

Extending equation-based congestion control to high-speed and long-distance networks

Lisong Xu *

Department of Computer Science and Engineering, University of Nebraska-Lincoln, Lincoln, NE 68588-0115, USA

Available online 27 November 2006

Abstract

TCP-friendly rate control (TFRC), an equation-based congestion control protocol, has been a promising alternative to TCP for multimedia streaming applications. However, TFRC using the TCP response function, has the same poor performance as TCP in high-speed and long-distance networks. In this paper, we propose high-speed equation-based rate control (HERC), as an extension of TFRC by replacing the TCP response function with a high-speed response function. HERC could be used for applications, such as high-definition video streaming, and remote collaboration involving high-resolution visualization, which prefer a high-speed and relatively smooth sending rate. The impact of a general high-speed response function on the throughput and smoothness of HERC is studied analytically and verified by using simulation. Our result indicates that by using the response function of a high-speed TCP variant and tuning HERC parameters accordingly, HERC can compete fairly with high-speed TCP flows in the same network, while maintaining the desired smoothness of TFRC.

© 2006 Elsevier B.V. All rights reserved.

Keywords: TCP-friendly rate control; Equation-based congestion control; High-speed TCP variants; Multimedia streaming

1. Introduction

TCP-friendly rate control (TFRC) [8], an equation-based congestion control protocol, is being adopted in Internet standards [9–11] for congestion control of media streaming applications. In order to compete fairly with TCP, TFRC adjusts its sending rate by using the TCP response function [16], which describes the sending rate of a TCP flow as a function of the loss rate, delay, and size of packets. TFRC is able to achieve approximately the same

long-term throughput as TCP under comparable network conditions, while providing a smoother sending rate.

Standard TCP has been working very well in traditional networks; however, it has been shown [7,5] that TCP cannot efficiently utilize the huge link capacity in high-speed and long-distance networks, and has an unsatisfactory performance. Recently, several high-speed TCP variants, such as HSTCP [7], STCP [14], BICTCP [26], FAST [13], HTCP [19], LTCP [4], TCP-Africa [15], and TCP-Westwood [6], have been proposed to achieve higher throughput than TCP in high-speed and long-distance networks, and be reasonably friendly to TCP traffic.

* Tel.: +1 402 472 1053; fax: +1 402 472 7767.

E-mail address: xu@cse.unl.edu

TFRC using the TCP response function, consequently, has the same poor performance as TCP in high-speed and long-distance networks. In this paper, we extend TFRC by *replacing the TCP response function with a high-speed response function*, and propose High-speed Equation-based Rate Control (HERC). By using the response function of a high-speed TCP variant, HERC can compete fairly with high-speed TCP flows in the same network, while providing a relatively smooth sending rate. For instance, HERC with HSTCP response function [7] can achieve approximately the same long-term throughput as HSTCP flows under comparable network conditions. The selection of the high-speed response function may depend on which high-speed TCP variant will finally become the dominant one in the Internet, and this is out of the scope of this paper. In this paper, we consider HERC with a general form of high-speed response functions. The proposed HERC could be used for applications, such as high-definition video streaming [17,1,2], and remote collaboration involving high-resolution visualization [20], which prefer a high-speed and relatively smooth sending rate.

Our early work [25] studies the rate smoothness of HERC. In this paper, we provide a more comprehensive study of the performance of HERC. Specifically, the *main contributions of our work* are as follows. *First*, we study the impact of a high-speed response function on the long-term throughput¹ of HERC. For the basic control (defined in [23] and described in the next section), the upper and lower bounds of the long-term throughput of HERC are derived. We find that with certain high-speed response functions, such as the one of STCP [14], the lower bound is indeed the long-term throughput of HERC with basic control. A closed-form expression of the lower bound is also obtained under some reasonable assumptions.

Second, we study the impact of a high-speed response function on the smoothness of HERC. Following the work of [8], we measure the smoothness by the coefficient of variation (CoV) of sending rates. Our result indicates that while HERC achieves higher throughput than TFRC, it has worse smoothness than TFRC, if they use the same loss history size. We obtain a closed-form expression for the CoV of the sending rates as a function

of the loss history size and response function coefficients. By adjusting the HERC loss history size accordingly, HERC can use any high-speed response function while maintaining the same smoothness as TFRC.

The rest of the paper is organized as follows. Section 2 gives an overview of TFRC. Section 3 presents HERC. Section 4 analyzes the long-term throughput of HERC, and Section 5 studies the smoothness of HERC. Section 6 shows the simulation results. Related work and conclusion can be found in Sections 7 and 8.

2. Background and notation

TFRC [9,8] is an equation-based congestion control protocol designed for best-effort multimedia streaming, by maintaining a relatively smooth sending rate, while providing approximately the same average sending rate as TCP under comparable network conditions.

The sending rate of TFRC [9] is calculated by using Eq. (1) below, which is a simplified version of the TCP response function proposed in [16] when assuming $p < 0.54$. *TCP response function* describes TCP throughput R as a function of packet size s , round-trip time rtt , loss event rate p , timeout period rto , and b that is the number of packets acknowledged by an ACK.

$$R_{\text{TCP}} = \frac{s}{rtt \cdot \sqrt{\frac{2bp}{3}} + rto \cdot 3\sqrt{\frac{3bp}{8}}p(1 + 32p^2)}. \quad (1)$$

When p is small and b is 1, Eq. (1) can be reduced to the following more simplified form, which is independent of rto .

$$R_{\text{TCP}} = \frac{s}{rtt} \frac{1.22}{p^{0.5}}. \quad (2)$$

TFRC defines a *loss event* as one or more packets lost (or marked with Explicit Congestion Notification (ECN)) within one round-trip time. Following the notation in [23], we denote θ_i to be the i th *loss interval*, which is the number of packets sent between the i th and $(i + 1)$ th loss events, and denote $\hat{\theta}_i$ to be the *weighted average loss interval* calculated by using the L most recent loss intervals as follows, where L is referred to as the *loss history size*.

$$\hat{\theta}_i = \sum_{l=1}^L w_l \theta_{i-l}, \quad (3)$$

where w_l is given by

¹ In this paper, we use throughput and sending rate interchangeably.

$$w_l = \begin{cases} \frac{1}{\frac{3}{4}l} & 0 < l \leq \frac{L}{2}, \\ \frac{L-l+1}{\frac{3}{4}l(\frac{L}{2}+1)} & \frac{L}{2} < l \leq L. \end{cases} \quad (4)$$

TFRC recommends $L = 8$, and then we have

$$\hat{\theta}_i = \frac{\theta_{i-1}}{6} + \frac{\theta_{i-2}}{6} + \frac{\theta_{i-3}}{6} + \frac{\theta_{i-4}}{6} + \frac{\theta_{i-5}}{7.5} + \frac{\theta_{i-6}}{10} + \frac{\theta_{i-7}}{15} + \frac{\theta_{i-8}}{30} \quad \text{if } L = 8. \quad (5)$$

TFRC estimates the loss event rate by the inverse of the weighted average loss interval. Let p_i denote the loss event rate estimated right after the i th loss event. We have

$$p_i = \frac{1}{\hat{\theta}_i}. \quad (6)$$

After estimating the loss event rate, TFRC then adjusts its sending rate. Let R_i denote the new TFRC sending rate calculated right after the i th loss event. We have

$$R_i = R_{\text{TCP}}(p_i), \quad (7)$$

where $R_{\text{TCP}}(\cdot)$ is either Eq. (1) or Eq. (2).

To simplify the analysis of TFRC, it is usually [23,18] assumed that TFRC keeps the same sending rate R_i between the i th and $(i+1)$ th loss events. This is called *basic control* [23]. The actual TFRC implementation [9], called *comprehensive control*, includes more sophisticated algorithms, such as the most recent loss interval calculation and history discounting, which address the impact of the most recent loss interval on the sending rate. It has been observed [23] that the long-term throughput of TFRC is mainly determined by the basic control, while the comprehensive control has a larger impact on the short-term throughput of TFRC.

3. High-speed equation-based rate control

Since TCP performs poorly in high-speed and long-distance networks [7], TFRC using the TCP response function has the same poor performance as TCP in these networks. In this section, we describe high-speed equation-based rate control (HERC) as an extension of TFRC for multimedia streaming over high-speed and long-distance networks. HERC maintains the overall design of TFRC (such as loss event rate estimation, basic control, and comprehensive control), but replaces the TCP response function with a high-speed response function. By replacing the TCP response function

with the response function of a high-speed TCP variant, HERC can compete fairly with high-speed TCP flows in the same network, while maintaining the desired smoothness of TFRC.

The response function of a high-speed TCP variant determines several important properties of the protocol, such as bandwidth scalability, and TCP friendliness. Fig. 1 shows the response functions (in Packets/RTT) of TCP, HSTCP [7], STCP [14], and BICTCP [26] as the loss event rate varies. (1) The bandwidth scalability of a protocol, defined as the ability of the protocol to achieve high throughput in a high-speed network, is typically determined by its sending rates under low loss event rates. In Fig. 1, a protocol becomes more bandwidth scalable as its sending rate gets higher under low loss event rates. Thus, HSTCP, STCP, and BICTCP all show better bandwidth scalability than TCP. (2) The TCP friendliness defines whether a protocol is being fair to TCP, and it is critical to the safety of deploying a protocol in the Internet. The TCP friendliness of a high-speed protocol is determined by the point where its response function intersects that of TCP. This is because a high-speed protocol usually runs its own algorithm only when the loss rate is smaller than the intersection point, and it acts the same as TCP when the loss rate is greater than the intersection point. Therefore, all these high-speed protocols are claimed to have reasonable TCP friendliness.

HERC may use the response function of any of these high-speed TCP variants, and then is able to achieve the same bandwidth scalability and TCP friendliness as the corresponding protocol. The selection of a high-speed response function may depend on which high-speed TCP variant will finally become the dominant one in the Internet, and this is out of the scope of this paper. Instead of using the response function of a specific protocol, in this paper, we study the performance of HERC with a general response function in the following form.




Fig. 1. Response functions of TCP and high-speed TCP variants.

$$R_{\text{HERC}} = \frac{s}{rtt} \frac{c}{p^d}, \quad (8)$$

where c and d are two protocol-specified parameters. TCP and its variants usually have constant c and d . For instance, TCP has $c = 1.22$ and $d = 0.5$, HSTCP [7] has $c = 0.12$ and $d = 0.835$, STCP [14] has $c = 0.08$ and $d = 1.0$, and BICTCP has [26] $c = 15.5$ and $d = 0.5$ for small loss event rate p .

4. Throughput analysis

In this section, we analyze the throughput of the proposed HERC. Vojnović and Le Boudec [23] analyze the throughput of TFRC by considering two control modes: basic control and comprehensive control. They find that the long-term throughput mainly depends on the basic control, and the comprehensive control has a larger impact on the short-term throughput. Therefore, in this section, we focus on the analytical study of long-term throughput of HERC with the basic control, and give the upper and lower bounds. Both of them together provide information on how well HERC tracks the target throughput.

We define S_j as the average sending rate within time interval j from t_j to $t_j + \Delta t$, where Δt is the duration of a time interval.

$$S_j = \frac{\text{number of bits sent between } t_j \text{ and } t_j + \Delta t}{\Delta t}. \quad (9)$$

The average sending rate $\mathbb{E}[S]$ of a protocol in n time intervals can be obtained by

$$\mathbb{E}[S] = \frac{\sum_{j=1}^n S_j}{n}. \quad (10)$$

Note that R_i defined in Section 2 is different from S_j , the former is the sending rate between the i th and $(i + 1)$ th loss events, whereas the later is the sending rate in time interval j . If every loss interval has the same duration, we have $\mathbb{E}[R] = \mathbb{E}[S]$.

4.1. Long-term throughput of HERC with basic control

Vojnović and Le Boudec [23] show that for the basic control, the long-term throughput of an equation-based congestion control protocol with response function $R(p)$ can be obtained by using Eq. (11), if successive loss intervals are independent of each other.

$$\mathbb{E}[S] = \frac{\mathbb{E}[\theta]}{\mathbb{E}\left[\frac{\theta}{R(\frac{1}{\theta})}\right]} = \frac{1}{\mathbb{E}\left[\frac{1}{R(\frac{1}{\theta})}\right]}. \quad (11)$$

There are indications [29,27] that the assumption that successive loss intervals are independent of each other is reasonable. It is well known [29] that the packet losses of a flow are not independently and identically distributed (IID), due to the strong correlation between consecutive packet losses. However, since all packet losses within the same RTT are counted as a single loss event, it is reasonable to consider that loss intervals are IID. Furthermore, several recent studies [29,27] of Internet traffic show that loss intervals can be approximately modeled as a Poisson process.

4.2. Upper and lower bounds of HERC with basic control

In this subsection, we obtain the upper and lower bounds of the long-term throughput of HERC with basic control by using Jensen's inequality. Similar methods have been used in [18,23] to obtain the upper and lower bounds of long-term throughput of TFRC.

Theorem 1. *For the basic control, the average sending rate $\mathbb{E}[S]$ of HERC with response function $R(p) = \frac{s}{rtt} \frac{c}{p^d}$ ($c > 0$, $1 \geq d \geq 0.5$) is bounded by the following inequality, if successive loss intervals are independent of each other. The equality with the lower bound holds if $d = 1$.*

$$R\left(\mathbb{E}\left[\frac{1}{\hat{\theta}}\right]\right) \leq \mathbb{E}[S] < R\left(\frac{1}{\mathbb{E}[\hat{\theta}]}\right).$$

Proof. Let $R_u(\hat{\theta}) = \frac{1}{R(\frac{1}{\hat{\theta}})}$. Its second derivative is given by $\frac{d^2(R_u(\hat{\theta}))}{d(\hat{\theta})^2} = \frac{rtt}{s \cdot c} \frac{d(d+1)}{\hat{\theta}^{d+2}}$, which is positive. Therefore, $R_u(\hat{\theta})$ is a convex function of $\hat{\theta}$. Thus, by Jensen's inequality, we have $\mathbb{E}[R_u(\hat{\theta})] > R_u(\mathbb{E}[\hat{\theta}])$.

Let $R_l(\frac{1}{\hat{\theta}}) = \frac{1}{R(\frac{1}{\hat{\theta}})}$. Its second derivative is given by $\frac{d^2(R_l(\frac{1}{\hat{\theta}}))}{d(\frac{1}{\hat{\theta}})^2} = \frac{rtt}{s \cdot c} \frac{d(d-1)}{\hat{\theta}^{d-2}}$, which is non-positive for any d between 0.5 and 1. Therefore, $R_l(\frac{1}{\hat{\theta}})$ is a concave function of $\frac{1}{\hat{\theta}}$, and by Jensen's inequality, we conclude $\mathbb{E}\left[R_l\left(\frac{1}{\hat{\theta}}\right)\right] \leq R_l\left(\mathbb{E}\left[\frac{1}{\hat{\theta}}\right]\right)$. The equality holds when the second derivative is zero (i.e., $d = 1$).

By using Eq. (11), we have $\mathbb{E}[S] = 1/\mathbb{E}[R_u(\hat{\theta})] = 1/\mathbb{E}[R_l(\frac{1}{\hat{\theta}})]$, and then we can prove the theorem after some straightforward manipulation. \square

Comment 1: The upper bound $R(\frac{1}{\mathbb{E}[\hat{\theta}]})$ is the target sending rate, since $\hat{\theta}$ is an unbiased estimator of the inverse of the actual loss event rate p (i.e., $\mathbb{E}[\hat{\theta}] = \frac{1}{p}$). Therefore, for the basic control, the average sending rate $\mathbb{E}[S]$ of HERC is less than the target sending rate. This is consistent with the TFRC analysis in [23,18]. Since the comprehensive control adds some throughput increase to the basic control, the final throughput of HERC may be higher than this upper bound.

Comment 2: The lower bound $R(\mathbb{E}[\frac{1}{\hat{\theta}}])$ is not the target sending rate, since $\frac{1}{\hat{\theta}}$ is a biased estimator of the loss event rate due to the convexity of $\frac{1}{\hat{\theta}}$ (i.e., $\mathbb{E}[\frac{1}{\hat{\theta}}] \neq p$). We note that if $d = 1$, the average sending rate $\mathbb{E}[S]$ of HERC with basic control is equal to the lower bound. Since the comprehensive control does not decrease the throughput of the basic control, the obtained lower bound is also a lower bound of HERC with comprehensive control.

4.3. Closed-form expression for the lower bound

The lower bound of HERC throughput obtained in the last subsection is an important property of HERC. (1) As we proved that if the exponent d of a high-speed response function is 1, such as the one of STCP [14], the lower bound is indeed the long-term throughput of HERC with basic control. (2) The lower bound instead of the upper bound (i.e., the target throughput) can be safely used as the maximum media playback rate of a media stream using HERC. Otherwise, the media playback rate may sometimes be faster than the media arrival rate, and then the player is forced to stop playback until enough data has been received to resume. In this subsection, we therefore obtain a closed-form expression for the lower bound of HERC throughput. We first calculate the probability density function $f_{\hat{\theta}}(k)$ of the weighted average loss interval $\hat{\theta}$, and then obtain a closed-form expression for $\mathbb{E}[\frac{1}{\hat{\theta}}]$.

4.3.1. Distribution of the weighted average loss interval

In the previous analysis, we only assume that successive loss intervals are independent of each other. In this subsection, we further assume that loss interval θ is exponentially distributed with an expected

value of $\frac{1}{p}$, which is a reasonable assumption as observed by several recent Internet measurement experiments [29,27].

The probability density function $f_{\hat{\theta}}(k)$ of $\hat{\theta}$ can be obtained by using the Laplace transform method. Since θ is exponentially distributed with an expected value of $\frac{1}{p}$, the Laplace transform of θ is given by $L_{\theta}(s) = \frac{p}{p+s}$. Since weighted average loss interval $\hat{\theta}_i = \sum_{l=1}^L \theta_{i-l} w_l$ is the sum of multiple independent random variables, the Laplace transform of $\hat{\theta}$ is given by $\prod_{l=1}^L L_{\theta}(w_l s)$. For example, if $L = 8$, we have

$$L_{\hat{\theta}}(s) = \left(\frac{p}{p+\frac{s}{6}}\right)^4 \left(\frac{p}{p+\frac{s}{7.5}}\right) \left(\frac{p}{p+\frac{s}{10}}\right) \left(\frac{p}{p+\frac{s}{15}}\right) \left(\frac{p}{p+\frac{s}{30}}\right) \quad \text{if } L = 8. \quad (12)$$

By inverting $L_{\hat{\theta}}(s)$, we can get probability density function $f_{\hat{\theta}}(k)$. For example, by inverting Eq. (12), we get $f_{\hat{\theta}}(k)$ as shown by Eq. (13). This probability density function is accurate; however, it is very complicated. In order to simplify our analysis, below we propose a simple function to approximate $f_{\hat{\theta}}(k)$.

$$f_{\hat{\theta}}(k) = 204.80e^{-0.075k} - 6.83e^{-0.1k} + 0.119e^{-0.15k} - 0.0002e^{-0.3k} + (0.000056k^3 - 0.018k^2 + 2.80k - 198.1)e^{-0.06k} \quad \text{if } L = 8. \quad (13)$$

Inspired by the fact that the sum of multiple independent and identical exponential random variables is an Erlang random variable [22], which has a relatively simple probability density function, we attempt to approximate $f_{\hat{\theta}}(k)$ with the probability density function of an Erlang random variable.

Define another weighted average loss interval $\hat{\theta}_i = \sum_{l=1}^M \theta_{i-l}/M$. For example, if $M = 7$, and we have

$$\hat{\theta}_i = \frac{\theta_{i-1}}{7} + \frac{\theta_{i-2}}{7} + \frac{\theta_{i-3}}{7} + \frac{\theta_{i-4}}{7} + \frac{\theta_{i-5}}{7} + \frac{\theta_{i-6}}{7} + \frac{\theta_{i-7}}{7} \quad \text{if } M = 7. \quad (14)$$

Based on the analysis of the moments of $\hat{\theta}$ and $\hat{\theta}_i$, we found that when $M = \frac{3}{4}L + 1$, the probability density function of $\hat{\theta}$ is very similar to that of $\hat{\theta}_i$. For example, when $L = 8$ and $M = 7$, $\hat{\theta}_i$ is given by Eq. (5), and $\hat{\theta}_i$ is given by Eq. (14). In this case, we can prove that random variables $\hat{\theta}$ and $\hat{\theta}_i$ have the same first moment (i.e., the expected value), and the difference ratios of their second and third moments are less than 0.2% and 0.6%,

respectively [27]. However, we do not have a rigorous proof that θ is the best approximation to $\hat{\theta}$ when $M = \frac{3}{4}L + 1$.

By using the probability density function of an Erlang random variable given in [22], we obtain $f_{\hat{\theta}}(k)$ as follows:

$$f_{\hat{\theta}}(k) = \frac{(Mp)^M}{(M-1)!} k^{M-1} e^{-Mpk}. \quad (15)$$

Fig. 2 shows the accurate $f_{\hat{\theta}}(k)$ obtained by inverting $L_{\hat{\theta}}(s)$, and $f_{\hat{\theta}}(k)$ given by Eq. (15). We can see that they are very close to each other, when $M = \frac{3}{4}L + 1$. Note that even though $f_{\hat{\theta}}(k) \approx f_{\hat{\theta}}(k)$, the value of $\hat{\theta}_i$ could be quite different from that of $\hat{\theta}_i$ given the same set of $\theta_{i-1}, \dots, \theta_{i-L}$. That is, variable $\hat{\theta}_i$ itself is not a good approximation to variable $\hat{\theta}_i$.

4.3.2. Closed-form expression for the lower bound

The lower bound of the average sending rate of HERC is $R(\mathbb{E}[\frac{1}{\hat{\theta}}])$, where $\mathbb{E}[\frac{1}{\hat{\theta}}]$ can be calculated by.

$$\begin{aligned} \mathbb{E}\left[\frac{1}{\hat{\theta}}\right] &= \int_0^{\infty} f_{\hat{\theta}}(k) \frac{1}{k} dk \approx \int_0^{\infty} f_{\hat{\theta}}(k) \frac{1}{k} dk \\ &= \frac{Mp}{M-1} \int_0^{\infty} \frac{(Mp)^{M-1}}{(M-2)!} k^{M-2} e^{-Mpk} dk \\ &= \frac{Mp}{M-1}. \end{aligned} \quad (16)$$

Finally, we have the following theorem.

Theorem 2. For both basic control and comprehensive control, the lower bound of the average sending rate of HERC with a loss history size of L and a loss event rate of p can be obtained by $R\left(\mathbb{E}\left[\frac{1}{\hat{\theta}}\right]\right) = R\left(\frac{Mp}{M-1}\right)$, where $M = \frac{3}{4}L + 1$, if the loss intervals are exponentially distributed.

4.4. Verifying upper and lower bounds of basic control

In this subsection, we verify our analysis of HERC basic control by using simulation. The performance of HERC with comprehensive control will be studied in Section 6.

We simulate a single HERC flow over a link with transmission errors. Transmission errors are generated so that loss intervals are IID with probability density function $f_{\theta}(k) = \lambda e^{-\lambda(k-\theta_{\min})}$, where $\lambda > 0$ and $\theta_{\min} > 0$ are two parameters of the function. A loss interval with this probability density function is actually the sum of a constant θ_{\min} and an exponential random variable with a mean of $1/\lambda$. The mean value and coefficient of variation (CoV) of loss intervals are given by $\mathbb{E}[\theta] = \theta_{\min} + 1/\lambda$, and $\text{CoV}[\theta] = \frac{1/\lambda}{\theta_{\min} + 1/\lambda}$. As described in [23], a desirable feature of function f is that we can vary $\text{CoV}[\theta]$ while keeping the mean and some high-order moments fixed.

First, we study the impact of HERC response function $R(p) = \frac{s}{ru p^d}$ on the upper and lower bounds. We fix c to 0.12, and vary d from 0.5 to 1. Loss history size L is set to 8. We set $\mathbb{E}[\theta]$ to 10^5 (i.e., $p = 10^{-5}$). The upper bound (or the target sending rate) is $R(p) = R(10^{-5})$, and the lower bound is obtained by measuring $\mathbb{E}[\frac{1}{\hat{\theta}}]$ directly in the simulation. Fig. 3 shows the measured lower bounds normalized to the upper bound, and the measured average sending rates also normalized to the upper bound. The left one shows the results with $\text{CoV}[\theta] = 1.0$ and the right one with $\text{CoV}[\theta] = 0.8$.

We have the following observations from Fig. 3. (1) The average sending rate of HERC with basic

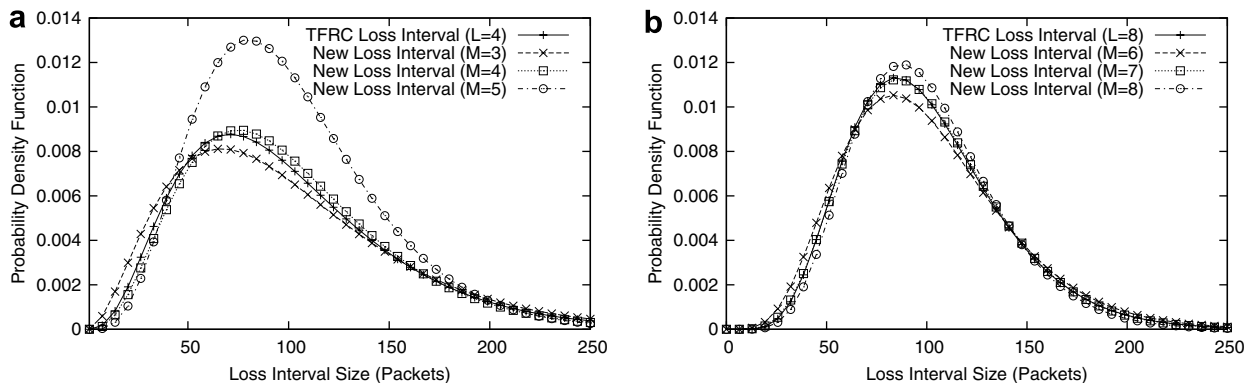


Fig. 2. The accurate probability density function of TFRC weighted average loss interval $\hat{\theta}$ obtained by inverting $L_{\hat{\theta}}(s)$, and that of new weighted average loss interval $\hat{\theta}$ given by Eq. (15). We can see that they are very close to each other when $M = \frac{3}{4}L + 1$. (a) $p = 0.01$, $L = 4$ and $M = 3, 4, 5$, (b) $p = 0.01$, $L = 8$ and $M = 6, 7, 8$.

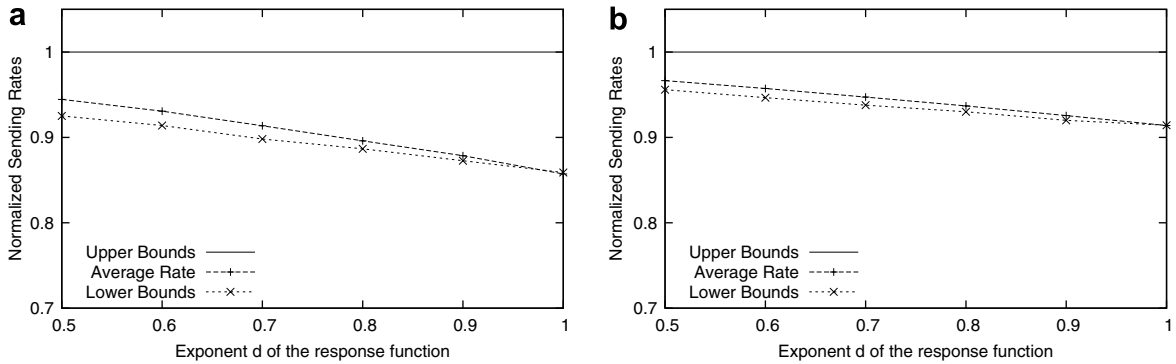


Fig. 3. The average sending rate of HERC with basic control is bounded by its lower and upper bounds. We can see that the larger the exponent d of the response function, the closer the sending rate to its lower bound: (a) $\text{CoV}[\theta] = 1.0$ and (b) $\text{CoV}[\theta] = 0.8$.

control is indeed bounded by the upper and lower bounds. (2) The gap between the upper and lower bounds depends not only on the response function, but also on the distribution of loss intervals. The gap reduces as the variance of loss intervals decreases. In addition to the above two observations that are consistent with previous TFRC analysis [23,18], we have another important observation. (3) The larger the exponent d of the response function, the closer the average rate to the lower bound. Especially, if $d = 1$, the average rate is the same as the lower bound as proved by Theorem 1. Since high-speed TCP variants have a larger d than the standard TCP, this implies that the throughput of HERC with a high-speed TCP response function is closer to its lower bound than that of TFRC with TCP response function to its lower bound.

Next, we study the accuracy of the closed-form expression for the lower bound, when loss intervals are exponentially distributed. In the simulation of Fig. 3a, the loss intervals are exponentially distributed with $p = 10^{-5}$, and then we can obtain the lower bound $R(\mathbb{E}[\frac{1}{\theta}]) = R(\frac{Mp}{M-1}) = R(\frac{7p}{6})$, based on Theorem 2. The closed-form lower bound $R(\frac{7p}{6})$ is shown in Fig. 4. We can see that the closed-form lower bounds closely match the measured lower bounds.

Finally, we evaluate the accuracy of the closed-form lower bound as we change loss history size L . Fig. 5 shows results with the same simulation parameters as Fig. 4, except that d is fixed to 1.0 and L varies from 4 to 64. We can see that there is some difference between the closed-form lower bound and the measured lower bound when $L = 4$, but in most cases, the difference is very small.

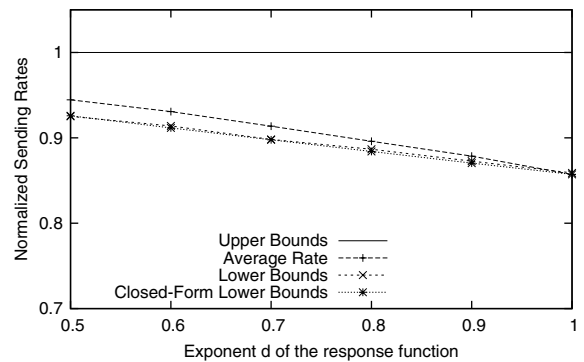


Fig. 4. The closed-form lower bound closely matches the measured lower bound for various response functions.

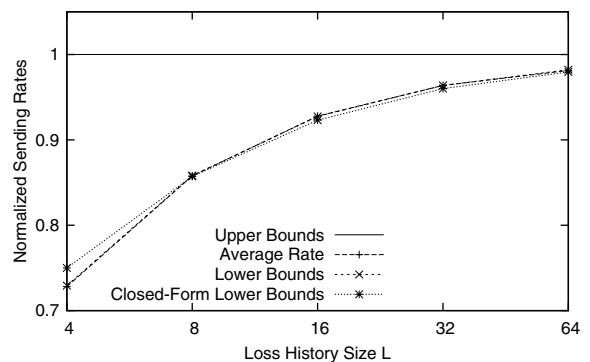


Fig. 5. The closed-form lower bound closely matches the measured lower bound for various loss history sizes.

5. Smoothness analysis

In addition to high throughput, HERC should provide a smooth sending rate, as big rate fluctuation




Fig. 6. HERC achieves higher throughput, and also shows bigger rate fluctuation than TFRC, when they have the same L .

may significantly affect the user-perceived multimedia quality. Fig. 6 shows the throughput curves of a HERC flow with $c = 0.12$, $d = 0.835$, and $L = 8$, and a TFRC flow with the same L , in an NS-2 simulation with a 1 Gbps bottleneck shared among HERC, TFRC, and other background traffic. We can see that HERC achieves higher throughput than TFRC; however, HERC shows larger rate fluctuation than TFRC.

Following the work of [8], we define the smoothness index of a protocol by the coefficient of variation (CoV) of sending rates. $\text{CoV}[S]$ can be obtained as follows:

$$\text{CoV}[S] = \frac{\sqrt{\frac{1}{n-1} \sum_{j=1}^n (S_j - \mathbb{E}[S])^2}}{\mathbb{E}[S]}, \quad (17)$$

where n is the number of time intervals, S_j is the sending rate in time interval j as defined by Eq. (9), and $\mathbb{E}[S]$ is the average sending rate as defined by Eq. (10). $\text{CoV}[S]$ is an indication of the rate smoothness of a protocol at time scale Δt . A smaller $\text{CoV}[S]$ value implies that the sending rate has smaller fluctuation, that is, better smoothness. For example, in Fig. 6, with 1-s measurement intervals, the $\text{CoV}[S]$ of HERC is 0.074, and that of TFRC is 0.037. That is, HERC has worse smoothness than TFRC, when they have the same loss history size.

5.1. Smoothness of HERC

In this subsection, we analytically study the smoothness of HERC with basic control. Our simulation results show that the obtained $\text{CoV}[S]$ of the basic control is also a good approximation to that of the comprehensive control.

We use $\text{CoV}[R]$ as an approximation to the smoothness index $\text{CoV}[S]$, where R_i is the sending rate between the i th and $(i + 1)$ th loss events calculated by using Eq. (8).

$$\text{CoV}[S] \approx \text{CoV}[R]. \quad (18)$$

Note that, only if every loss interval has the same duration, $\text{CoV}[S]$ is exactly equal to $\text{CoV}[R]$. Our simulation results show that in most cases, $\text{CoV}[R]$ is a good approximation to $\text{CoV}[S]$.

Because round-trip delay rtt is an exponentially weighted moving average with a big filter constant 0.9 [9], its variance is typically small. Therefore, it is reasonable to assume that rtt is a constant. Considering that s , c , and d are all constants, $\text{CoV}[R]$ can be obtained as follows, where the last step is based on [21].

$$\begin{aligned} \text{CoV}[R] &= \text{CoV}\left[\frac{s}{rtt} \frac{c}{p^d}\right] = \text{CoV}\left[\frac{1}{p^d}\right] \\ &= \text{CoV}[\hat{\theta}^d] \approx d \cdot \text{CoV}[\hat{\theta}]. \end{aligned} \quad (19)$$

We also assume that θ is independently and identically distributed, and then $\text{CoV}[\hat{\theta}]$ can be calculated as follows:

$$\begin{aligned} \text{CoV}[\hat{\theta}] &= \text{CoV}\left[\sum_{i=1}^L w_i \theta_{i-1}\right] \\ &= \sqrt{\sum_{i=1}^L w_i^2 \text{CoV}[\theta]} \\ &= \sqrt{\frac{32L + 56}{27L^2 + 54L}} \text{CoV}[\theta]. \end{aligned} \quad (20)$$

For any $L \geq 1$, we can approximate $\text{CoV}[\hat{\theta}]$ by

$$\text{CoV}[\hat{\theta}] \approx \frac{1.09}{\sqrt{L}} \text{CoV}[\theta]. \quad (21)$$

Putting Eqs. (18), (19) and (21) together, we get the following theorem.

Theorem 3. For the basic control, the $\text{CoV}[S]$ of HERC with response function $\frac{s}{rtt} \frac{c}{p^d}$ and loss history size L is approximated by $\text{CoV}[S] \approx \frac{1.09d}{\sqrt{L}} \text{CoV}[\theta]$, if the weighted moving average of RTT is a constant and loss intervals are independently and identically distributed.

Therefore, the smoothness of HERC is linearly proportional to d , but independent of c . Typically, a high-speed response function has a larger d than the TCP response function, in order to have better

bandwidth scalability. If HERC has the same L and experience the same $\text{CoV}[\theta]$ as TFRC, but has a larger d , then HERC has worse smoothness than TFRC. This is the reason for the worse smoothness of HERC in Fig. 6, where the d of HERC is larger than that of TFRC (i.e., $0.835 > 0.5$).

Now, we consider a few common loss models used in the literature [28].

- HERC with a deterministic periodic loss model: Under a deterministic periodic loss model, every loss interval has the same number of packets. Therefore $\text{CoV}[\theta] = 0$, and then $\text{CoV}[S] = 0$.
- HERC with a Bernoulli loss model: Under a Bernoulli loss model, a packet is lost with probability p , and loss intervals are geometrically distributed. Thus, $\text{CoV}[\theta] = \sqrt{1-p}$, and then $\text{CoV}[S] = 1.09d\sqrt{\frac{1-p}{L}}$.
- TFRC with a Bernoulli loss model: Considering that the response function of TFRC as described in Eq. (2) is also in the form of $\frac{s}{ru} \frac{c}{p^d}$, Theorem 3 can be applied to TFRC as well. The d of TFRC is 0.5, and the recommended L of TFRC is 8. Therefore, under a Bernoulli loss model, the $\text{CoV}[S]$ of TFRC is $1.09d\sqrt{\frac{1-p}{L}} = 0.193\sqrt{1-p} \approx 0.193$ for small p . This value is close to the numerical result shown in [28].

5.2. Verifying smoothness analysis

We have conducted various NS-2 simulations of HERC with comprehensive control to verify Theorem 3, and the simulation results show that our analysis has a very good accuracy even with comprehensive control.

First, we check the relation between $\text{CoV}[S]$ and $\text{CoV}[\theta]$. In order to accurately control $\text{CoV}[\theta]$, we simulate a single HERC flow over a link with transmission errors. Transmission errors are generated so that loss intervals are IID with probability density function $f_{\theta}(k) = \lambda e^{-\lambda(k-\theta_{\min})}$. We fix $\mathbb{E}[\theta]$ to 10^5 (i.e., the loss event rate is 10^{-5}), and vary $\text{CoV}[\theta]$ from 10^{-7} to 1.0. We set $c = 0.12$, $d = 0.835$ (corresponding to the HSTCP response function [7]), and $L = 8$ (recommended value for TFRC). We measure $\text{CoV}[S]$ by using Eq. (17) with a time interval of 1 s (referred to as simulation result in Fig. 7), and also calculate $\text{CoV}[S]$ by using Theorem 3 (referred to as analytical result in Fig. 7). The results with 95% confidence intervals are shown in Fig. 7. We

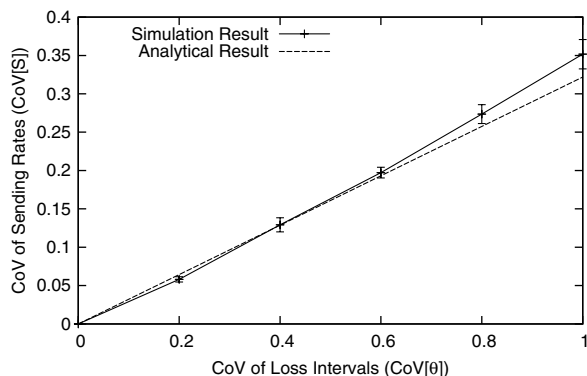


Fig. 7. Smoothness of HERC with $d = 0.835$ and $L = 8$, as $\text{CoV}[\theta]$ increases.

can see that our analysis has good accuracy. We also notice that as $\text{CoV}[\theta]$ increases, the simulation result becomes greater than the analytical result, even though the relative difference is very small ($< 10\%$). We find out that this difference is caused by the comprehensive control of HERC, since Theorem 3 considers only the basic control of HERC.

Next, we test the relation between $\text{CoV}[S]$ and d , as d varies over the suggested range $[0.5, 1]$. We also set $\mathbb{E}[\theta]$ to 10^5 , $\text{CoV}[\theta]$ to 0.8, and L to 8. The simulation result with 95% confidence intervals and the analytical result are shown in Fig. 8. We observe that our analysis has good accuracy. Also we can see that as d increases, the smoothness becomes worse. This is consistent with our finding in Fig. 6.

Finally, we analyze the relation between the smoothness $\text{CoV}[S]$ and the loss history size L . We vary L from 4 to 512. We set $\mathbb{E}[\theta]$ to 10^5 , $\text{CoV}[\theta]$ to 0.8, c to 0.12, and d to 0.835. The simulation result with 95% confidence intervals and the analytical result are shown in Fig. 9. We observe that our analysis has good accuracy for L between 4 and 128. After L becomes greater than 128, our analytical result continues decreasing, however, the simulation result keeps almost constant. This is consistent with the results shown in [24]. We find out that the difference is due to the comprehensive control, which we do not model in the analysis. This means the comprehensive control leads to a minimum $\text{CoV}[S]$, which cannot be mitigated by increasing L .

5.3. Improving smoothness

Theorem 3 shows that if both HERC and TFRC have the same L and $\text{CoV}[\theta]$, and HERC has a larger d than TFRC, then HERC has worse smoothness

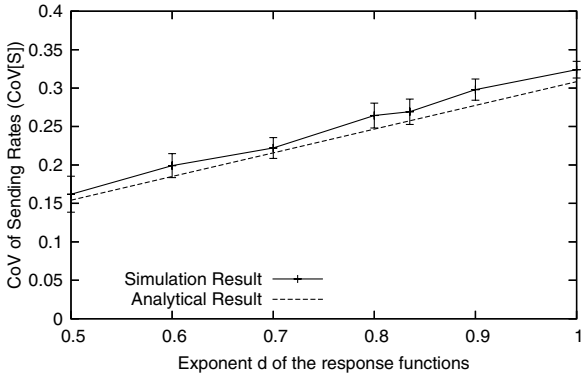


Fig. 8. Smoothness of HERC with $L = 8$ and $\text{CoV}[\theta] = 0.8$, as the exponent d of the response function increases.

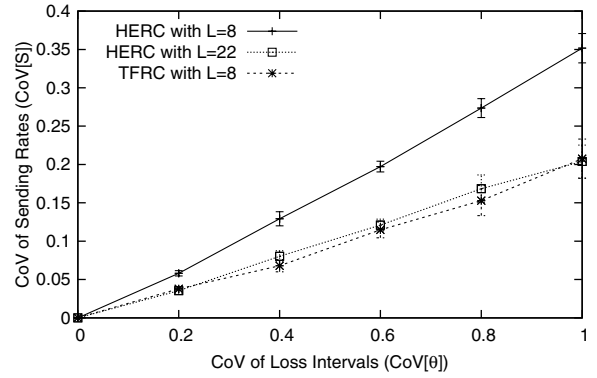


Fig. 10. HERC with $L = 32d^2$ achieves the same smoothness as TFRC with $L = 8$.

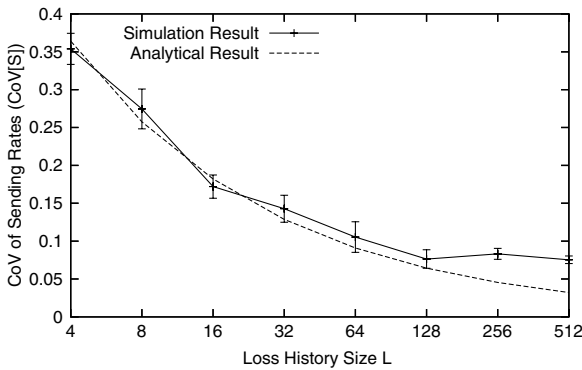


Fig. 9. Smoothness of HERC with the HSTCP response function ($d = 0.835$) and $\text{CoV}[\theta] = 0.8$, as loss history size L increases.

than TFRC. In this subsection, we present a method to improve the smoothness of HERC, so that HERC is able to achieve better bandwidth scalability than TFRC, while maintaining the same smoothness as TFRC under the same loss interval distribution. We define the *smoothness scaling factor* F of a protocol as follows:

$$F = \frac{\text{CoV}[S]}{\text{CoV}[\theta]} = \frac{1.09d}{\sqrt{L}}. \quad (22)$$

If two protocols have the same F , then with the same $\text{CoV}[\theta]$, they have the same smoothness index $\text{CoV}[S]$.

The recommended values for TFRC are $L = 8$ and $d = 0.5$, and then the smoothness scaling factor is $F_{\text{TFRC}} = \frac{1.09 \times 0.5}{\sqrt{8}} = 0.193$. For a HERC flow with response function $\frac{s}{rt} \frac{c}{p^2}$ and loss history size L , the smoothness scaling factor is $F_{\text{HERC}} = \frac{1.09 \times d}{\sqrt{L}}$. By setting $F_{\text{HERC}} = F_{\text{TFRC}}$, we get

$$\frac{1.09 \times d}{\sqrt{L}} = 0.193 \Rightarrow L = 32d^2. \quad (23)$$

Therefore, in order for HERC to achieve the same smoothness as TFRC under the same loss interval distribution, HERC should set its L to $32d^2$. For example, if d is 0.835, then L should be set to 22.

Fig. 10 shows the smoothness of HERC with the HSTCP response function ($d=0.835$) and $L = 8$, HERC with the HSTCP response function ($d=0.835$) and $L = 22$, and TFRC with the TCP response function ($d=0.5$) and $L = 8$. We measure their smoothness as we vary $\text{CoV}[\theta]$ from 10^{-7} to $1 - 10^{-7}$, with $\mathbb{E}[\theta] = 10^5$. We can see that even with different response functions, HERC with $L = 22$ and TFRC with $L = 8$ achieve the same smoothness, while HERC with $L = 8$ has worse smoothness than TFRC.

6. Simulation result

In Sections 4 and 5, we have verified our analysis result by simulation where loss intervals are IID random variables. In this section, we evaluate the performance of HERC by NS-2 simulation with more realistic background traffic.

Fig. 11 shows that our simulation uses a typical dumbbell topology. The one-way delay of the bottleneck link is 50 ms. Each source and sink are con-




Fig. 11. Simulation network topology.

nected to the bottleneck link through different access links with delays randomly varied from 0.1 ms to 0.9 ms to mitigate phase effect. To increase traffic dynamics and further reduce phase effect, various kinds of background traffic are simulated in both directions, including Web traffic, random burst UDP traffic with the Pareto distribution, short-term TCP flows with limited congestion window sizes, and long-lived TCP flows without window constraint. All background traffic consumes a minimum of 25% and 40% of the bottleneck bandwidth in the forward and backward directions, respectively. The packet size is set to 1500 bytes for all connections. We test adaptive RED (Random Early Detectic) routers with suggested parameters for high-speed links [12].

We measure the throughput and smoothness of four forward connections: HSTCP, HERC with HSTCP response function ($d=0.835$) and $L=22$, TCP, and TFRC with TCP response function ($d=0.5$) and $L=8$. We run each simulation for 500 s and take the measurement after the first 250 s. The results are calculated with 95% confidence intervals by repeating the simulation for 10 times. Note that HSTCP could be replaced by other high-speed TCP variants, but the response function and L of HERC should be adjusted accordingly.

Fig. 12 shows the throughput of these four protocols as the bottleneck capacity increases from 100 Mbps to 500 Mbps. We can see that HERC with HSTCP response function achieves similar throughput as HSTCP, and as expected TFRC with TCP response function achieves similar throughput as TCP. Even though the absolute difference between HERC and HSTCP is larger than that between TFRC and TCP, the relative difference between HERC and HSTCP is comparable to that

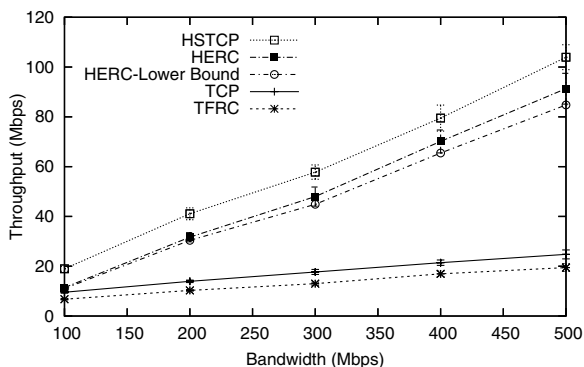


Fig. 12. HERC with the HSTCP response function is reasonably fair with HSTCP, and TFRC is reasonably fair with TCP.

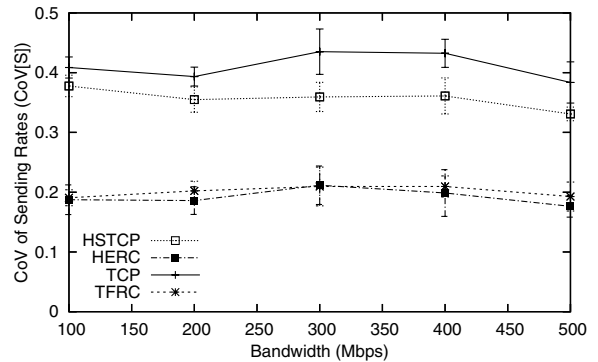


Fig. 13. HERC with the HSTCP response function achieves the same smoothness as TFRC.

between TFRC and TCP. In addition, we measure the loss event rate experienced by HERC in the simulation, and then calculate the closed-form lower bound by using Theorem 2. We can see that the obtained lower bound in this simulation has a good accuracy, even though the loss intervals are not independently and exponentially distributed.

Fig. 13 shows the smoothness of these protocols. We can see that both HERC and TFRC achieve better smoothness than HSTCP and TCP. We also observe that HERC has similar smoothness as TFRC, and this verifies our analysis in Section 5 that by setting $L=32d^2$, HERC can have the same smoothness as TFRC.

Overall, we can see that by using the response function of a high-speed TCP variant and tuning HERC parameters accordingly, HERC can compete fairly with high-speed TCP flows in the same network, while maintaining the same smoothness as TFRC.

7. Related work

Equation-based congestion control, such as TFRC [8], is based on pioneering work by Padhye et al. [16] that models the throughput of TCP using packet loss, delay, and sizes. Most TFRC studies [8,3,28] evaluate the performance of TFRC by simulation, and there are only a small number of analytical studies [23,18,27]. The work [23] by Vojnović and Le Boudec study theoretically the case when TFRC is TCP friendly, and they show that in most cases, the TCP sending rate is the upper bound of the TFRC throughput. Rhee and Xu [18] analytically and experimentally identify the reasons why TFRC may not give the same throughput as TCP, by

examining how the three main factors that determine TFRC throughput, namely, the TCP friendly equation, loss event rate estimation and delay estimation, can influence the long-term throughput imbalance between TFRC and TCP. Xu and Helzer [27] propose an analytical model for TFRC traffic and a queueing model for a TFRC client buffer, in order to study the user-perceived media quality.

8. Conclusion

In this paper, we present HERC as an extension of TFRC for multimedia streaming over high-speed and long-distance networks. By using the response function of a high-speed TCP variant and tuning HERC parameters accordingly, HERC can compete fairly with high-speed TCP flows in the same network, while maintaining the desired smoothness as TFRC. This makes HERC more suitable for applications, such as high-definition video streaming, and remote collaboration involving high-resolution visualization, which prefer a high-speed and relatively smooth sending rate.

In this paper, we analyze the throughput and smoothness of HERC by assuming IID loss intervals, fixed round-trip times, and without considering the impact of the most recent loss interval, and history discounting. Developing more accurate models is of future interest. In addition, we are interested in more realistic Internet experiments to investigate the performance of HERC.

Acknowledgement

We would like to thank Dr. Injong Rhee for his generous help and comments. The work reported in this paper is supported in part by Maude Hammond Fling Faculty Research Fellowship and Layman Fund Award.

References

- [1] Internet 2 and Research Channel Working Group. <<http://www.researchchannel.org/projects/i2wg/index.asp>>.
- [2] Internet HDTV Project. <<http://www.washington.edu/hdtv>>.
- [3] D. Bansal, H. Balakrishnan, S. Floyd, S. Shenker. Dynamic behavior of slowly-responsive congestion control algorithms, in: Proceedings of ACM SIGCOMM, San Diego, CA, August 2001.
- [4] S. Bhandarkar, S. Jain, A.L.N. Reddy, Improving TCP performance in high bandwidth high RTT links using layered congestion control, in: Proceedings of the Third PFLDNet Workshop, France, February 2005.
- [5] H. Bulot, R. Les Cottrell, R. Hughes-Jones, Evaluation of advanced TCP stacks on fast long-distance production networks, in: Proceedings of the Second PFLDNet Workshop, Argonne, Illinois USA, February 2004.
- [6] C. Casetti, M. Gerla, S. Mascolo, M.Y. Sanadidi, R. Wang, TCP Westwood: bandwidth estimation for enhanced transport over wireless links, in: Proceedings of ACM Mobicom, Rome, Italy, July 2001.
- [7] S. Floyd, High speed TCP for large congestion windows, RFC 3649 (December) (2003).
- [8] S. Floyd, M. Handley, J. Padhye, J. Widmer, Equation-based congestion control for unicast applications, in: Proceedings of ACM SIGCOMM, Stockholm, Sweden, August 2000, pp. 43–56.
- [9] S. Floyd, M. Handley, J. Padhye, J. Widmer, TCP friendly rate control (TFRC): protocol specification, RFC 3448 (January) (2003).
- [10] S. Floyd, E. Kohler, TCP friendly rate control (TFRC): the small-packet variant, Internet Draft (March) (2006).
- [11] S. Floyd, E. Kohler, J. Padhye, Profile for DCCP congestion control ID 3: TFRC congestion control, Internet Draft (March) (2005).
- [12] S. Floyd, S. Ratnasamy, S. Shenker, Modifying TCP's congestion control for high speeds, Rough Draft (May) (2002).
- [13] C. Jin, D.X. Wei, S.H. Low, FAST TCP: motivation, architecture, algorithms, performance, in: Proceedings of IEEE INFOCOM, Hong Kong, March 2004.
- [14] T. Kelly, Scalable TCP: improving performance in highspeed wide area networks, ACM SIGCOMM Computer Communication Review 33 (2) (2003) 83–91.
- [15] R. King, R. Riedi, R. Baraniuk. Evaluating and improving TCP-Africa: an adaptive and fair rapid increase rule for scalable TCP, in: Proceedings of the third PFLD Workshop, France, February 2005.
- [16] J. Padhye, V. Firoiu, D. Towsley, J. Kurose. Modeling TCP throughput: a simple model and its empirical validation, in: Proceedings of the ACM SIGCOMM, 1998, pp. 303–314.
- [17] C. Perkins, L. Gharai, T. Lehman, A. Mankin. Experiments with delivery of HDTV over IP networks, in: Proceedings of the 12th International Packet Video Workshop, Pittsburgh, PA, April 2002.
- [18] I. Rhee, L. Xu, Limitations in equation-based congestion control, in: Proceedings of ACM SIGCOMM, Philadelphia, PA, August 2005, pp. 49–60.
- [19] R.N. Shorten, D.J. Leith, H-TCP: TCP for high-speed and long-distance networks, in: Proceedings of the Second PFLDNet Workshop, Argonne, Illinois, February 2004.
- [20] R. Singh, B. Jeong, L. Renambot, A. Johnson, J. Leigh. TeraVision: a distributed, scalable, high resolution graphics streaming system, in: Proceedings of Cluster, San Diego, CA, 2004.
- [21] J.R. Taylor, An Introduction to Error Analysis, University Science Books, Sausalito, California, 1997.
- [22] K. Trivedi, Probability and Statistics with Reliability, Queueing, and Computer Science Applications, John Wiley and Sons, New York, 2001.
- [23] M. Vojnović, J. Le Boudec, On the long run behavior of equation-based rate control, IEEE/ACM Transaction on Networking (2005).

- [24] Z. Wang, S. Banerjee, S. Jamin, Media-friendliness of a slowly-responsive congestion control protocol, in: Proceedings of ACM NOSSDAV, Kinsale, Ireland, June 2004.
- [25] L. Xu. Extending equation-based congestion control to high-speed long-distance networks: smoothness analysis, in: Proceedings of IEEE GLOBECOM, St. Louis, MO, November 2005.
- [26] L. Xu, K. Harfoush, I. Rhee, Binary increase congestion control for fast long-distance networks, in: Proceedings of IEEE INFOCOM, Hong Kong, March 2004.
- [27] L. Xu, J. Helzer. Media streaming via TFRC: an analytical study of the impact of TFRC on user-perceived media quality, in: Proceedings of IEEE INFOCOM, Barcelona, Spain, April 2006.
- [28] Y. Yang, M. Kim, S. Lam, Transient behaviors of TCP-friendly congestion control protocols, *Computer Networks* 41 (2) (2003) 193–210.
- [29] Y. Zhang, N. Duffield, V. Paxson, S. Shenker. On the constancy of Internet path properties, in: Proceedings of

ACM SIGCOMM Internet Measurement Workshop, November 2001.



Lisong Xu received his B.E. and M.E. degrees in Computer Science from the University of Science and Technology Beijing in 1994 and 1997, respectively. He received his Ph.D. degree in Computer Science from North Carolina State University in 2002. From 2002 to 2004, he was a Post-Doctoral Research Fellow at North Carolina State University. He is currently an assistant professor in Computer Science and Engineering at the

University of Nebraska-Lincoln. His research interests include congestion control for multimedia streaming and for fast long-distance networks.