

A NOTE ON THE STOCHASTIC BIAS OF SOME INCREASE-DECREASE CONGESTION CONTROLS: HIGHSPEED TCP CASE STUDY

MILAN VOJNOVIĆ, JEAN-YVES LE BOUDEC, DON TOWSLEY, AND VISHAL MISRA

ABSTRACT. We consider increase-decrease congestion controls, a formulation that accommodates many known congestion controls. There have been many works that aim to obtain relation between the loss event rate p and time-average window \bar{w} for some known particular instances of increase-decrease controls. In contrast, in this note, we study the inverse problem where one is given a target response function $x \rightarrow f(x)$ and the design problem is to construct an increase-decrease control such that, ideally, $\bar{w} = f(p)$, or at least $\bar{w} \leq f(p)$. One common method for solving this is to design a control that satisfies the requirements in a reference system, and then try to evaluate the behavior in a general system. In this note, we consider that the reference is for deterministic constant inter-loss times. Our findings are as follows. We show that for AIMD, in the long-run, determinism minimizes time-average window over the entire set of sequences of inter-loss times with an arbitrary fixed mean. For a broader subset of increase-decrease controls, we identify conditions under which if $\bar{w}' \geq f(p')$ in the reference system (i.e. the control may be non conservative), then for any independent identically distributed (i.i.d.) random inter-loss times, we have $\bar{w} \geq \frac{1}{1+\varepsilon} f(\frac{1}{1+\varepsilon} p)$, for some $\varepsilon \geq 0$ specified in this note. In other words, moving from the reference system to the more general case of i.i.d. losses will not eliminate any non conservativeness. We apply our results to a stochastic fluid version of HighSpeed TCP [6]. We show that for this idealized HighSpeed TCP our result applies with ε not larger than 0.0012. This implies that for idealized HighSpeed TCP \bar{w} is almost lower bounded by $f(p)$ under the hypotheses above. Our general analysis result rises the issue whether it is a good practice to design congestion controls by taking deterministic constant inter-loss times as a reference system, given that we demonstrate that this reference system is, in some sense explained in the paper, in fact a worst-case, rather than a best-case, as would be more desirable.

1. INTRODUCTION

We study the steady-state, in particular the time-average window, of window congestion controls that we call increase-decrease controls. We say a control is increase-decrease if it operates as follows. In absence of loss events, an increase-decrease window control increases the window with the rate that is a function of the current window, else, if a loss event happen, the window is decreased to a function

Key words and phrases. increase-decrease, congestion control, stochastic bias, determinism, extremal property, AIMD, HighSpeed TCP.

M. Vojnović and J.-Y. Le Boudec are with School of Computer and Communications Sciences, EPFL, CH-1015 Lausanne, milan.vojnovic@epfl.ch, jean-yves.leboudec@epfl.ch, D. Towsley is with Department of Computer Science, University of Massachusetts, Amherst MA 01003, towsley@cs.umass.edu, V. Misra is with Department of Computer Science, Columbia University, New York, NY 10027, misra@cs.columbia.edu.

of the window just before the loss event. We claim that this is a formulation that would encompass many known congestion controls. Some particular instances of increase-decrease controls are well-known additive-increase multiplicative-decrease (AIMD) [4], congestion avoidance of TCP [10], and HighSpeed TCP [6], which is taken as a case study in this note.

In this note, we study a stochastic fluid increase-decrease control that is defined in a compact way as: given an initial window $W(0) \geq 0$, the window evolves as

$$(1) \quad W(t) = W(0) + \int_0^t a(W(s))ds - \int_0^t [W(s-) - b(W(s-))]N(ds), \quad t \geq 0.$$

Where $N(t)$ is number of the loss events observed by the sender in the time interval $(0, t]$, $t \in \mathbb{R}_+$, $a : \mathbb{R}_+ \rightarrow (0, \infty)$ is an increase function, $b : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, $0 \leq b(w) < w$, all $w \geq 0$, is a decrease function.¹ In particular, for an AIMD control with some fixed parameters $\alpha > 0$ and $0 < \beta < 1$, $a(w) := \alpha$ and $b(w) := \beta w$, $w \geq 0$.

We consider p and \bar{w} defined as

$$p = \lim_{t \rightarrow \infty} \frac{N(t)}{\int_0^t W(s)ds},$$

and

$$\bar{w} = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t W(s)ds.$$

The quantity p is the long-run loss event ratio, \bar{w} is the *time-average* window. We also consider an *event-average* window \bar{w}_0 , the event-average of the window sampled just after the window reduction instants,

$$\bar{w}_0 = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t W(s)N(ds).$$

An implicit assumption in the above limits is that the limits exist. This is an issue in its own right. One would need to prove that the system defined by (1) is stable, that is that exists unique stationary ergodic limit, see [3] for a textbook account of this matter. This is out of the scope of this note, we take the stability as a premise.

Many works have studied the relation $p \rightarrow \bar{w}$ of the loss event rate p and *time-average* window \bar{w} for some particular instants of increase-decrease controls; mostly AIMD, e.g. [1, 5, 12, 13, 14]. In this context, the control is known, that is the functions $w \rightarrow a(w)$ and $w \rightarrow b(w)$ are known, and the goal is to obtain $p \rightarrow \bar{w}$. We call this the *direct problem*. Our results are related to an *inverse problem*². We are given the target response function $p \rightarrow f(p)$ and aim to design an increase-decrease control, that is to identify the functions $w \rightarrow a(w)$ and $w \rightarrow b(w)$, such that, ideally, $\bar{w} = f(p)$.

The design goal $\bar{w} = f(p)$ may be feasible to achieve for a given loss event process, but it would not be possible to fix some $w \rightarrow a(w)$ and $w \rightarrow b(w)$ such that $\bar{w} = f(p)$ for *all* loss event processes. A design method is to solve the inverse problem for a reference system of loss events. Often, the reference system is taken to be a deterministic system with inter-loss times fixed to a constant; we call it *deterministic constant inter-loss times*. This reference system is simple to study, which

¹Technicality note, it is implicit above that $t \rightarrow W(t)$ is right-continuous with left-hand limits, and for any $t \in \mathbb{R}_+$, $N(dt) \in \{0, 1\}$, that is the point process of loss events is simple.

²Two problems are said to be inverses of one another if the formulation of each involves all or part of the solution of the other [11].

may have been the primary reason of its appearance in many other works. The generic design method can be summarized as follows: one first chooses a reference system of loss events, and then, for a given function $p \rightarrow f(p)$, finds $w \rightarrow a(w)$ and $w \rightarrow b(w)$ such that $\bar{w}' = f(p')$, where \bar{w}' and p' are, respectively, the time-average window and loss event rate attained under the reference system.

In reality, the inter-loss times are never fixed, but variable. By common sense, one would like that the time-average window of a control does not excessively deviate from its original target response function. If we know that $\bar{w}' \leq f(p')$ for a reference system, then there is no reason to expect that the same inequality would be preserved for an arbitrary system, that is, that it would follow $\bar{w} \leq f(p)$. We would like to understand are there conditions under which if $\bar{w}' \leq f(p')$ (resp. $\bar{w}' \geq f(p')$), then we can conclude that $\bar{w} \leq f(p)$ (resp. $\bar{w} \geq f(p)$). In other words, we aim to understand the way the stochastic bias would act.

Summary of our Results. We show that for an AIMD control, in the long-run, the minimum time-average window over the entire set of inter-loss times with an arbitrary fixed mean $1/\lambda$, is attained with the inter-loss times fixed to $1/\lambda$. The result establishes an *extremal*³ property of the deterministic constant inter-loss times. The extremal property is in the sense that the time-average window is minimized; we call it the *worst-case*. The result is established in a sample-path framework, for any sequence of inter-loss times with an arbitrary fixed mean $1/\lambda$.

The last result is confined to AIMD controls, but under a mild assumption on the process of loss events. We show a result that holds for a broader set of increase-decrease controls, but a smaller set of loss event processes. We identify conditions on the increase and decrease functions, and a condition on the target response function, under which, if $\bar{w}' \geq f(p')$ in the reference system, then for any independent and identically distributed (i.i.d.) random inter-loss times, $\bar{w} \geq \frac{1}{1+\varepsilon} f(\frac{1}{1+\varepsilon} p)$, for some $\varepsilon \geq 0$ that depends on a functional of $w \rightarrow a(w)$. If $\varepsilon = 0$, then the result tells us that for any i.i.d. random inter-loss times, the time-average window \bar{w} is lower-bounded by $f(p)$. If, in contrast, ε is strictly positive, but small, then we expect \bar{w} would be almost lower bounded by $f(p)$. The last result is obtained as a conjunction of two results discussed next.

Under the aforementioned conditions on the increase and decrease functions, for an increase-decrease control, the time-average window \bar{w} attained for any i.i.d. random inter-loss event times with mean $1/\lambda$ and the time-average window \bar{w}' attained under the reference system are related as $\bar{w} \geq \frac{1}{1+\varepsilon} \bar{w}'$. For all increase-decrease controls for which $\varepsilon = 0$, the last result establishes that the deterministic constant inter-loss times are *extremal*, more specifically, the *worst-case*, over the set of i.i.d. random inter-loss times. If $\varepsilon > 0$, but small, than the deterministic constant inter-loss times are *almost* the worst-case.

If the control is designed such that it verifies $\bar{w}' \geq f(p')$ under the reference system, and under the additional condition on the target response function referred to earlier, from the last above result, we can conclude $\bar{w} \geq \frac{1}{1+\varepsilon} f(\frac{1}{1+\varepsilon} p)$.

³An extremal problem (see [15], Sec. 2.3) in our context can be thought as a variational problem, where given the time-average window $\bar{w}(F)$ attained by i.i.d. random inter-loss times with distribution F , for any fixed $0 < \lambda < \infty$, we solve $\inf_{F \in \mathcal{D}} \bar{w}(F)$ over the set \mathcal{D} of distributions on \mathbb{R}_+ such that $\int_0^\infty (1 - F(x)) dx = 1/\lambda$. The result in this note tells us that the infimum is attained for $F(x) = 1_{x \geq 1/\lambda}$, that is for the reference system with inter-loss times fixed to $1/\lambda$.

As a by-product of our analysis, under the present assumptions on the increase and decrease functions, we show that $\bar{w}_0 \geq \frac{1}{1+\varepsilon} \bar{w}'_0$, for some $\varepsilon \geq 0$ defined earlier, over the entire set of stationary random inter-loss times.

We can rephrase our finding as follows. Assume $\varepsilon = 0$. If one observes the time-average window \bar{w} of an increase-decrease control, which verifies conditions identified in this note, and it is driven by any i.i.d. random inter-loss times with mean $1/\lambda$, then we can conclude that $\bar{w} \geq \bar{w}'$, where \bar{w}' is the time-average window that would be attained by the same control if we fix the inter-loss times to the mean $1/\lambda$. This is an effect of the stochastic bias. Now, in turn, given that we designed our control at the first place such that $\bar{w}' = f(p')$ (or it turns out to hold, $\bar{w}' \geq f(p')$), and given that it happens that the target response function f is such that $x \rightarrow xf(x)$ is non-decreasing, then we can conclude that $\bar{w} \geq f(p)$. The observation that we make here rises an issue:

Is it in general a good practice to design increase-decrease controls by taking as a reference system the one with deterministic constant inter-loss times, given that we demonstrate that there exist elements under which this reference is a worst-case?

We apply our analysis results to HighSpeed TCP proposed in [6]; see also some complementary discussions in [7]. HighSpeed TCP can be taught of as a special instant of increase-decrease controls. Moreover, the design problem found in [6] can be seen as an instance of the generic design method posed in this note. We study HighSpeed TCP by instancing the functions $w \rightarrow a(w)$ and $w \rightarrow b(w)$ in (1) to those defined in [6]. We refer to this system as *idealized* HighSpeed TCP.

Our results are as follows. We show that for idealized HighSpeed TCP our main analysis result discussed above holds for some $\varepsilon \in (0, 0.0012)$. In fact, an implication is that for idealized HighSpeed TCP and any i.i.d. random inter-loss times, it always holds $\bar{w} \geq (1 - \varepsilon'')f(p)$, for some $\varepsilon'' \in (0, 0.0023)$. Hence, \bar{w} is almost lower bounded by $f(p)$. As an aside result, we show that for idealized HighSpeed TCP and any fixed inter-loss times, it holds $f(p') \leq \bar{w}' \leq (1 + \varepsilon')f(p')$, for some $\varepsilon' \in (0, 0.06)$. The maximum deviation happens to be at the knee point of HighSpeed TCP response function. It remains open to evaluate how much the time-average window \bar{w} of HighSpeed TCP would deviate from its target $f(p)$. We conjecture that the stochastic bias would tend to make this deviation larger.

The remainder of the note is organized as follows. In Section 2, we give some further preliminaries. Our main results are shown in Section 3. In Section 4 we display some concluding remarks. Most of the proofs are deferred to Appendix.

2. ADDITIONAL NOTATIONS AND ASSUMPTIONS

We first assume a bit more about N , the counting process of loss events. Assume N to be with non-null finite intensity λ , that is for any $t \geq 0$, $\mathbb{E}[N(t)] = \lambda t$. Let $T_0 \leq 0 < T_1 < T_2 < \dots$ be a sequence of loss event instants. Define $S_n = T_{n+1} - T_n$, $n = 0, 1, 2, \dots$ to be sequence of inter-loss times. We will find handy to use this additional notation, $W_n = W(T_n)$, $n = 0, 1, 2, \dots$, which is a sequence of windows embedded at the loss events.

It is readily checked from (1) that for any $n = 0, 1, 2, \dots$,

$$W(t) = \phi^{-1}(\phi(W_n) + t), \quad t \in [T_n, T_{n+1}),$$

where ϕ is a primitive of $1/a$, that is

$$\int \frac{dx}{a(x)} = \phi(x) + C,$$

for an arbitrary constant $C \in \mathbb{R}$, and ϕ^{-1} is the inverse of ϕ . The formula above, given the embedded windows, defines window dynamics in-between successive loss events. Now, it is an easy step to note that for any $n = 0, 1, 2, \dots$,

$$(2) \quad W_{n+1} = b(\phi^{-1}(\phi(W_n) + S_n)).$$

The last formula is a stochastic recursive sequence⁴ that defines the governing dynamics of the windows embedded at the loss events.

We now formulate the direct and inverse problem for the reference system of deterministic constant inter-loss times to which we alluded insofar in the text at many instances.

Direct Problem: given $a(\cdot)$ and $b(\cdot)$, find p and \bar{w} that solve:

$$(3) \quad \begin{aligned} \bar{w}_0 &= b(\phi^{-1}(\phi(\bar{w}_0) + 1/\lambda)) \\ \frac{1}{p} &= \int_0^{1/\lambda} \phi^{-1}(\phi(\bar{w}_0) + s) ds. \end{aligned}$$

We have 2 equations and 3 unknowns \bar{w}_0 , λ , and p . We can always take λ as a parameter, solve for p and \bar{w}_0 . It is only left to observe that $\bar{w} = \lambda/p$, hence, we obtain a parametric solution for any fixed parameter λ .

Inverse Problem: given $f(\cdot)$, find $a(\cdot)$, $b(\cdot)$, \bar{w}_0 , λ , and p that solve:

$$(4) \quad \begin{aligned} \bar{w}_0 &= b(\phi^{-1}(\phi(\bar{w}_0) + 1/\lambda)) \\ \frac{1}{p} &= \int_0^{1/\lambda} \phi^{-1}(\phi(\bar{w}_0) + s) ds \\ \frac{\lambda}{p} &= f(p). \end{aligned}$$

The system is under-dimensioned, we have 3 equations, 5 unknowns λ , p , \bar{w}_0 , $a(\cdot)$, and $b(\cdot)$. With 2 left degrees of freedom, for a fixed λ , one may fix $w \rightarrow a(w)$ (resp. $w \rightarrow b(w)$), and then, at least in principle, solve $w \rightarrow a(w)$ (resp. $w \rightarrow b(w)$).

The design problem in HighSpeed TCP context [6] can be seen as solving the inverse problem above. There, for a fixed $\lambda > 0$ and $w \rightarrow b(w)$, one is left with 3 equations and 3 unknowns, a well-posed problem. A difficulty arises with solving the inverse problem because it requires to solve a functional with respect to $\phi^{-1}(\cdot)$. Knowing $\phi^{-1}(\cdot)$ is in principle equivalent to knowing $a(\cdot)$ up to an additive constant.

In the remainder of this section, we introduce some aspects of HighSpeed TCP that are of interest within the scope of this note. The construction of the control found in the context of HighSpeed TCP can be seen as applying the generic design method that we defined earlier, but solving the inverse problem approximately, as we outline next. In [6], the author imposes the target response function

$$f(p) = \frac{K(p)}{p^{\gamma(p)}},$$

where $K(p) = \sqrt{3/2}$, $\gamma(p) = 1/2$, for $p \geq p_*$, $p_* := 3/2/w_*^2$, $w_* := 38$, else $K(p) = K_1$ and $\gamma(p) = \gamma_1$ ($K_1 \simeq 0.12$, $\gamma_1 \simeq 1/1.2$).

⁴It is a special stochastic recursive sequence known as generalized autoregression [3]. In this note, we do not further exploit this observation, but display it here for a mathematically inclined reader.

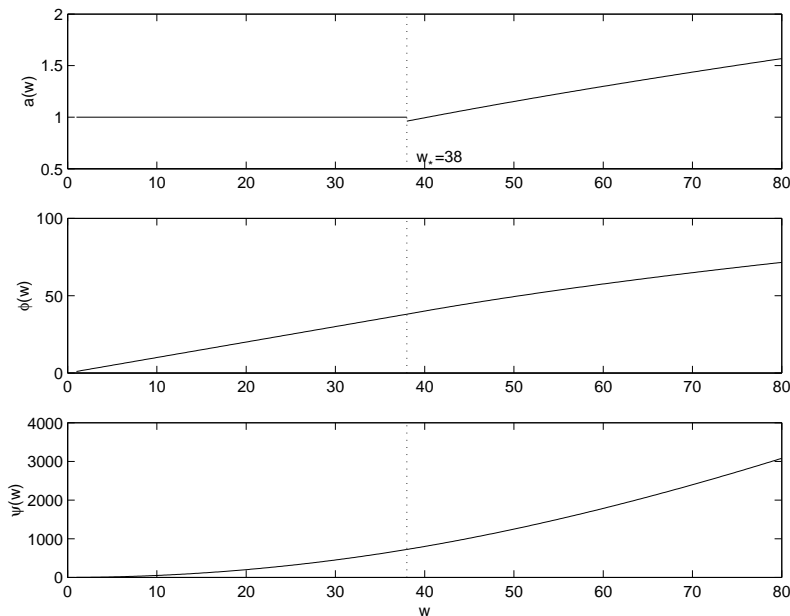


FIGURE 1. (Top) Function a , (Middle) ϕ (primitive of $1/a$), and (Bottom) ψ (primitive of ϕ) for HighSpeed TCP.

The decrease function in [6] is fixed to, $b(w) = w/2$, $w < w_*$, else,

$$b(w) = (1 - \beta(w))w, \quad \beta(w) := c \ln w + d, \quad w \geq w_*$$

where c and d are some constants specified in [6] ($c \simeq -0.052$, $d \simeq 0.689$).

Then, by some arguments found in [6], the increase function is derived to be, $a(w) = 1$, $w < w_*$, else

$$a(w) = K_1^{1/\gamma_1} w^{2-1/\gamma_1} \frac{2\beta(w)}{2 - \beta(w)}, \quad w \geq w_*$$

The argument used in [6] to derive $w \rightarrow a(w)$ for the fixed $w \rightarrow b(w)$ and $p \rightarrow f(p)$ can be seen as approximately solving the inverse problem posed above by assuming that $a(w)$ and $\beta(w)$ are almost constants, that is, that the control is almost an AIMD. The argument in [6] would yield an exact solution if $a(\cdot)$ and $\beta(\cdot)$ would be functions of the time-average window, hence, constants, for a given steady-state. If $w \rightarrow a(w)$ and $w \rightarrow \beta(w)$ are slowly-varying, that is almost constants on the time-scale of the control, then, accuracy of the approximation would be good. We make no attempt to solve the inverse problem. Rather, we take $a(\cdot)$ and $b(\cdot)$ as defined in [6] and solve the direct problem, that is compute $p' \rightarrow \bar{w}'$. We then compare \bar{w}' with the original target response function $f(p')$. If a would be an exact solution then ideally we would find a match between \bar{w}' and f . We show later that this is not the case.

In Figure 1 we show plots of $w \rightarrow a(w)$, $w \rightarrow \phi(w)$ and $w \rightarrow \psi(w)$, the last function is a primitive of $w \rightarrow \phi(w)$. We show analytical expressions for the last two functions in Appendix E.




FIGURE 2. (Left) The solid line is the minimum time-average window obtained by numerically solving the problem $(P_{m,\lambda,0})$ for fixed $\lambda = 1$. The dotted line is the time-average window attained for the deterministic constant inter-loss times with $\lambda = 1$. (Right) The same data as on the left-side, but plotted as the relative difference of the minimum time-average window and the reference time-average window of the deterministic constant inter-loss times.

3. MAIN RESULTS

An Extremal Property of AIMD. We show an extremal property that holds for AIMD under very mild assumptions. Consider an AIMD control with parameters $\alpha > 0$ and $0 < \beta < 1$. Locally to this section, we work in a sample-path framework. Let, for some fixed m , $(s_0, s_1, s_2, \dots, s_{m-1})$ be a sequence of non-negative inter-loss times subject to the only constraint that $\frac{1}{m} \sum_{n=0}^{m-1} s_n = 1/\lambda$, $\lambda > 0$. Let $w_0 \geq 0$ be initial window. We define

$$h_{\lambda,w_0}(s_1, s_2, \dots, s_{m-1}) = \frac{\lambda w_0}{m} \sum_{n=0}^{m-1} \beta^n s_n + \frac{\lambda \alpha}{2m} \sum_{n=0}^{m-1} s_n^2 + \frac{\lambda \alpha}{m} \sum_{n=1}^{m-1} \sum_{k=0}^{n-1} \beta^{n-k} s_k s_n.$$

One can readily check that $h_{\lambda,w_0}(s_1, s_2, \dots, s_{m-1})$ is the time-average window attained by the AIMD control with initial window w_0 and the sequence of inter-loss times s_1, s_2, \dots, s_{m-1} .

Now, imagine an adversarial whose goal is to beat down the time-average window of our AIMD source by choosing freely any sequence of the inter-loss times subject to the only constraint that the sequence has arithmetic mean $1/\lambda$. The adversary would solve the quadratic constrained optimization problem, for some fixed m , $\lambda > 0$, $w_0 \geq 0$:

$$\begin{aligned} & \mathbf{P}_{m,\lambda,w_0}: \\ & \text{minimize} \quad h_{\lambda,w_0}(s_0, s_1, \dots, s_{m-1}) \\ & \text{subject to} \quad s_0 \geq 0, s_1 \geq 0, \dots, s_{m-1} \geq 0 \\ & \quad \quad \quad \frac{1}{m} \sum_{n=0}^{m-1} s_n = \frac{1}{\lambda}. \end{aligned} \tag{5}$$

We define $H_m^{\lambda,w_0} = h_{\lambda,w_0}(s_0^*, s_1^*, \dots, s_{m-1}^*)$, where $s_0^*, s_1^*, \dots, s_{m-1}^*$ is a solution of the problem (P_{m,λ,w_0}) .

Now, let us redefine the goal of our adversary, and consider that we would like to obtain a lower bound on H_m^{λ, w_0} under the constraints in (P_{m, λ, w_0}) . An elegant way to obtain the lower bound is to consider the special case $(P_{m, \lambda, 0})$. In other words, we consider (P_{m, λ, w_0}) for the zero initial window. Note that, for any $\lambda > 0$, $w_0 \geq 0$, it indeed holds

$$H_m^{\lambda, w_0} \geq H_m^{\lambda, 0}, \text{ any } m = 1, 2, \dots$$

In general, the problem (P_{m, λ, w_0}) can be solved by the method of Lagrange multipliers. For $(P_{m, \lambda, 0})$, it amounts to solving the system of linear equations

$$(6) \quad \sum_{k=0}^{n-1} \beta^{n-k} s_k + s_n + \sum_{k=n+1}^{m-1} \beta^{k-n} s_k = \frac{\gamma}{\alpha}, \quad n = 0, 1, \dots, m-1,$$

where γ is the Lagrange multiplier. Now, one can readily note that the matrix of the system (6), for $0 < \beta < 1$, has rank m , the augmented matrix has the same rank, thus there exists unique solution that is displayed next

$$s_0 = s_{m-1} = \frac{1}{\lambda} \frac{m}{2 + (1 - \beta)(m - 2)}, \quad s_k = (1 - \beta)s_0, \text{ else.}$$

The solution is clearly non-negative, hence, the simple method of Lagrange multipliers provides the solution. By plugging this result into $h_m^{\lambda, 0}(\cdot)$, we obtain

$$(7) \quad H_m^{\lambda, 0} = \frac{\alpha}{2\lambda} \frac{(1 + \beta)m}{2\beta + (1 - \beta)m}.$$

We in fact have shown the following result:

Lemma 1. $H_m^{\lambda, 0}$ defined by (7) is a lower bound on the objective function in (P_{m, λ, w_0}) , for any m , $\lambda > 0$ and $w_0 \geq 0$.

In other words, the time-average window attained by an AIMD source, with an arbitrary fixed initial window $w_0 \geq 0$, and any non-negative sequence of inter-loss times with mean $1/\lambda$, cannot be smaller than $H_m^{\lambda, 0}$.

Now, let us define $s_m^\lambda := (1/\lambda, 1/\lambda, \dots, 1/\lambda)$ of length m , which corresponds to the deterministic constant inter-loss times. We obtain

$$h_{\lambda, w_0}(s_m^\lambda) = \frac{\alpha}{2\lambda} \frac{1 + \beta}{1 - \beta} + \left(w_0 - \frac{\alpha\beta}{\lambda(1 - \beta)} \right) \frac{1 - \beta^m}{(1 - \beta)m}.$$

From the above computations, we directly obtain the following result.

Theorem 1. For an AIMD control with any fixed $\lambda > 0$ and $w_0 \in [0, \infty)$, it holds

$$\frac{h_{\lambda, w_0}(s_m^\lambda)}{H_m^{\lambda, 0}} \downarrow 1, \text{ as } m \rightarrow \infty.$$

In other words, the lower bound $H_m^{\lambda, 0}$ is asymptotically tight; it is asymptotically attained by the sequence of inter-loss times fixed to $1/\lambda$.

The theorem tells us that, in the long-run, the sequence of deterministic constant inter-loss times is *extremal*. In fact, it is a *worst-case*, that is, in the long-run, it attains the minimum possible time-average window over the entire set of non-negative sequences of inter-loss times with arithmetic mean $1/\lambda$. We show a pictorial illustration of the theorem in Figure 2.




FIGURE 3. The curve is the time-average window attained for period-two loss events with the mean $1/\lambda$, which is divided by the reference time-average window of deterministic constant inter-loss times fixed to $1/\lambda$, versus the parameter η .

Remark 1. *The above result may perhaps appear to be suggestive from the exact time-average send rate obtained in an enlightening work [1], Proposition 2 therein. The results in this reference are obtained in a stationary and ergodic framework. It indeed follows from the result in [1] that the deterministic constant inter-loss times is the worst-case over the set of i.i.d. random inter-loss times. However, for more general, stationary ergodic inter-loss times, the conclusion of the above theorem does not seem to directly follow from [1].*

We demonstrate our result with an example, period-two loss events, which, apart from the deterministic constant inter-loss times, may be the simplest deterministic case to consider.

Example 1 (period-two loss events). *Consider the sequence of inter-loss times defined as, for some fixed $0 \leq \eta \leq 1$, and an even m ,*

$$s_n = \frac{2}{\lambda}(\eta + (1 - 2\eta)1_{\{n \text{ not even}\}}), \quad n = 0, 1, \dots, m - 1.$$

In perhaps simpler terms, an equivalent description is

$$s_m^{\lambda, \eta} := (2\eta/\lambda, 2(1 - \eta)/\lambda, 2\eta/\lambda, \dots) \text{ of length } m.$$

It requires a little effort to compute

$$h_{\lambda, w_0}(s_{\infty}^{\lambda, \eta}) = \frac{\alpha}{\lambda(1 - \beta^2)} [1 + \beta^2 - 2(1 - \beta(2 - \beta))\eta(1 - \eta)].$$

It is readily seen that $h_{\lambda, w_0}(s_{\infty}^{\lambda, \eta}) \geq h_{\lambda, w_0}(s_{\infty}^{\lambda})$, where recall s_{∞}^{λ} is the infinite sequence of inter-loss times fixed to $1/\lambda$. Moreover, the equality is attained for $\eta = 1/2$, at which $s_{\infty}^{\lambda, \eta}$ degenerates to s_{∞}^{λ} .

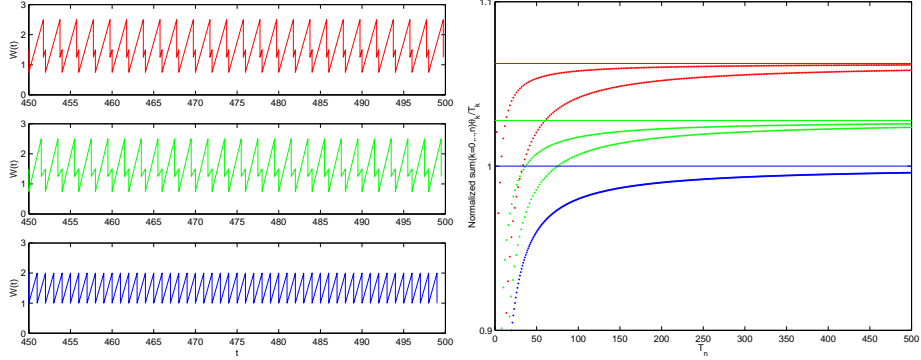


FIGURE 4. (Left) The window evolution of the period-two example with $w_0 = 0$. (Right) the time-average window evaluated at the loss events; the solid lines are asymptotic values obtained by analysis. The curves are, from the top to bottom, for $\eta = 1/8, 1/4, 1/2$. The numerical example confirms the extremal property of the deterministic constant inter-loss times.

In Figure 3 we show a plot of the ratio between the time-average window for period-two inter-loss times with the mean $1/\lambda$ and the time-average window attained for the inter-loss times fixed to $1/\lambda$. The result nicely illustrates that, $\eta = 1/2$, that is the deterministic inter-loss times, is the worst-case. The maximum possible overshoot is

$$2 \frac{1 + \beta^2}{(1 + \beta)^2}$$

which for $\beta = 1/2$ amounts to $10/9 = 1.11\bar{1}$, hence not more than 12%. We give a further numerical example in Figure 4.

Related Results for More General Increase-Decrease Controls. We consider a subset of increase-decrease controls that verify the following assumptions.

(A1): $x \rightarrow \phi(b(\phi^{-1}(x)))$ is increasing convex.

(A2): There exists a convex function $x \rightarrow \varphi(x)$ such that, for some $\varepsilon \geq 0$,

$$\varphi(x) \leq \phi^{-1}(x) \leq (1 + \varepsilon)\varphi(x), \text{ all } x \geq 0.$$

(A3): $x \rightarrow xf(x)$ is non-decreasing.

The assumptions (A1) and (A2) define a continuous subset of increase-decrease controls. We show next our main analysis result.

Theorem 2. Consider any increase-decrease control that obeys to (A1) and (A2). If the control is designed such that for deterministic constant inter-loss times the time-average window \bar{w}' and loss event rate p' are related as $\bar{w}' \geq f(p')$, and if $x \rightarrow f(x)$ obeys to (A3), then for any i.i.d. random inter-loss times $\bar{w} \geq \frac{1}{1+\varepsilon} f(\frac{1}{1+\varepsilon} p)$.

Remark 2. An obvious implication from the statement of the theorem is

$$\bar{w} \geq \left(\frac{1}{1 + \varepsilon} \inf_{x \in [0,1]} \frac{f(\frac{1}{1+\varepsilon}x)}{f(x)} \right) f(p).$$

The theorem follows as a conjunction of two lemmas that we show next.

Lemma 2. *Consider any increase-decrease control that obeys to (A1) and (A2). For any $0 < \lambda < \infty$, the time-average window \bar{w} attained by the control with any i.i.d. random inter-loss times with mean $1/\lambda$, and the time-average window \bar{w}' attained by the control with inter-loss times fixed to $1/\lambda$, are related as $\bar{w} \geq \frac{1}{1+\varepsilon} \bar{w}'$.*

The result implies that, if we know that $\bar{w}' \geq f(p')$, then we know that under the assumptions of Lemma 2, we have $\bar{w} \geq \frac{1}{1+\varepsilon} f(p')$. It still remains to show when the last implies $\bar{w} \geq \frac{1}{1+\varepsilon} f(\frac{1}{1+\varepsilon} p)$.

Lemma 3. *If $\bar{w} \geq \frac{1}{1+\varepsilon} \bar{w}'$, $\bar{w}' \geq f(p')$, and $x \rightarrow xf(x)$ is non-decreasing, then $\bar{w} \geq \frac{1}{1+\varepsilon} f(\frac{1}{1+\varepsilon} p)$.*

Remark 3. *The hypothesis $x \rightarrow xf(x)$ non-decreasing is verified for many known loss-throughput functions f , but not all. It is indeed verified by $f(p) = K/p^\gamma$, $K > 0$, for any $\gamma \leq 1$, which encompasses the square-root function with $\gamma = 1/2$. However, one may check that the hypothesis is not verified by a loss-throughput formula in [13].*

We next consider another assumption,

(A2'): $x \rightarrow a(x)$ is non-decreasing.

Note that (A2) is weaker than (A2'). To see this note that $x \rightarrow \phi^{-1}(x)$ is convex iff $x \rightarrow a(x)$ is non-decreasing. Hence, if (A2') holds, then (A2) indeed holds with $\varepsilon = 0$. Hence, we have the following corollary from Theorem 2.

Corollary 1. *Replace (A2) in Theorem 2 with (A2'), then with all the rest remaining unchanged, the statement of the theorem reads as $\bar{w} \geq f(p)$.*

Note that under (A2') one can analogously adapt Lemma 2 and Lemma 3.

From the proof of Theorem 2 we obtain almost for free the following result, which may have an interest in its own right. Note that the next theorem applies for any stationary random inter-loss times, hence, for a weaker hypothesis than in Theorem 2.

Theorem 3. *Consider any increase-decrease control that obeys to (A1) and (A2). Then, the window event-average \bar{w}_0 attained by the control with any stationary random inter-loss times with mean $1/\lambda$, and the window event-average \bar{w}'_0 attained by the control with inter-loss times fixed to $1/\lambda$ are related as $\bar{w}_0 \geq \frac{1}{1+\varepsilon} \bar{w}'_0$.*

Corollary 2. *Replace (A2) in Theorem 3 with (A2'), then the statement of the theorem would read as $\bar{w}_0 \geq \bar{w}'_0$.*

Application to HighSpeed TCP. We show in Appendix D that idealized High-Speed TCP verifies the hypotheses of Theorem 2. Given that the verification is done by numerical computations, we pose the result as a claim.

Claim 1. *For idealized HighSpeed TCP, the statement of Theorem 2 is true for some $\varepsilon \in (0, 0.0012)$. Moreover, for any i.i.d. random inter-loss times, it holds $\bar{w} \geq (1 - \varepsilon'')f(p)$, for some $\varepsilon'' \in (0, 0.0023)$.*

The above result is based on the following fact concluded by numerical computations.




FIGURE 5. (Top) $p' \rightarrow \bar{w}'$ for HighSpeed TCP, the dotted line is the target response function f , (Bottom) relative difference of \bar{w}' from the target f , that is $(\bar{w}' - f)/f$; we observe that \bar{w}' is never smaller than f , and never larger than f by the factor 1.006.

Claim 2. *The time-average window \bar{w}' and loss event rate p' attained by idealized HighSpeed TCP with fixed inter-loss times, are related as $f(p') \leq \bar{w}' \leq (1 + \varepsilon')f(p')$, for some $\varepsilon' \in (0, 0.06)$. See Fig. 5.*

4. CONCLUDING REMARKS

The extremal property of the deterministic reference system pointed out in this note may not come as a surprise, in particular, to those familiar with a folk theorem of queueing theory that says that under some conditions determinism minimizes waiting time in queues, e.g. see [9], [2] (Chapter 4).

We showed in this note that, in the long-run, determinism minimizes time-average window of AIMD controls over the entire set of non-negative sequences of inter-loss times with an arbitrary fixed mean. We identified a broader subset of increase-decrease controls for which determinism minimizes or almost minimizes the time-average window over a set of i.i.d. random inter-loss times.

We next point to some further examples in the context of network congestion controls, where determinism is an extremal. Thus far in this note, we assumed that

the point process of loss events is simple, that is at any instant of time there exists at most one loss event. Consider, now, an AIMD control with parameters α and β , but with loss events that may arrive in batches. Let $\dots, Z_{n-1}, Z_n, Z_{n+1}, \dots$ be a stationary sequence of batch sizes. It is a straightforward exercise to extend the result in [1] to the system with batch loss events

$$\bar{w} = \alpha\lambda \left(\frac{1}{2}\mathbb{E}[S_0^2] + \sum_{k=1}^{\infty} \mathbb{E}[\beta^{Z_{-1}+Z_{-2}+\dots+Z_{-k}} S_0 S_{-k}] \right).$$

Now, if we assume that the sequence of batch sizes $\dots, Z_{n-1}, Z_n, Z_{n+1}, \dots$ is independent of the point process of batch arrivals, then by convexity of $x \rightarrow \beta^x$, it follows

$$\bar{w} \geq \alpha\lambda \left(\frac{1}{2}\mathbb{E}[S_0^2] + \sum_{k=1}^{\infty} \beta^{k\mathbb{E}[Z_0]} \mathbb{E}[S_0 S_{-k}] \right).$$

In other words, under the present assumptions, the worst-case is deterministic with batch sizes fixed to the mean value. A related result was obtained in [8].

In [16, 17], the authors identify conditions under which a model of equation-based rate control formulated in [16, 17] is resulting in $\bar{x} \leq f(p)$, where \bar{x} is the time-average send rate. If the function $p \rightarrow f(p)$ is such that in the reference system of deterministic constant inter-loss times, $\bar{x}' = f(p')$, then, it follows that for any stationary random loss process with the loss event rate p' , and under the conditions in [16], it holds $\bar{x} \leq \bar{x}'$. Note that, if the last holds, then determinism is an extremal, but in contrast to the results established for the increase-decrease controls in this note, it is a *best-case*.

5. ACKNOWLEDGMENTS

The first author would like to thank Prof. P. A. Humblet for his kindness and promptness in sending a copy of his manuscript [9]. We thank anonymous reviewers for their suggestions, in particular, adding Example 1 was motivated by a construction of one of the reviewers.

REFERENCES

- [1] Eitan Altman, Konstantin Avrachenkov, and Chadi Barakat. A stochastic model of TCP/IP with stationary random losses. In *Proc. of ACM SIGCOMM 2000*, pages 231–242, 2000.
- [2] Francois Baccelli and Pierre Brémaud. *Elements of Queueing Theory*, volume 26. Applications of Mathematics, Springer-Verlag, 1991.
- [3] A. A. Borovkov. *Ergodicity and Stability of Stochastic Processes*. John Wiley & Sons, Wiley Series in Probability and Statistics, 1998.
- [4] D. Chiu and R. Jain. Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks. *Computer Networks and ISDN Systems*, 17:1–14, June 1989.
- [5] Vincent Dumas, Fabrice Guillemin, and Philippe Robert. A Markovian Analysis of Additive-Increase Multiplicative-Decrease Algorithms. *Adv. Appl. Prob.* 34:85–111, 2002.
- [6] Sally Floyd. Highspeed TCP for large congestion windows, August 2002. Internet Engineering Task Force, INTERNET-DRAFT, draft-floyd-tcp-highspeed-01.txt.
- [7] Sally Floyd, Sylvia Ratnasamy, and Scott Shenker. Modifying TCP’s congestion control for high speeds, May 2002. preprint, <http://www.icir.org>.
- [8] Fabrice Guillemin, Philippe Robert, and Bert Zwart. AIMD Algorithms and Exponential Functionals. To Appear in *Annals of Applied Probability*, 2002.
- [9] Pierre A. Humblet. Determinism minimizes waiting time in queues. Technical Report LIDS-P-1207, MIT, May 1982.
- [10] Van Jacobson and Michael J. Karels. Congestion Avoidance and Control. In *Proc. of ACM SIGCOMM 1988*, 314–329, Stanford, CA.

- [11] J. B. Keller. Inverse problems. *Am. Math. Mon.*, 83:107–118, 1976.
- [12] Teunis J. Ott and J. H. B. Kemperman and Matt Mathis. The Stationary Behavior of Ideal TCP Congestion Avoidance. Unpublished Manuscript, <http://web.njit.edu/~ott/Papers>, Aug, 1996.
- [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.
- [14] Serap Savari and Emre Telatar. The Behavior of Certain Stochastic Processes Arising in Window Protocols. In *Proc. of IEEE Globecom'99*, 1B:791–795, Rio de Janeiro, Brazil, 1999.
- [15] Dietrich Stoyan. *Comparison Methods: for Queues and Other Stochastic Models*. Wiley Series in Probability and Mathematical Statistics, 1983.
- [16] Milan Vojnović and Jean-Yves Le Boudec. On the long-run behavior of equation-based rate control. In *Proc. of ACM SIGCOMM 2002*, pages 173–184, Pittsburgh, PA, August 2002.
- [17] Milan Vojnović and Jean-Yves Le Boudec. Some Observations on Equation-based Rate Control. In *Proc. of ITC-17*, 173–184, Salvador da Bahia, Brazil, Dec, 2001.

APPENDIX A. PROOF OF LEMMA 2

The proof is based on standard tools of stochastic orderings for stochastic recursive sequences, see for instance [2], Chapter 4. Let \leq_{icx} be binary relation of increasing convex ordering; for any two cumulative distribution functions F and G on \mathbb{R} , it is defined as

$$F \leq_{icx} G \Leftrightarrow \int_{\mathbb{R}} f(x)F(dx) \leq \int_{\mathbb{R}} f(x)G(dx), \text{ for any } f \in \mathcal{L},$$

where \mathcal{L} is the set of all increasing convex functions. An ordering for sequences holds, if it holds component-wise.

Let $\{W_0, W_1, W_2, \dots\}$ be a sequence of embedded windows that obey to the recurrence (2) with driving sequence of inter-loss times $\{S_0, S_1, S_2, \dots\}$, $E[S_n] = 1/\lambda$, all $n = 0, 1, 2, \dots$. Similarly, let $\{W'_0, W'_1, W'_2, \dots\}$ be a sequence of embedded windows that obey to (2), but with inter-loss times $\{1/\lambda, 1/\lambda, 1/\lambda, \dots\}$. Assume also $Y'_0 \leq_{icx} Y_0$ (say, $Y'_0 = Y_0$).

We use a convenient transformation $Y_n = \phi(W_n)$, $n = 0, 1, 2, \dots$. It indeed holds

$$\{Y'_0, 1/\lambda, 1/\lambda, 1/\lambda, \dots\} \leq_{icx} \{Y_0, S_0, S_1, S_2, \dots\}.$$

Under (A1), it follows from a known result of stochastic orderings for stochastic recursive sequences (see Section 2.2, Chapter 4, [2]) that if the the last ordering holds, then

$$(8) \quad \{Y'_1, Y'_2, Y'_3, \dots\} \leq_{icx} \{Y_1, Y_2, Y_3, \dots\}.$$

Note that by Palm inversion formula, we have

$$(9) \quad E[W(\infty)] = \lambda E[z(Y_\infty, S_\infty)],$$

where, by definition,

$$z(y, s) = \Phi(y + s) - \Phi(y),$$

and $x \rightarrow \Phi(x)$ is a primitive of ϕ^{-1} . $W(\infty)$ is a random variable with distribution of the steady-state window. $Y_\infty := \phi(W_\infty)$, where W_∞ is a random variable with distribution equal to the Palm distribution with respect to the point process of loss events.

Hence, it is readily seen that $E[W(\infty)] \geq E[W'(\infty)]$ is equivalent to

$$(10) \quad E[z(Y_\infty, S_\infty)] \geq E[z(Y'_\infty, 1/\lambda)] = z(Y'_\infty, 1/\lambda).$$

We first note, for any fixed $n = 0, 1, 2, \dots$,

$$\begin{aligned}
\mathbb{E}[z(Y_n, S_n)] &= \mathbb{E}[\Phi(Y_n + S_n) - \Phi(Y_n)] \\
&= \mathbb{E}[\mathbb{E}[\Phi(Y_n + S_n)|Y_n] - \Phi(Y_n)] \\
(11) \quad &\geq \mathbb{E}[\Phi(Y_n + \mathbb{E}[S_n|Y_n]) - \Phi(Y_n)] \\
&= \mathbb{E}[\Phi(Y_n + 1/\lambda) - \Phi(Y_n)] \quad (\text{independence}) \\
&= \mathbb{E}[z(Y_n, 1/\lambda)].
\end{aligned}$$

The inequality above follows from the following property,

(P1): $x \rightarrow \Phi(x)$ is convex iff $x \rightarrow \phi^{-1}(x)$ is non-decreasing,

and the fact that by hypothesis $a(\cdot)$ is positive-valued, hence, $x \rightarrow \phi(x)$ is non-decreasing, and thus $x \rightarrow \phi^{-1}(x)$ is non-decreasing as well.

By (A2), we have

$$(12) \quad z^*(y, s) \leq z(y, s) \leq (1 + \varepsilon)z^*(y, s),$$

where, by definition,

$$z^*(y, s) = \int_y^{y+s} \varphi(x) dx.$$

Note that for $h_s(y) := z^*(y, s)$, any fixed $s \geq 0$, $h_s''(y) = \varphi'(y + s) - \varphi'(y)$. Hence, for any fixed $s \geq 0$, $y \rightarrow z^*(y, s)$ is convex iff $x \rightarrow \varphi'(x)$ is non-decreasing. The last property is implied by $x \rightarrow \varphi(x)$ convex.

From the inequalities in (12), the property that for any fixed $s \geq 0$, $y \rightarrow z^*(y, s)$ is convex, and the increasing convex order (8), it follows

$$\mathbb{E}[z(Y_n, 1/\lambda)] \geq \mathbb{E}[z^*(Y_n, 1/\lambda)] \geq \mathbb{E}[z^*(Y'_n, 1/\lambda)] \geq \frac{1}{1 + \varepsilon} \mathbb{E}[z(Y'_n, 1/\lambda)].$$

Lastly, recall from (10) that $\mathbb{E}[z(Y_\infty, S_\infty)] \geq \mathbb{E}[z(Y_\infty, 1/\lambda)]$. Combining the last with the above display, we have $\mathbb{E}[z(Y_\infty, S_\infty)] \geq \mathbb{E}[z(Y'_\infty, 1/\lambda)]$, which we already noted is equivalent to $\mathbb{E}[W(\infty)] \geq \mathbb{E}[W'(\infty)]$.

APPENDIX B. PROOF OF LEMMA 3

Firstly, from the hypothesis $\bar{w}' \geq f(p')$ and the inversion formula $\bar{w}' = \lambda/p'$, we conclude $\lambda \geq p'f(p')$. Hence,

$$(13) \quad \bar{w} = \frac{\lambda}{p} \geq \frac{p'f(p')}{p}.$$

Secondly, from the inversion formulas $\bar{w} = \lambda/p$ and $\bar{w}' = \lambda/p'$, and the hypothesis $\bar{w} \geq \frac{1}{1+\varepsilon}\bar{w}'$, we conclude

$$\frac{1}{1 + \varepsilon} p \leq p'.$$

Lastly, by a hypothesis, $x \rightarrow xf(x)$ is non-decreasing, and hence, we have $p'f(p') \geq \frac{1}{1+\varepsilon}pf(\frac{1}{1+\varepsilon}p)$. From the last inequality, and (13), we conclude that $\bar{w} \geq \frac{1}{1+\varepsilon}f(\frac{1}{1+\varepsilon}p)$.

APPENDIX C. PROOF OF THEOREM 3

By a hypothesis, $1/a$ is strictly-positive, and thus $x \rightarrow \phi(x)$ is non-decreasing. Hence, $x \rightarrow \phi^{-1}(x)$ is non-decreasing as well. From the last property and the ordering (8), it follows

$$(14) \quad \mathbb{E}[W'_\infty] = \phi^{-1}(\mathbb{E}[Y'_\infty]) \leq \phi^{-1}(\mathbb{E}[Y_\infty]) = \phi^{-1}(\mathbb{E}[\phi(W_\infty)]).$$




FIGURE 6. (Top) $x \rightarrow g(x)$ for HighSpeed TCP, $g(x) := \phi(b(\phi^{-1}(x)))$, (Bottom) $x \rightarrow g'(x)$. The graphs demonstrate $x \rightarrow g(x)$ is convex.

Now recall the inequality in (A2) that reads as $\varphi(x) \leq \phi^{-1}(x) \leq (1 + \varepsilon)\varphi(x)$, all $x \geq 0$. From the former inequality we can conclude $\phi \leq \varphi^{-1}$. Hence,

$$(15) \quad \phi^{-1}(\mathbf{E}[\phi(W_n)]) \leq (1 + \varepsilon)\varphi(\varphi^{-1}(\mathbf{E}[W_n])) = \mathbf{E}[W_n].$$

Conjunction of (14) and (15) completes the proof.

APPENDIX D. ANALYTICAL AND NUMERICAL ARGUMENTS THAT CONFIRM CLAIM 1

We only have to verify that (A1), (A2) with $\varepsilon \in (0, 0.0012)$, (A3), and $\bar{w}' \geq f(p')$, are true, then the first assertion of the claim follows from Theorem 2. The second assertion follows by carrying on the step in Remark 2.

Step 1: (A1) is true, that is $x \rightarrow g(x)$ is convex. Here, by definition, $g(x) = \phi(b(\phi^{-1}(x)))$. We do not give a rigorous proof here, but rely on a direct numerical computation. For $x \leq w_*$, $g'(x) = 1/2$, else

$$g'(x) = \frac{b'(\phi^{-1}(x))a(\phi^{-1}(x))}{a(b(\phi^{-1}(x)))}.$$

$x \rightarrow g'(x)$ is non-decreasing as shown in Fig. 6. Combining this with the fact that $x \rightarrow g(x)$ is continuous, any chord on $g(\cdot)$ lies above $g(\cdot)$, hence, $x \rightarrow g(x)$ is convex.

Step 2: (A2) is true for some $\varepsilon \in (0, 0.0012)$. Let κ be such that

$$\kappa x \leq \phi^{-1}(x), \text{ all } x \geq 0.$$

We take

$$(16) \quad \kappa = \inf_{x \geq 0} \frac{\phi^{-1}(x)}{x}.$$




FIGURE 7. An exaggerated illustration to show how for HighSpeed TCP $x \rightarrow \phi^{-1}(x)$ deviate from its convex closure.

Let $x_* > 0$ be such that $\kappa x_* = \phi^{-1}(x_*)$. Define

$$\varphi(x) = \begin{cases} \kappa x, & x \leq x_*, \\ \phi^{-1}(x), & \text{else.} \end{cases}$$

The function $x \rightarrow \varphi(x)$ is the convex closure of $\phi^{-1}(x)$, see Fig. 7 for a pictorial illustration. Next, consider

$$\phi^{-1}(x) \leq \varphi(x) \sup_{y \geq 0} \frac{\phi^{-1}(y)}{\varphi(y)}.$$

Now, note, in our particular instance (see Fig. 7),

$$\sup_{y \geq 0} \frac{\phi^{-1}(y)}{\varphi(y)} = \frac{\phi^{-1}(w_*)}{\kappa w_*} = \frac{1}{\kappa}.$$

From (16), and the substitution $x = \phi(y)$,

$$\frac{1}{\kappa} = \sup_{x \geq 0} \frac{x}{\phi^{-1}(x)} = \sup_{y \geq 0} \frac{\phi(y)}{y}.$$

By direct numerical computation we obtain $1/\kappa < 1 + \varepsilon$, where $0 < \varepsilon < 0.0012$.

Step 3: (A3) is true. For HighSpeed TCP, $pf(p) = \sqrt{3/2}p^{1/2}$, for $p \leq 0.001$, else, $pf(p) = K_1 p^{1-1/\gamma_1}$. Given the values of K_1 and γ_1 , $p \rightarrow pf(p)$ is indeed non-decreasing on $[0, 0.001]$ and $[0.001, 1]$. Given also that the values of K_1 and γ_1 are set in [6] such that $p \rightarrow f(p)$ is continuous, $p \rightarrow pf(p)$ is continuous as well. Hence, $p \rightarrow pf(p)$ is non-decreasing on the entire domain $[0, 1]$.

Step 4: $\bar{w}' \geq f(p')$. See Appendix E.

Step 5: Second assertion holds. Recall Remark 2. It is readily seen

$$\inf_{x \in [0, 1]} \frac{f(\frac{1}{1+\varepsilon}x)}{f(x)} = (1 + \varepsilon)^{-1/\gamma_1}.$$

Hence, we have

$$\bar{w} \geq (1 + \varepsilon)^{-(1+1/\gamma_1)} f(p).$$

By numerical computation for $\gamma_1 = 1/1.2$, $(1 + \varepsilon)^{-(1+1/\gamma_1)} \geq 1 - \varepsilon''$, for some $\varepsilon'' \in (0, 0.0023)$, and the assertion follows.

APPENDIX E. NUMERICAL COMPUTATION THAT CONFIRMS CLAIM 2

We solve the direct problem (3). Define ψ as a primitive of ϕ . Note that for $w \in [0, w_*]$, $a(w) = 1$ and $b(w) = w/2$. Hence, $\phi(w) = w$, and $\psi(w) = w^2/2$, for $w \in [0, w_*]$.

We have

$$\phi(w) = \phi_0(w) + w_* - \phi_0(w_*), \quad w \geq w_*,$$

where $\phi_C(x)$ is the primitive of $1/a$ with the integration constant $C \in \mathbb{R}$. A bit of elementary integral calculus, reveals

$$\phi_0(w) = K_1^{-\frac{1}{\gamma_1}} \left[\frac{1}{c} e^{-\frac{d}{c} \frac{1-\gamma_1}{\gamma_1}} e\left(-\frac{1-\gamma_1}{\gamma_1} (\ln w + \frac{d}{c})\right) - \frac{\gamma_1}{2(1-\gamma_1)} w^{\frac{1}{\gamma_1}-1} \right].$$

Where, we define

$$e(x) = \int \frac{e^{-x}}{x} dx = \ln x + \sum_{n=1}^{\infty} \frac{(-1)^n}{n \cdot n!} x^n.$$

Further, we calculate $w \rightarrow \psi(w)$, as

$$\psi(w) = \psi_0(w) + (w_* - \phi_0(w_*))(w - w_*) + \frac{1}{2} w_*^2 - \psi_0(w_*), \quad w \geq w_*,$$

where, ψ_C is the primitive of ϕ with the integration constant $C \in \mathbb{R}$. A bit more of elementary integral calculus yields

$$\psi_0(w) = K_1^{-\frac{1}{\gamma_1}} \left\{ \frac{1}{c} e^{-\frac{d}{c} \frac{1-\gamma_1}{\gamma_1}} \left[e^{\frac{d}{c} w \ln(\frac{\gamma_1-1}{\gamma_1} (\ln w + \frac{d}{c}))} - e(-\ln w - \frac{d}{c}) - \frac{\gamma_1}{1-\gamma_1} s(\frac{\gamma_1-1}{\gamma_1} (\ln w + \frac{d}{c})) \right] - \frac{\gamma_1^2}{2(1-\gamma_1)} w^{\frac{1}{\gamma_1}} \right\},$$

where

$$s(x) = \sum_{n=1}^{\infty} \frac{(-1)^n}{n \cdot n!} \left(\frac{1-\gamma_1}{\gamma_1} \right)^{n+1} e_n\left(\frac{\gamma_1}{1-\gamma_1} x\right),$$

and $e_n(x) := \int x^n e^{-x} dx$. In numerical solving, we use the elementary recursion $e_n(x) = -x^n e^{-x} + n \cdot e_{n-1}(x)$, $n > 0$.

By above calculations, we have $w \rightarrow \phi(w)$ and $w \rightarrow \psi(w)$. Now solving the direct problem (3) corresponds to, for any fixed parameter λ , solve \bar{w}_0 and p that obey to

$$\bar{w}_0 = b(\phi^{-1}(\phi(\bar{w}_0) + 1/\lambda))$$

$$\frac{1}{p} = [\phi(\bar{w}_0) + 1/\lambda] \phi^{-1}(\phi(\bar{w}_0) + 1/\lambda) - \bar{w}_0 \phi(\bar{w}_0) - \psi(\phi^{-1}(\phi(\bar{w}_0) + 1/\lambda)) + \psi(\bar{w}_0).$$

Then, by observing $\bar{w} = \frac{\lambda}{p}$, we obtain $p \rightarrow \bar{w}$, for any fixed parameter λ . We show a numerical result in Fig. 5, which confirms Claim 2.