_T^0[Y(0)],$$

be covariance of X and Y with respect to P_T^0 . Likewise, let

$$\text{var}_T^0[X(0)] := E_T^0[X(0)^2] - E_T^0[X(0)]^2,$$

be variance of X with respect to P_T^0 .

The long-run loss event ratio as observed by the source is

$$\bar{p} = \frac{1}{E_T^0[\theta_0]}. \quad (1)$$

Let $\hat{\theta}_n$ be the loss event interval estimator, computed at T_n . We assume

(E) $\hat{\theta}_n$ is an unbiased estimator² of $1/\bar{p}$.

Moreover, we assume that $\hat{\theta}_n$ is defined as a moving-average of the

²Note that, by Jensen's inequality $E_T^0[\frac{1}{\theta_0}] \geq \bar{p}$, and thus $1/\hat{\theta}_n$ is a biased estimator of $\bar{p}$.

Figure 1: Some functionals of interest for our functions f : SQRT, PFTK-standard, and PFTK-simplified (resp. labeled as SQRT, PFTK and PFTK’); $r = 100$ ms, $q = 4r$. Values of x close to 0 correspond to heavy losses. The top figure tends to indicate that the convexity condition (F1) in Theorem 1 would be satisfied in all three cases, but this is strictly true only for SQRT and PFTK-simplified; it also illustrates that convexity is much more pronounced for PFTK-simplified than for SQRT. The bottom figure illustrates that the concavity condition (F2) of Theorem 2 is true for SQRT; for PFTK-standard and PFTK-simplified it holds only for small loss event ratios; for heavy loss (x small), the curves are convex and thus the opposite condition (F2c) holds.

loss event intervals:

$$\hat{\theta}_n = \sum_{l=1}^L w_l \theta_{n-l}, \quad (2)$$

for some positive valued weights $(w_l)_{l=1}^L$; note by (E) we require $\sum_{l=1}^L w_l = 1$. TFRC uses this type of loss event interval estimator, for a particular setting of the weights such that w_l are equal for $1 \leq l \leq L/2$, and otherwise w_l linearly decrease with l .

2.2 Basic Control

The basic control is defined as follows. For $t \in [T_n, T_{n+1})$,

$$X(t) = f\left(\frac{1}{\hat{\theta}_n}\right). \quad (3)$$

Function f is the loss-throughput formula assumed to be positive-valued and non-increasing.

2.3 Comprehensive Control

Here we add an additional mechanism to the basic control in (3), and call the resulting system the comprehensive control. The mechanism reflects a response to positive feedback as found in TFRC [10].

Let $\theta(t)$ be the number of packets sent since the most recent loss event that was observed before t . Then we define the comprehensive control as follows, for $t \in [T_n, T_{n+1})$,

$$X(t) = f\left(\frac{1}{\hat{\theta}(t)}\right), \quad (4)$$

$$\hat{\theta}(t) = (w_1 \theta(t) + \sum_{l=1}^{L-1} w_{l+1} \theta_{n-l}) 1_{A_t} + \hat{\theta}_n 1_{\bar{A}_t}.$$

Where

$$A_t = \left\{ \theta(t) > \frac{1}{w_1} \left[\hat{\theta}_n - \sum_{l=1}^{L-1} w_{l+1} \theta_{n-l} \right] \right\},$$

and $1_{A_t} = 1$ if A_t is true, else $1_{A_t} = 0$.

In other words, at time t , the loss event interval estimator $\hat{\theta}(t)$ is updated with $\theta(t)$, if that increases the value of the estimator. If this is not the case, then $\hat{\theta}(t)$ is kept to $\hat{\theta}_n$. Note that once the condition A_t is true ($\theta(t)$ sufficiently large), the control (4) responds to the positive feedback by increasing the send rate.

Notice that the send rate dynamics is such that, if $\hat{\theta}_{n+1} \leq \hat{\theta}_n$, then $X(t) = f(1/\hat{\theta}_n)$ for all $t \in [T_n, T_{n+1})$. Else, for $\hat{\theta}_{n+1} > \hat{\theta}_n$ the send rate is $X(t) = f(1/\hat{\theta}_n)$, for $t \in [T_n, U_n]$, and then the rate increases according to (4) for $t \in (U_n, T_{n+1})$. Here, from the definition of A_t ,

$$U_n = \frac{1}{w_1 f\left(\frac{1}{\hat{\theta}_n}\right)} \left[\hat{\theta}_n - \sum_{l=1}^{L-1} w_{l+1} \theta_{n-l} \right].$$

2.4 Functions f Used in This Paper

We use the following loss throughput formulae. We first consider perhaps the simplest one, “the square-root”, which we call SQRT [12]:

$$f(p) = \frac{1}{c_1 r \sqrt{p}}, \quad (5)$$

where c_1 is some constant, and r the average round-trip time.

We refer to another well-known function f (Eq. (30) in [13]) as PFTK-standard:

$$f(p) = \frac{1}{c_1 r \sqrt{p} + q \min[1, c_2 \sqrt{p}](p + 32p^3)}, \quad (6)$$

for some constant c_2 ; here q is TCP retransmit timeout. A variant of the above formula is referred to in TFRC specification [10]). We call it PFTK-simplified.

$$f(p) = \frac{1}{c_1 r \sqrt{p} + q c_2 (p^{3/2} + 32p^{7/2})}. \quad (7)$$

Note that (7) is equal to (6) for $p \leq \frac{1}{c_2^2}$, and otherwise it is less.

Here $c_1 = \sqrt{2b/3}$ and $c_2 = 3/2\sqrt{3b/2}$, where b is the number of packets acknowledged by a single ack; typically $b = 1$ [10].

Also note that most of our findings apply to other functions f as well.

3. WHAT MAKES THE CONTROL CONSERVATIVE OR NOT

We first give the core mathematical expressions that are used to compute, in theory and in numerical experiments, the throughput of our control.

3.1 Representation of Throughput

PROPOSITION 1. *The throughput of the basic control (3) is*

$$E[X(0)] = \frac{E_T^0[\theta_0]}{E_T^0\left[\frac{\theta_0}{f\left(\frac{1}{\theta_0}\right)}\right]}. \quad (8)$$

For the comprehensive control, we do not have a closed form expression. However, for both PFTK-simplified and SQRT we have the following approximation, which is an upper bound.

PROPOSITION 2. *For PFTK-simplified and SQRT, the throughput of the comprehensive control (4) is approximated by the following upper bound*

$$E[X(0)] \leq \frac{E_T^0[\theta_0]}{E_T^0[\frac{\theta_0}{f(\frac{1}{\theta_0})}] - E_T^0[V_0 1_{\hat{\theta}_1 > \hat{\theta}_0}]}, \quad (9)$$

where

$$V_n = \frac{1}{w_1} \left[-2c_1 r (\hat{\theta}_{n+1}^{-\frac{1}{2}} - \hat{\theta}_n^{-\frac{1}{2}}) + 2c_2 q (\hat{\theta}_{n+1}^{-\frac{1}{2}} - \hat{\theta}_n^{-\frac{1}{2}}) - \frac{2}{5} c_3 q (\hat{\theta}_{n+1}^{-\frac{5}{2}} - \hat{\theta}_n^{-\frac{5}{2}}) + (\hat{\theta}_{n+1} - \hat{\theta}_n) \frac{1}{f(1/\theta_n)} \right],$$

and $c_3 := 32c_2$ (for SQRT, let $c_2 = c_3 = 0$).

Note that, in view of (2), the throughput of both basic and comprehensive control is expressed in terms of the expected values of some functions of the loss interval intervals $\theta_n, \theta_{n-1}, \dots, \theta_{n-L}$. Thus knowing the joint probability law of $\theta_n, \theta_{n-1}, \dots, \theta_{n-L}$ would, at least in theory, enable one to compute the throughput, and explain how the correlation structure of the loss process plays a role.

3.2 Conditions for the Basic Control to be Conservative

We temporarily focus on the basic control. We give exact sufficient conditions for conservativeness, or non-conservativeness. The results have interest of their own, and they also suggest the key factors that drive conservativeness.

3.2.1 A Sufficient Condition for the Basic Control to be Conservative

THEOREM 1. *Assume that*

(F1) *the loss-throughput formula f is such that $\frac{1}{f(1/x)}$ is convex with x ,*

(C1) $\text{cov}_T^0[\theta_0, \hat{\theta}_0] \leq 0$.

Then the basic control (3) is conservative.

Interpretation. The convexity condition (F1) is satisfied by the SQRT loss throughput formula, and by PFTK-simplified; it is not satisfied by PFTK-standard, but almost (we will come back to this in a few lines). This is straightforward to demonstrate, and can also be seen on Figure 1. The figure also shows that convexity is much more pronounced for PFTK formulae, and thus, we should expect more conservativeness with PFTK than with the square root formula (this is confirmed numerically in Section 1).

Condition (C1) is true in particular when the covariance is 0, which happens when successive loss event intervals are (stochastically) independent. There are indications in [20] that this may be true, and the theorem says that this would lead to a conservative behavior. We show in appendix the following more explicit statement, which gives a bound on long-run throughput:

$$E[X(0)] \leq f(\bar{p}) \frac{1}{1 + \frac{f'(\bar{p})\bar{p}^3}{f(\bar{p})} \text{cov}_T^0[\theta_0, \hat{\theta}_0]}. \quad (10)$$

This shows that, in most cases, if the covariance is positive but small, there cannot be any significant non-conservativeness of the basic protocol.

Figure 2: The top figure shows $g(x) := 1/f(1/x)$ when $f(\cdot)$ is PFTK-standard and its convex closure (dotted line). On the interval shown in the top figure, g^{**} is equal to the tangent common to both ends of the graph. Outside the interval it is equal to g . $g(\cdot)$ is not strictly speaking convex, but almost. The bottom figure shows the ratio g/g^{**} , which is bounded by $r = 1.0026$.

The theorem says more. Remember that $\hat{\theta}_n$ is an incremental estimator of the loss event interval $1/\bar{p}$, built on the information available up to the loss event $n - 1$, while θ_n is the true next loss event interval. Both have the same expectation, as we assumed that $\hat{\theta}_n$ is unbiased. However, this does not mean that $\hat{\theta}_n$ is a good predictor of θ_n . This depends on the joint statistics, in particular the autocovariance, of the loss process. The covariance of θ_n and $\hat{\theta}_n$ reflects how good a predictor $\hat{\theta}_n$ is. Condition (C1) means that $\hat{\theta}_n$ is a bad predictor, and, maybe surprisingly, the theorem suggests that this leads to a conservative behavior. Conversely, consider now a hypothetical case where the loss process goes into phases, with slow transitions. Then the loss event interval becomes highly predictable; the theorem does not say that this alone will make the control non-conservative. However, this may really happen, as we find in Section 3.4. We give another, perhaps more realistic example in Section 3.2.2.

Note that $\hat{\theta}_n$ is the moving-average estimator in (2), and thus

$$\text{cov}_T^0[\theta_0, \hat{\theta}_0] = \sum_{l=1}^L w_l \text{cov}_T^0[\theta_0, \theta_{-l}], \quad (11)$$

in other words, it depends only on the spectral properties of the loss event interval.

The following corollary was shown in the discussion above.

COROLLARY 1. *If the convexity condition (F1) holds and the loss event intervals are independent then the basic control (3) is conservative.*

When Convexity is Almost True. The convexity condition (F1) is not true for PFTK-standard (because of the min term), but almost, as we see now. For a function $g(x)$, we quantify its deviation from convexity by the ratio to its convex closure

$$r = \sup_x \left[\frac{g(x)}{g^{**}(x)} \right].$$

The convex closure $g^{**}(x)$ is the largest convex function that lower bounds $g(x)$; it is obtained by applying convex conjugation twice [16]. Fig. 2 shows $g(x) = 1/f(1/x)$ for PFTK-standard and its convex closure; here, we have $r = 1.0026$.

PROPOSITION 3. *Assume that the loss-throughput formula f is such that $\frac{1}{f(1/x)}$ deviates from convexity by a ratio r , and that (C1) holds. Then the basic control (3) cannot overshoot by more than a factor equal to r .*

Thus, considering that a fraction of a percent is more than reasonable accuracy, we can conclude that for practical purposes, we can act as if PFTK-standard would satisfy the convexity condition F1.

3.2.2 When the sufficient conditions do not hold

We give a different set of conditions, which provide additional insights. The first of these sets was found, in a restricted form, in our previous work [17]. The second set applies to the cases where Theorem 1 does not apply.

THEOREM 2. *Assume that*

(F2) *the loss-throughput formula f is such that $f(1/x)$ is concave with x ,*

(C2) $\text{cov}_T^0[X_0, S_0] \leq 0$.

Then the basic control (3) is conservative.

Conversely, if

(F2c) *the loss-throughput formula f is such that $f(1/x)$ is strictly convex with x ,*

(C2c) $\text{cov}_T^0[X_0, S_0] \geq 0$,

(V) *the loss event estimator $\hat{\theta}_n$ has non-zero variance.*

Then the basic control (3) is non-conservative.

Interpretation. The concavity condition (F2) is true for the square root formula. In contrast, PFTK-standard and PFTK-simplified are such that concavity (F2) is true for rare losses, but convexity (F2c) is true for frequent losses (see Figure 1, lower graph). The covariance condition (C2) is between X_n , the rate set at the occurrence of the n -th loss event, and S_n , the time until the next loss event. If the loss process is memoryless and independent of the activity of our source, then the duration S_n of the loss interval is negatively correlated with the send rate X_n in the given interval (since S_n is counted in real time, not per packet); in such cases, condition (C2) is true, and the basic control is conservative as long as losses are rare to moderate (or if the SQRT formula is used). This part of Theorem 2 complements Theorem 1.

Consider now the second part of Theorem 2. Assume that $(S_n)_n$, the sequence of loss event intervals counted in real time, is independent of the send rate. This may happen for example for an audio source that modulates its send rate by varying the packet size rather than the packet send rate, and if the packet dropping probability in RED routers is independent of packet size. Then (C2c) holds, with equality. Now assume also that PFTK-standard is used, and the network setting happens to be such that the loss event interval θ_n is mostly in the region where PFTK-standard is convex (i.e. heavy losses). The theorem says that the basic control is non-conservative, except in the degenerate case where there is no randomness in the system, i.e. the loss estimator has converged to a constant value. We show simulations that illustrate this case in Section 4.

Another example is for a more traditional source such as TFRC, but when the loss process goes through phases (for example, the network paths used by the flow oscillate between congestion and no congestion), and the sending rate roughly follow the phases (i.e., is responsive at the time scale of the loss process). Then when the network is in the congestion phase, X_n is most often small, and because of congestion, S_n is small. In such a case, condition (C2c) may be true and the basic control may *not* be conservative. In Section 3.4 we show such cases.

Comments. From a methodology viewpoint, the first part of Theorem 2 illustrates well the importance of the Feller paradox-type of the arguments used in this paper. The send rate $X(t)$ is updated only at the loss events times $(T_n)_n$. Consider an observer who picks up a point in time at random; she is more likely to fall in a large loss event time S_n . Given that S_n is negatively correlated with X_n , it is thus more likely that on average she will observe a smaller rate than another observer that would sample the rate at $(T_n)_n$. From this we conclude $E[X(0)] \leq E_T^0[X_0]$ (The concavity assumption (F2), by Jensen's inequality, shows in turn that $E_T^0[X_0] \leq f(\bar{p})$, thus finally $E[X(0)] \leq f(\bar{p})$ and the control is conservative).

The main result in our previous work [17] is similar to the first part of Theorem 2, but with the correlation condition (C2) replaced by the condition that the expected duration S_n , conditional to the rate X_n , decreases with X_n :

(C3) $E_T^0[S_0|X_0 = x]$ is non-increasing with x .

It is a direct consequence of Harris' inequality³ (see for example [1], p. 225) that (C3) implies the negative correlation condition (C2); thus our previous result is a special case of the first part of Theorem 2.

Of course, we should expect that the combination of (C2c) and (V) implies that (C1) does not hold. This indeed holds and is shown in the appendix.

It is legitimate to wonder whether Theorem 1 is derived from Theorem 2 or vice versa. It does not seem to be the case (we discuss this in the appendix). Note however that if the concavity condition (F2) holds, then the convexity condition (F1) necessarily also holds (but the converse is not true).

3.3 What This Tells Us

The analytical results in the previous section are for the basic control. We expect the comprehensive control to give a slightly higher throughput, since it differs by an additional increase during long loss event intervals. This motivates us to pose as assumptions the following analysis, which we will confirm later in the paper by numerical experiments and *ns* simulations.

CLAIM 1. *Assume that the loss event interval θ_n and the loss event interval estimator $\hat{\theta}_n$ are lightly or negatively correlated. Consider the region where the loss event interval estimator $\hat{\theta}_n$ takes its values. The more convex $\frac{1}{f(1/x)}$ is in this region, the more conservative the control is. The more variable $\hat{\theta}_n$ is, the more conservative the control is.*

Application. For protocols like TFRC, we expect the condition to hold in many practical cases [20]. For the three functions we consider in this paper, $\frac{1}{f(1/x)}$ is more convex for small x , that is, for large loss event ratios $\bar{p}$. Thus, the control should be more conservative with heavy loss than low loss. This effect is more pronounced for PFTK-standard (6) and PFTK-simplified (7), which

³Harris' inequality says that if $f(x)$ and $g(x)$ are non-decreasing functions, and X is one random variable, then the covariance of $f(X)$ and $g(X)$ is non-negative.

are convex and very steep for large $\bar{p}$, than for SQRT. This explains the observed drop in throughput for the control, with PFTK and heavy losses.

The variability of $\hat{\theta}_n$ depends on the variability of $(\theta_n)_n$, and can be controlled by the length of the moving-average estimator $\hat{\theta}_n$ (2). With some proper setting of the weights $(w_l)_{l=1}^L$, the larger the length of the estimator L , the smaller the variability of the estimator $\hat{\theta}_n$ (for instance, for finite-variance $(\theta_n)_n$, and uniform weights $w_l = 1/L$, $l = 1, 2, \dots, L$, variance of $\hat{\theta}_n$ scales as $O(1/L)$). We should find that for larger L the control becomes less conservative.

The second claim concerns a case where the conditions in Claim 1 do not hold.

CLAIM 2.

- Assume that duration in real time of the loss event interval S_n and the send rate X_n are negatively or non correlated. If $f(1/x)$ is concave in the region where the loss event interval estimator $\hat{\theta}_n$ takes its values, the control tends to be conservative.
- Conversely, if S_n and X_n are positively or non correlated, and if $f(1/x)$ is strictly convex in the region where the loss event interval estimator $\hat{\theta}_n$ takes its values, and $(\theta_n)_n$ is not fixed to some constant, the control is non-conservative.

In both cases, the more variable $\hat{\theta}_n$ is, the more pronounced the effect is.

Application. We expect to have a close to zero correlation for adaptive audio applications such as [4] when packet losses in RED routers are independent of packet size. Thus, depending on which convexity condition holds, we will find one or the other outcome. For SQRT, the control should always be conservative. The same holds for PFTK with light to moderate losses. The *opposite* holds for either PFTK formulae with heavy losses (loss event ratio larger than 0.1).

3.4 Numerical Examples

We now support some of the observations we made from our analytical study by numerical examples. Such a numerical study enables us to isolate individual factors that we expect to contribute to either conservative or non-conservative behavior. We show later in Section 4 the results of *ns* simulations.

All results in this section are based on numerical investigations of the basic control and the comprehensive control, with functions SQRT or PFTK-simplified. For PFTK-standard, we rely on *ns* simulations shown in Section 4; in view of the claims, the results do not differ significantly.

3.5 Validation of Claim 1

We consider independent and identically distributed (i.i.d) $(\theta_n)_n$ with marginal density function $\mu(x) = \lambda \exp(-\lambda(x - x_0))$, for $x \geq x_0$, and $\lambda, x_0 \geq 0$; μ is known as generalized exponential distribution.

We explain why we choose the density function μ . μ has some nice properties: $E_T^0[\theta_0] = x_0 + \lambda$, $cv_T^0[\theta_0] = \frac{\lambda}{\sqrt{\lambda + x_0}}$, $S_T^0[\theta_0] = 2$, and $K_T^0[\theta_0] = 6$. Here

$$cv_T^0[\theta_0]^2 = \frac{\text{var}_T^0[\theta_0]}{E_T^0[\theta_0]^2},$$

is the coefficient of variation; and

$$S_T^0[\theta_0] = \frac{E_T^0[(\theta_0 - E_T^0[\theta_0])^3]}{\text{var}_T^0[\theta_0]^{3/2}}, \quad K_T^0[\theta_0] = \frac{E_T^0[(\theta_0 - E_T^0[\theta_0])^4]}{\text{var}_T^0[\theta_0]^2},$$

Figure 3: Normalized throughput $E[X(0)]/f(\bar{p})$ versus $\bar{p}$ for the basic control; $cv_T^0[\theta_0] = 1 - 10^{-3}$; SQRT (Top) and PFTK-simplified with $q = 4r$ (Bottom). The estimator weights are as with TFRC of length L .

are the skewness and kurtosis parameters, respectively.⁴ We note that μ gives us a freedom to vary either $E_T^0[\theta_0]$ or $cv_T^0[\theta_0]$, while the other of these two parameters is kept fixed. At the same time, skewness and kurtosis parameters remain unchanged. Thus μ enables us to separate the effects due to convexity of $1/f(1/x)$ and variability of $\hat{\theta}_n$. With some other distributions, for instance, the geometrical distribution $P(\theta_0 = k) = \bar{p}(1 - \bar{p})^{k-1}$, $k \geq 1$, we would have $cv_T^0[\theta_0] = \sqrt{1 - \bar{p}}$. In this case, the variability of $(\hat{\theta}_n)_n$ would decrease as we increase $\bar{p}$. For some other discrete-valued distributions we may be able to fix $cv_T^0[\theta_0]$, but not at the same time the kurtosis parameter that reflects the peakedness of the distribution, and thus the variability of $(\hat{\theta}_n)_n$.

We compute the throughput $E[X(0)]$ numerically for the basic and comprehensive control from Equation (8) and (9), respectively. The results are obtained by Monte Carlo simulation; out of 5 independent simulations runs each with 10000 samples, and 0.95 confidence intervals.

Our first objective is to evaluate the impact of convexity of the function $1/f(1/x)$. To that end, we fix $cv_T^0[\theta_0] = 1 - 10^{-3}$. In Fig. 3 we show the normalized throughput $\bar{x} = E[X(0)]/f(\bar{p})$ versus $\bar{p}$ for the basic control with SQRT and PFTK-simplified functions f . The values $\bar{x} \leq 1$ correspond to the conservative behavior (resp. $\bar{x} > 1$ to non-conservative behavior). Ideally, we would have $\bar{x} = 1$. For SQRT function, we observe, for each fixed value of length L of the estimator $\hat{\theta}_n$, $\bar{x}$ is around the same value irrespective of $\bar{p}$. This constancy of $\bar{x}$ with respect to $\bar{p}$ is to be expected. For exponentially distributed θ_n and uniform weights $w_l = 1/L$,

⁴Skewness and kurtosis parameters quantify skewness and sharpness of a probability distribution.

Figure 4: Normalized throughput $E[X(0)]/f(\bar{p})$ versus $\bar{p}$ for the comprehensive control; $cv_T^0[\theta_0] = 1 - 10^{-3}$; SQRT (Top) and PFTK-simplified with $q = 4r$ (Bottom). The estimator weights are as with TFRC of length L .

$l = 1, 2, \dots, L$, a simple calculation reveals

$$\bar{x} = \frac{(L-1)!}{\sqrt{L}\Gamma(L-\frac{1}{2})},$$

which does not depend on $\bar{p}$; here Γ is the gamma function. We also come to the same conclusion by computing $\bar{x}$ for TFRC weights, which is lengthy, and thus not shown here.

On the other hand, for PFTK-simplified function, we observe that $\bar{x}$ decreases towards 0 as we increase $\bar{p}$. This explains the known throughput drop for heavy losses.

In Fig. 4 we show the corresponding results for the comprehensive control. The results are qualitatively the same as the respective results for the basic control shown in Fig. 3. For SQRT function, the normalized throughputs are less, but fairly close, to the ideal value 1. For PFTK-simplified function, the results are somewhat less conservative than for the basic control.

Next we investigate the impact of the variability of $\hat{\theta}_n$. To that end, we consider the normalized throughput $\bar{x}$ as we vary the coefficient of variation of θ_n , while we keep the equilibrium point fixed at $\bar{p} = 0.01$ and 0.1 ; see Fig. 5. We show numerical results only for the basic control with PFTK-simplified formula. We observe that the larger the variability of $\hat{\theta}_n$, the more conservative the control is. This is indeed more pronounced for larger $\bar{p}$ due to the larger convexity and steepness of $1/f(1/x)$ for small x (large $\bar{p}$) with PFTK-simplified function.

Also observe how the throughput depends on L , the window size used for the estimator $\hat{\theta}_n$, in Fig. 3, Fig. 4 and Fig. 5. A large L reduces the variability of $\hat{\theta}_n$ and thus increases the throughput, as predicted in Claim 1.

Lastly, we briefly mention the results obtained for θ_n geometrically distributed with parameter $\bar{p}$ (not shown in the paper). In this case, the same qualitative statements also hold, but with a slight

Figure 5: Normalized throughput $E[X(0)]/f(\bar{p})$ versus the coefficient of variation of $(\theta_n)_n$ for the basic control; $\bar{p} = 0.01$ (Top) and $\bar{p} = 0.1$ (Bottom). The function f is PFTK-simplified with $q = 4r$. The estimator weights set as TFRC weights.

trend to non-conservative direction for large $\bar{p}$ due to the reason explained earlier.

3.6 Validation of Claim 2

We do additional experiments to verify Claim 2, which, incidentally, also provide some examples of non-conservative behavior. Consider $(\theta_n)_n$ that we define as follows. We suppose there exists a hidden Markov chain (HMC) that governs the loss events. We define this HMC $(Z_n)_n$ to be discrete time clocked at the loss events. Assume $(Z_n)_n$ takes values on a finite state space E ; call $\mathbf{P} = [p_{ij}]$ the matrix of transition probabilities, and π its stationary probability. Take as a modeling assumption:

$$P(\theta_n = m | Z_n = i, Z_k, \theta_k, k < n) = P(\theta_n = m | Z_n = i).$$

In other words, given that, at the n -th loss event, the HMC is in the state i , θ_n is independent of anything, but Z_n . Let $g_i(m) := P(\theta_n = m | Z_n = i)$. Notice, $(Z_n)_n$ is a semi-Markov process with $P(Z_{n+1} = j, \theta_n = m | Z_n = i) = p_{ij}g_i(m)$.

For the basic control, from (8),

$$E[X(0)] = \frac{\sum_{i \in E} e(i_0)\pi_i}{\sum_{i \in E^{L+1}} e(i_0)g(i_1, \dots, i_L)p_{i_0 i_1} \dots p_{i_{L-1} i_L} \pi_{i_L}},$$

where

$$g(i_1, \dots, i_L) = E_T^0 \left[\frac{1}{f(1/\hat{\theta}_0)} | Z_{-1} = i_1, \dots, Z_{-L} = i_L \right],$$

and $e(i) = E_T^0[\theta_0 | Z_n = i]$. Likewise, one obtains the throughput expression for the comprehensive control.

We next consider a simple, but instructive case: the 2-state HMC with $L = 1$. Without loss of generality, we call one state *good*, and other *bad*; we label the state space as $E = \{g, b\}$. Moreover, we

Figure 6: The graph shows contour plot of $\text{cov}_T^0[X_0, S_0]$ versus p_{gb} and p_{bg} ; $n_g = 200$ and $n_b = 50$. The function f is PFTK-simplified with $r = 100$ ms and $q = 4r$.

assume there are two fixed $n_g \geq n_b$ such that $g_a(n_a) = 1$, and $g_b(n_b) = 1$. In other words, when the HMC is in the good state (resp. bad), then the loss event interval is fixed to n_g (resp. n_b).

Under the above assumptions, we obtain:

$$E[X(0)] = \frac{p_{bg}n_g + p_{gb}n_b}{p_{bg} \frac{n_g}{f(1/n_g)} + p_{gb} \frac{n_b}{f(1/n_b)} + p_{gb}p_{bg}h(n_g, n_b)},$$

where, for the basic control,

$$h(n_g, n_b) = \left(\frac{1}{f(1/n_b)} - \frac{1}{f(1/n_g)} \right) (n_g - n_b),$$

and, for the comprehensive control,

$$h(n_g, n_b) = 2c_1r(n_g^{\frac{1}{2}} - n_b^{\frac{1}{2}}) - 2c_2q(n_g^{-\frac{1}{2}} - n_b^{-\frac{1}{2}}) - \frac{2}{5}c_3q(n_g^{-\frac{5}{2}} - n_b^{-\frac{5}{2}}) - \frac{n_g - n_b}{f(1/n_g)}.$$

We next discuss the covariance $\text{cov}_T^0[X_0, S_0]$ of our 2-state HMC. Notice that X_n and S_n being negatively correlated or non-correlated is equivalent to $\text{cov}_T^0[X_0, S_0] \leq 0$. We show a plot of $\text{cov}_T^0[X_0, S_0]$ versus the transition probabilities p_{gb} and p_{bg} in Fig. 6. Observe that the covariance is positive for small values of p_{gb} and p_{bg} , which corresponds to the slow dynamics of our 2-state HMC. It is simple to show, and perhaps instructive to note, that for the slow HMC limit, and f being SQRT function, we would have $\text{cov}_T^0[X_0, S_0] \rightarrow \text{var}_T^0[\sqrt{\theta_0}]$; thus a positive value, increasing in the variability of $(\theta_n)_n$. In view of our Claim 2, we expect to find non-conservative behavior when the dynamics of the HMC is slow, which we do confirm next.

We first consider the basic control with PFTK-simplified formula. In Fig. 7, we show the normalized throughput $E[X(0)]/f(\bar{p})$ versus the transition probabilities p_{gb} and p_{bg} of our HMC. n_g and

Figure 7: Normalized throughput $E[X(0)]/f(\bar{p})$ versus p_{gb} and p_{bg} for the basic control; $n_g = 200$ and $n_b = 50$. The function f is PFTK-simplified with $r = 100$ ms and $q = 4r$.

n_b are set to 200 and 50, which correspond to the loss event ratios 0.005 and 0.02, while in a good and bad state, respectively. In Fig. 8 we show numerical results for the comprehensive control obtained under the same setting. Note that we do find some slight overshoot in the lower left corner of the figures (normalized throughput greater than unity).

Note that for the given values of n_g and n_b the function $f(1/x)$ is concave with x in the region where x takes its values. Further, observe from Fig. 6 and Fig. 7 (resp. Fig. 8) that whenever $\text{cov}_T^0[X_0, S_0]$ is not positive, the control is conservative. The last two observations together confirm the first statement of Claim 2. The second statement of Claim 2 we do not verify here numerically, but by *ns* simulation in Section 4. Further numerical examples, with another model, that support Claim 2 can be found in [17].

We give some further observations. By Corollary 1 we should find the conservative behavior for $p_{gb} + p_{bg} = 1$ (note that this is a degenerate case such that $(\theta_n)_n$ is i.i.d.), which we confirm to be the case. We note that very conservative behavior occurs for $p_{gb} + p_{bg} > 1$, where $\text{cov}_T^0[X_0, S_0]$ is negative, but also $\text{cov}_T^0[\theta_0, \hat{\theta}_0]$ may be negative; this is to be expected from the bound on throughput in Equation (10). Another observation is that the non-conservative behavior happens for positively correlated $(\theta_n)_n$ ($p_{gb} + p_{bg} \leq 1$), in particular, for small values of p_{gb} and p_{bg} (slow dynamics of the HMC). In the remainder of this section, we discuss this limit case in some more detail.

We show that for the slow dynamics of the HMC the control may have a substantial overshoot, as opposed to a modest overshoot observed in Fig. 7 and Fig. 8. We define the slow HMC limit as $p_{gb}, p_{bg} \rightarrow 0$, and $p_{gb} = up_{bg}$, for some fixed $u > 0$. Then, for both the basic and comprehensive control we obtain:

Figure 8: Normalized throughput $E[X(0)]/f(\bar{p})$ versus p_{gb} and p_{bg} for the comprehensive control; $n_g = 200$ and $n_b = 50$. The function f is PFTK-simplified with $r = 100$ ms and $q = 4r$.

$$E[X(0)] \rightarrow \frac{p_{bg}n_g + p_{gb}n_b}{p_{bg} \frac{n_g}{f(1/n_g)} + p_{gb} \frac{n_b}{f(1/n_b)}} = \frac{E_T^0[\theta_0]}{E_T^0[\frac{\theta_0}{f(1/\theta_0)}]}.$$

Thus, $E[X(0)]/f(\bar{p})$ for our slow HMC limit is:

$$\bar{x}^0(u) = \frac{1}{f(\frac{u+1}{un_g+n_b})} \frac{un_g + n_b}{u \frac{n_g}{f(1/n_g)} + \frac{n_b}{f(1/n_b)}}. \quad (12)$$

For a given function f , one may compute u^* at which the global maximum of $\bar{x}^0(u)$ is attained. For SQRT function, the calculation is tractable and we obtain $u^* = \sqrt{n_b/n_g}$. That is

$$\frac{p_{bg}}{p_{gb}} = \sqrt{\frac{n_b}{n_g}}.$$

Notice that, the larger the discrepancy of the good and bad states (larger n_g/n_b), the smaller the relative number of transitions from the bad to good state. The last implies that the HMC is most of the time in the bad state, with occasional short excursions to the good state. It is for this dynamics of our HMC, when we may expect significant non-conservative behavior (overshoot).

For SQRT function f , the maximum value of $\bar{x}^0$ is

$$\bar{x}^* = \frac{1}{2} \sqrt{2 + \sqrt{\frac{n_g}{n_b}} + \frac{1}{\sqrt{\frac{n_g}{n_b}}}}. \quad (13)$$

which we note is monotonically increasing with n_g/n_b .

We show in Fig. 9 numerical values of $\bar{x}^*$ (13) versus the ratio n_g/n_b , which we recall is for SQRT function. We also show the results for PFTK-simplified function by numerical computation of the maximum of (12). We observe that for sufficiently large values of n_g/n_b we may have a substantial non-conservative behavior.

Figure 9: Maximum normalized throughput $E[X(0)]/f(\bar{p})$ attained in the slow HMC limit versus n_g/n_b ; thick line is for SQRT; thin lines are for PFTK-simplified ($r = 100$ ms, $q = 4r$); n_b is set as indicated in the graph.

4. VALIDATION BY SIMULATION

We conduct *ns* simulation experiments to validate the claims made in Section 3.3. Unless otherwise indicated, we consider a link shared by TFRC and TCP Sack1 connections. The link implements RED queue management of the rate of 15 Mb/s; we set the buffer length, min_thresh, and max_thresh to 2.5, 0.25 and 1.25 times the bandwidth delay product, respectively. The round-trip time is about 50 ms. We mimic this setting from [2].

4.1 Validation of Claim 1

In Fig. 9 we show the normalized throughput for PFTK-standard formula. We verify, the larger the loss event ratio is, the more conservative the control is. We also note that the larger the smoothing of the loss event interval estimator (larger L) is, the less conservative the control is. Next, for PFTK-simplified (Fig. 10) we observe the results are very close to those with PFTK-standard. We verify in Fig. 11, the conservativeness with SQRT formula is less pronounced, and less dependent on L . In all the cases, covariance of the instantaneous value and the estimator of the loss event interval is small; indicating low autocorrelation of the loss event intervals.

4.2 Validation of Claim 2

We consider a source that sends packets at regular time intervals (20 ms), but controls packet sizes. The source has a connection established through a loss module that allows us to tune the packet drop rate. For such a source, we have the covariance of the send rate and the interval between two loss events equal to zero. Thus, by Claim 2 we expect our source to be conservative for $f(1/x)$ concave with x ; conversely, non-conservative for $f(1/x)$ convex with x . We show the results for two lengths of the loss event interval estimator $L = 4$ and 8 (resp. Fig. 12 and 13). We verify,

Figure 10: Function f is PFTK-standard. The upper graph shows the normalized throughput $E[X(0)]/f(\bar{p})$ attained by TFRC versus the loss event ratio $\bar{p}$. The lower graph shows $\text{cov}_T^0[\hat{\theta}_0, \theta_0]$ normalized by $E_T^0[\theta_0]E_T^0[\hat{\theta}_0]$.

the control with SQRT is always conservative. For PFTK-standard and PFTK-simplified the same holds for a low loss rate, however, for high loss rate the functions are convex, and thus the control exhibits non-conservative behavior in this region. Observe from Fig. 12 and Fig. 13, as the loss event ratio increases, the coefficient of variation of $\hat{\theta}_n$ gets smaller. Smaller variability of the loss event estimator makes the control to be either less conservative or less non-conservative, depending on which behavior is in action. Contrary, larger variability of $\hat{\theta}_n$ exaggerates either conservative or non-conservative behavior.

5. CONSERVATIVE DOES NOT MEAN TCP-FRIENDLY

We have focused so far on conservativeness, i.e., whether an adaptive source does indeed satisfy its equation in the long-run. In this section, we point out that this is not the same as TCP-friendliness, for two reasons.

- The loss event ratio experienced by TCP and our adaptive source may differ.
- TCP may not satisfy its own equation.

We address the former point with some detail (within the space limitation of this article), and illustrate the latter on an example.

Figure 11: Same setting as with Fig. 10, but function f is PFTK-simplified.

5.1 Loss Event Ratio seen by Various Sources

We first perform a simple analysis, as follows. Assume that the sources in the network are driven by the hidden congestion process $Z(t)$ that evolves in real time, $t \in \mathbb{R}$, and takes values on a countable state space E . The state transitions are clocked by the point process $(T'_n)_n$; assumed to be stationary with finite intensity λ' . This is an approximation which fits with the case of a source with negligible influence on the global network [15].

Let $\pi_i := P(Z(0) = i)$ denote the steady-state probability that the congestion process is in the state $i \in E$. Define

$$\bar{p}_i = \frac{1}{E_T^0[\theta_0 | Z(0) = i]}.$$

Define, also, $\bar{x}_i = E[X(0) | Z(0) = i]$. Note that $\bar{p}_i$ and $\bar{x}_i$ are the loss event ratio and the average send rate while the congestion process is in the state i . We show in the appendix

$$\bar{p} = \frac{\sum_{i \in E} b_i \bar{p}_i \bar{x}_i \pi_i}{\sum_{i \in E} b_i \bar{x}_i \pi_i}, \quad (14)$$

where

$$b_i = \frac{E_{T'}^0[\sum_{n \in \mathbb{Z}} \theta_n 1_{T_n \in [0, S'_0)} | Z(0) = i]}{E_{T'}^0[\int_0^{S'_0} X(s) ds | Z(0) = i]}.$$

We would have, $b_i \rightarrow 1$, as $\frac{\lambda'}{\lambda_i} \rightarrow 0$, for $i \in E$ (here, by definition, $\lambda_i = 1/E_T^0[S_0 | Z(0) = i]$). The limit corresponds to the sepa-

Figure 12: Same setting as with Fig. 10, but function f is SQRT.

ration of timescales argument; we assume the congestion process evolves more slowly than the timescale of the control (remember that the control is clocked by the loss events). We base our further discussion on this limit loss event ratio, i.e.

$$\bar{p} \rightarrow \frac{\sum_{i \in E} \bar{p}_i \bar{x}_i \pi_i}{\sum_{i \in E} \bar{x}_i \pi_i}. \quad (15)$$

If our source is non-adaptive (call it “Poisson”) then $\bar{x}_i = \bar{x}$ is independent of i . The resulting loss event ratio $\bar{p}_P = \sum_{i \in E} \pi_i \bar{p}_i$ can be thought of as the time-average of the network loss event ratio; except for possible aliasing effects, it should be close to what a constant bit rate (CBR) source would experience. Now if, like TCP, our source is very responsive, i.e., follows the hidden congestion process pretty closely, then $\bar{x}_i$ depends on i in the following way: $\bar{x}_i$ is large for “good” states ($\bar{p}_i$ small) and small for bad states ($\bar{p}_i$ large). Thus we should have a smaller $\bar{p}$. For TCP, this is confirmed by measurements in [14]. The more responsive the source is, the more pronounced this should be; now TCP is expected to be more responsive than our adaptive source, whose responsiveness depends on the averaging window L . We summarize this as follows (see Fig. 15 for an illustration).

CLAIM 3. *The loss event ratios for TCP ($\bar{p}_T$), our adaptive equation based rate controlled source ($\bar{p}_E$), and a non-adaptive source (Poisson) ($\bar{p}_P$) should be in the relation*

$$\bar{p}_T \leq \bar{p}_E \leq \bar{p}_P.$$

The more responsive source E is, the closer $\bar{p}_E$ should be to $\bar{p}_T$.

Figure 13: Upper graph shows normalized throughput versus the loss event ratio as attained by a source with constant packet rate, but controlled packet sizes. The connection goes through a loss module where packet loss occurs. The length of the loss event interval estimator is $L = 4$. The lower graph shows squared coefficient of variation of $(\hat{\theta}_n)_n$.

Figure 14: Same setting as in Fig. 13, but with the length of the loss event interval estimator $L = 8$.

5.2 Putting Things Together

Claim 3 tells us that our adaptive source sees a higher loss event ratio than TCP, which drives it in the TCP-friendly direction, on top and above the factors mentioned earlier. Assuming (as is most common) that the conditions for conservativeness in Section 3 apply, we would have $\bar{x}_A \leq f(\bar{p}_E) \leq f(\bar{p}_T)$, (the latter is because f is decreasing). This makes our adaptive source TCP-friendly under the assumption that TCP does satisfy its equation. Unfortunately, this is only approximately true. Fig. 16 shows an experiment where TCP is below the formula PFTK-standard for light load and above for high loads. Fig. 17 shows that, as a result, TFRC flows have higher throughput for medium load than TCP. This is in spite of TFRC being conservative (Fig. 9) and experiencing higher loss than TCP (Fig. 15), as predicted by our theory. This illustrates the importance of separating the 3 factors identified in this paper.

Figure 15: Loss event ratio as experienced by TFRC, TCP, and Poisson connections versus N (number of TFRC and number TCP connections in one bottleneck). We have $\bar{p}_T \leq \bar{p}_E \leq \bar{p}_P$ as expected. Also, the smoother the TFRC flows (larger L), the higher the loss event ratio.

6. CONCLUSION

Our study should help designers of TCP-friendly equation-based rate control better understand the trade-offs that have to be taken. First, it is important to separately verify the three elements: (1) conservativeness, (2) TCP loss event ratio versus this protocol's loss event ratio, and (3) TCP's obedience to its own formula. Failing to do so blurs the setting of parameters. Second, one should be aware of the strong dependency on the nature of function f ; SQRT behaves differently than PFTK. If PFTK is used, and under some conditions on the loss process defined in Claim 1, very pronounced conservativeness should be expected for heavy loss. Under some other conditions (Claim 2), the opposite may hold. In any case, the more variable the estimator is, the more pronounced the effect is. This might lead a protocol designer to change some parameters of her protocol, in order to correct either effect. Understanding why and when the effects occur is essential to avoid undesired corrections. By their general nature, and because they are based on analysis, our results should apply to a large class of situations including dynamically varying environments.

There are three directions for further work. First, our findings should be confronted with measurements; in particular, the autocovariance property of loss event intervals will be of interest. Second, the impact of the variation of round trip time needs to be incorporated; this can be done using the same approach as we used for the other factors. Third, we focused in this paper on the relation between loss event ratio and throughput, leaving aside any prediction of which values these variables may take, in a given setting; it will be interesting to study this in more detail, in particular the existence of stable points.

7. ACKNOWLEDGEMENTS

We would like to thank Chadi Barakat and Catherine Boutremans for sharing with us their *ns* simulator expertise. We also thank anonymous reviewers for their valuable comments.

Figure 16: $f(\bar{p})$ for PFTK-standard formula versus $\bar{x} = E[X(0)]$ for TCP Sack1 sources. For medium throughputs, TCP Sack1 is below the formula.

8. REFERENCES

- [1] Francois Baccelli and Pierre Brémaud. *Elements of Queueing Theory*, volume 26. Applications of Mathematics, Springer-Verlag, 1991.
- [2] Deepak Bansal, Hari Balakrishnan, Sally Floyd, and Scott Shenker. Dynamic Behavior of Slowly-responsive Congestion Control Algorithms. In *Proc. of ACM Sigcomm'01*, San Diego, California, USA, August 2001.
- [3] C. Barakat. TCP/IP Modeling and Validation. *IEEE Network*, 15(3):38–47, May-June 2001.
- [4] C. Boutremans and J.-Y. Le Boudec. Adaptive Delay Aware Error Control for Internet Telephony. In *Proc. of 2nd IP-Telephony Workshop*, pages 81–92, Columbia University, New York, April 2001.
- [5] Neal Cardwell, Stefan Savage, and Thomas Anderson. Modeling TCP Latency. In *Proc. of the IEEE INFOCOM'2000*, Tel-Aviv, Israel, March 2000.
- [6] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 2. John-Wiley & Sons, 2 edition, 1966.
- [7] Sally Floyd. Best Control Principles, September 2000. rfc-2914, www.ietf.org.
- [8] 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.
- [9] 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.
- [10] Mark Handley, Jitendra Padhye, Sally Floyd, and Jörg Widmer. TCP Friendly Rate Control (TFRC) Protocol Specification, IETF internet-draft, July 2001. <http://search.ietf.org/internet-drafts/draft-ietf-tsvwg-tfrc-02.txt>.
- [11] Jamshid Mahdavi and Sally Floyd. TCP-Friendly Unicast Rate-based Flow Control. Technical note sent to end-2-end interest mailing list, http://www.psc.edu/networking/papers/tcp_friendly.html, January 1997.
- [12] 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.
- [13] 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.

Figure 17: The top graph shows throughput attained by TFRC and TCP Sack1. The lower graphs shows ratio of TFRC and TCP Sack1 throughputs versus the number of connections.

- [14] Vern Paxson. End-to-end Internet Packet Dynamics. *IEEE/ACM Trans. on Networking*, pages 277–292, June 1999.
- [15] Salamatian M. R. and Vaton S. Hidden Markov Modeling for Network Communication Channels. In *Proceedings of ACM Sigmetrics / Performance 200*, Cambridge, Massachusetts, USA, June 2001.
- [16] Rockafellar R. T. *Convex Analysis*. Princeton University Press, Princeton, 1970.
- [17] Milan Vojnović and Jean-Yves Le Boudec. Some Observations on Equation-based Rate Control. In *Proc. of ITC-17*, pages 173–184, Salvador, Bahia, Brazil, Dec. 2001.
- [18] J. Widmer, R. Denda, and M. Mauve. A Survey on TCP-Friendly Congestion Control. *IEEE Network*, 15(3):28–37, May-June 2001.
- [19] 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.
- [20] Yin Zhang, Nick Duffield, Vern Paxson, and Scott Shenker. On the Constancy of Internet Path Properties. In *Proc. of ACM Sigcomm Internet Measurement Workshop*, Nov. 2001.

APPENDIX

A. APPENDIX

Proof of Proposition 1

PROOF. The starting point is Palm inversion formula [1], which relates time average to event average with respect to some point process. In our setting it reads as:

$$E[X(0)] = \frac{E_T^0[\int_0^{S_0} X(s)ds]}{E_T^0[S_0]}. \quad (16)$$

We can think of (16) as the ratio of the expected number of packets sent in-between two successive loss events and the expected loss event inter-arrival time. However, it is important to remember the expected values are with respect to the Palm probability that is as seen at the loss event instants.

For the basic control this gives

$$E[X(0)] = \frac{E_T^0[X_0 S_0]}{E_T^0[S_0]}. \quad (17)$$

By (3), $\theta_n = X_n S_n$, $X_n = f(1/\hat{\theta}_n)$, and $S_n = \frac{\theta_n}{f(1/\hat{\theta}_n)}$. Combining the last three identities into (17) we obtain (8). $\square$

Proof of Proposition 2

PROOF. Note that if $\hat{\theta}_{n+1} \leq \hat{\theta}_n$, then the comprehensive control in Equation (9) behaves the same way as the basic control (3). In this case $\theta_n = X_n S_n$, and thus $S_n = \frac{\theta_n}{f(1/\hat{\theta}_n)}$. Next, we consider the other case, i.e. $\hat{\theta}_{n+1} > \hat{\theta}_n$.

Suppose $U_n < S_n$. Then, for $U_n \leq t < S_n$,

$$\theta(t) = \theta(U_n) + \int_{U_n}^t X(s - \tau_s) ds,$$

where τ_s is the feedback delay. Using the definition of the comprehensive control (Equation (4)), we obtain the following delayed differential equation (DDE):

$$\frac{d\theta(t)}{dt} = f\left(\frac{1}{w_1\theta(t - \tau_t) + W_n}\right),$$

where $W_n = \sum_{l=1}^{L-1} w_{l+1}\theta_{n-l}$.

Shift the time origin to T_n and solve $\theta(S_n) = \theta_n$ for S_n . For this we need first to solve the above DDE for a given function f .

We neglect the feedback delay and thus consider the ordinary differential equation (ODE):

$$\frac{d\theta(t)}{dt} = f\left(\frac{1}{w_1\theta(t) + W_n}\right). \quad (18)$$

By doing so, note that we approximate S_n with a smaller value (this is due to $f(1/x)$ non-decreasing with x). Given that S_n acts in the denominator of the throughput expression, we in fact obtain an upper bound on the throughput.

Finally, solving (18) for PFTK-simplified formula, we come to the expression given in Equation (9). By the same argument as above we claim Equation (9) is an upper bound on the throughput for any function f that is less than or equal to PFTK-simplified (or SQRT by setting $c_2 = c_3 = 0$). $\square$

Proof of Theorem 1

PROOF. Define $g(x) := \frac{1}{f(\frac{1}{x})}$. Also call $m = \frac{1}{p}$, thus $E_T^0[\theta_0] = E_T^0[\hat{\theta}_0] = m$. From Equation (8), conservativeness is equivalent to

$$E_T^0[\theta_0 g(\hat{\theta}_0)] \geq m g(m) \quad (19)$$

Function g is convex, thus is above its tangents:

$$g(x) \geq (x - m)g'(m) + g(m).$$

Apply the above to $x = \hat{\theta}_0$, multiply by θ_0 and take the expectation. After some calculus, this shows Equation (10).

Now f is decreasing. Since $\text{cov}_T^0[\theta_0, \hat{\theta}_0] \leq 0$, it follows from Equation (10) that the control is conservative. $\square$

Proof of Proposition 3

PROOF. Use the same notation as in the proof of Theorem 1. By Equation (8) the ratio of throughput to $f(\bar{\rho})$ is equal to

$$\rho := \frac{mg(m)}{\mathbb{E}_T^0 \left[\theta_0 g(\hat{\theta}_0) \right]}. \quad (20)$$

Now we have

$$g^{**}(x) \leq g(x) \leq rg^{**}(x).$$

The same argument as in the proof of Theorem 1, applied to g^{**} instead of g , shows that and thus $\rho \leq r$. $\square$

Proof of Theorem 2

PROOF. Use the same notation as in the proof of Theorem 1.

Part 1. By (C2)

$$\mathbb{E}_T^0[\theta_0 g(\hat{\theta}_0)] \geq \frac{m}{\mathbb{E}_T^0 \left[\frac{1}{g(\hat{\theta}_0)} \right]}, \quad (21)$$

now (F2) means that $\frac{1}{g}$ is concave, thus by Jensen's inequality:

$$\mathbb{E}_T^0 \left[\frac{1}{g(\hat{\theta}_0)} \right] \leq \frac{1}{g(\mathbb{E}_T^0[\hat{\theta}_0])}, \quad (22)$$

which combined with the previous equation shows that the control is conservative.

Part 2. By (C2c) and (F2c) we have the reverse inequalities in Equation (21) and Equation (22), but the inequality is strict in Equation (22) because convexity is strict and $\hat{\theta}_n$ is not a degenerate random variable. $\square$

Derivation of Equation (14)

PROOF. We start from Equation (1). By Neveu's exchange formula ([1], Sec. 3.3.4) and simple conditioning

$$\begin{aligned} \bar{\rho} = \frac{1}{\mathbb{E}_T^0[\theta_0]} &= \frac{\mathbb{E}_{T'}^0[\sum_{n \in \mathbb{Z}} \mathbf{1}_{T_n \in [0, S'_0]}]}{\mathbb{E}_{T'}^0[\sum_{n \in \mathbb{Z}} \theta_n \mathbf{1}_{T_n \in [0, S'_0]}]} \\ &= \frac{\sum_{i \in E} \mathbb{E}_{T'}^0[\sum_{n \in \mathbb{Z}} \mathbf{1}_{T_n \in [0, S'_0]} | Z(0)=i] \mathbb{P}_{T'}^0(Z(0)=i)}{\sum_{i \in E} \mathbb{E}_{T'}^0[\sum_{n \in \mathbb{Z}} \theta_n \mathbf{1}_{T_n \in [0, S'_0]} | Z(0)=i] \mathbb{P}_{T'}^0(Z(0)=i)}. \end{aligned} \quad (23)$$

We show that the above is equivalent to Equation (14).

As an application of Palm inversion formula to $X(0)\mathbf{1}_{Z(0)=i}$, we obtain

$$\bar{x}_i = \mathbb{E}[X(0) | Z(0) = i] = \frac{\mathbb{E}_{T'}^0[\int_0^{S'_0} X(s) ds | Z(0) = i]}{\mathbb{E}_{T'}^0[S'_0 | Z(0) = i]},$$

where we also use (obtained by another application of Palm inversion formula to $\mathbf{1}_{Z(0)=i}$)

$$\pi_i = \mathbb{P}(Z(0) = i) = \frac{\mathbb{E}_{T'}^0[S'_0 | Z(0) = i]}{\mathbb{E}_{T'}^0[S'_0]} \mathbb{P}_{T'}^0(Z(0) = i).$$

By a similar argument, from Neveu's exchange formula applied to $\theta_0 \mathbf{1}_{Z(0)=i}$, we have

$$\frac{1}{\bar{\rho}_i} = \mathbb{E}_T^0[\theta_0 | Z(0) = i] = \frac{\mathbb{E}_{T'}^0[\sum_{n \in \mathbb{Z}} \theta_n \mathbf{1}_{T_n \in [0, S'_0]} | Z(0) = i]}{\mathbb{E}_{T'}^0[\sum_{n \in \mathbb{Z}} \mathbf{1}_{T_n \in [0, S'_0]} | Z(0) = i]},$$

where we use the identity obtained by Neveu's exchange formula applied to $\mathbf{1}_{Z(0)=i}$,

$$\mathbb{P}_T^0(Z(0) = i) = \frac{\mathbb{E}_{T'}^0[\sum_{n \in \mathbb{Z}} \mathbf{1}_{T_n \in [0, S'_0]} | Z(0) = i]}{\mathbb{E}_{T'}^0[\sum_{n \in \mathbb{Z}} \mathbf{1}_{T_n \in [0, S'_0]}]} \mathbb{P}_{T'}^0(Z(0) = i).$$

Finally, by plugging the above expressions for $\bar{x}_i$, π_i , and $\bar{\rho}_i$ into Equation (14) we recover Equation (23). $\square$

A.1 Comparison of conditions in Theorem 1 and Theorem 2

Use the same notation as in the proof of Theorem 1. Notice that by the assumption that f is non-increasing (Section 2.2), g is non-increasing as well. For technical convenience, suppose g is strictly decreasing at m , that is $g'(m) < 0$.

PROPOSITION 4. Assume (F2c), (C2c), and (V) hold, i.e., the second part of Theorem 2 applies. Then, in Theorem 1, if (F2) is true, it must be that (C1) does not hold.

PROOF. Note the equivalence

$$\text{cov}_T^0[X_0, S_0] > 0 \Leftrightarrow \mathbb{E}_T^0[\theta_0 g(\hat{\theta}_0)] < \frac{m}{\mathbb{E}_T^0[f(1/\hat{\theta}_0)]}.$$

Under (F2), by the same argument as in Theorem 1,

$$\mathbb{E}_T^0[\theta_0 g(\hat{\theta}_0)] \geq g'(m) \text{cov}_T^0[\theta_0, \hat{\theta}_0] + mg(m).$$

Suppose (C2c) and (F2) are true, then from the last two inequalities, we conclude that the following is implied:

$$\text{cov}_T^0[\theta_0, \hat{\theta}_0] > \frac{m}{g'(m)} \left(\frac{1}{\mathbb{E}_T^0[f(1/\hat{\theta}_0)]} - \frac{1}{f(1/m)} \right).$$

Finally, if $f(1/x)$ is strictly convex with x , that is (F2c) holds, and (V) holds, then the right-hand side in the above inequality is strictly positive, and thus (C1) does not hold. $\square$

An Intermediate Property between Theorem 1 and Theorem 2

The following theorem is intermediate between Theorem 1 and Theorem 2.

THEOREM 3. If (F1) and

$$(C3) \quad \text{cov}_T^0[X_0 S_0, \frac{1}{X_0}] \geq 0,$$

the basic control is conservative.

The proof is similar to that of Theorem 2 and is not given here. If the convexity condition (F1) is almost true, then the same as in Proposition 3 holds.

This theorem is intermediate between Theorem 1 and Theorem 2. Indeed (F2) $\Rightarrow$ (F1) and (C3) $\Rightarrow$ (C2). The former is straightforward; a proof of the latter implication uses the convexity of $1/x$. Thus Theorem 3 is with a weaker condition on the function f than Theorem 2, but this comes at the expense of having a stronger condition on the statistics of $(\theta_n)_n$. A natural question is whether both Theorem 3 and the first part of Theorem 2 derive from a more general theorem, which would state that under the combination of the less restrictive conditions (F1) and (C2), the control would be conservative. But this is not true; a counter-example is the case presented in the second paragraph of the interpretation of Theorem 2.