

An extension of the square root law of TCP

Krishanu Maulik* and Bert Zwart†

* Stat-Math Unit
Indian Statistical Institute
203 Barrackpore Trunk Road Kolkata 700 108 INDIA
krishanu@isical.ac.in

† Georgia Institute of Technology
H. Milton Stewart School of Industrial and Systems Engineering
Ferst Drive 765, 30332, Atlanta, USA
bertzward@gatech.edu

Abstract

Using probabilistic scaling methods, we extend the square root law of TCP to schemes which may not be of the AIMD type. Our results offer insight in the relationship between throughput and loss rate, and the time scale on which losses take place. Similar results are shown to hold in scenarios where dependencies between losses occur.

1 Introduction

Since the introduction and emergence of TCP [12], a large number of studies have been devoted to the modeling and performance analysis of TCP. The focus of most modeling studies is on an essential feature of current implementations of TCP: the window size is regulated by an *additive increase multiplicative decrease* (AIMD) algorithm. An emerging theme in this line of research is a *square root law*, (see, for example, [16]) relating the throughput ω of a connection with the rate α at which packets get lost: $\omega \approx K/\sqrt{\alpha}$. The constant K depends on the version of TCP, but also on the models. Several different approaches have been used in this line of work; for example direct modeling and performance analysis [2, 3], simulation [7], approximations [21] and probabilistic scaling [5, 11, 19, 20].

In this work, we combine probabilistic scaling and performance analysis methods to analyze a class of congestion avoidance algorithms which we (following [3]) call *general increase multiplicative decrease* (GIMD) algorithms. As argued in [2, 3, 14], the recent explosion of bandwidth on the Internet raises the question whether the conservative AIMD scheme is (still) optimal. In particular, the *multiplicative increase multiplicative decrease* scheme, also called *scalable TCP* has been proposed; cf. Tom Kelly [14] (see also Floyd [8]). The results in [14] indicate that $\omega \approx K/\alpha$ in this case, which connects MIMD to the proportional fairness allocation mechanism in flow level models introduced by Frank Kelly *et al.* [13]; see for example Massoulié [15] and references therein.

The findings in [14] are confirmed in this paper, although we observe that scalable TCP is not the only scheme leading to a throughput of the order $1/\alpha$. In fact, the main topic of this work is a general investigation of the relation between the throughput and the loss rate in situations where the loss rate is small. As can be expected (see also [19]), the relation $\omega \approx K\alpha^{-b}$ prevails in general, but b is in general not equal to $1/2$. As our results indicate, the exponent b is strongly dependent on the rate at which the window size increases if no loss occurs. In addition, we observe that the rate at which a loss occurs as function of the window size has a major impact as well. This rate is often assumed to be linear in the window size, but this may not need to be the case, see [2] for a situation in which the number of losses is independent of the window size.

In addition, we allow losses to occur in clumps, or batches. Clumping of losses is observed by Paxson [22], and modelled in an AIMD setting in Guillemin *et al.* [11]. In that paper it is shown that several characteristics, such as slow start and variable RTT, do not play a role on the critical time scale if losses are rare. A similar observation can be made for the time between two losses within a single clump.

Motivated by these findings, we consider the following model of the evolution of the window size of a TCP connection. We consider a single lossy data stream controlled by TCP, which behavior does not influence the state of the entire network. Let $\kappa \in [0, 1]$, $h \geq 0$ and consider, for fixed $\alpha > 0$, the Markov chain $\{W_n, n \geq 1\}$, given by

$$W_{n+1} = \begin{cases} W_n + \lceil c(\alpha)W_n^\kappa \rceil, & \text{with probability } e^{-\alpha W_n^h}, \\ \lceil \delta^{B_n} W_n \rceil, & \text{with probability } 1 - e^{-\alpha W_n^h}, \end{cases} \quad (1)$$

where $\{B_n\}$ is an i.i.d. sequence independent of W_0 . B_n represents the number of losses in the n th clump. We let B be an independent copy of B_1 and assume that B is strictly positive with probability 1. For practical purposes B will be positive integer valued, but we do not put any such restriction, since it is not important for our analysis. As before, α is the loss rate.

If $h = 1$, the probability that a loss occurs is approximately αW_n , implying that (clumps of) losses occur independently of one another. If it is reasonable to assume that the probability of a loss is not dependent on the window size, then one can take $h = 0$ (again, we refer to [2] for such situations). $c(\alpha)$ may be a constant or may vary with α ; we give explicit assumptions in Section 2.

When $\kappa = 0$, the model (1) reduces to the usual AIMD scheme, while scalable TCP is obtained by taking $\kappa = 1$. Also, note that in the AIMD case, the second term in the increase part is constant $c(\alpha)$. So it makes sense to choose $c(\alpha)$ as an integer itself and then round-up is not required. In this case, for fixed $\alpha > 0$, the Markov chain simplifies to

$$W_{n+1} = \begin{cases} W_n + c(\alpha), & \text{with probability } e^{-\alpha W_n^h}, \\ \lceil \delta^{B_n} W_n \rceil, & \text{with probability } 1 - e^{-\alpha W_n^h}. \end{cases} \quad (2)$$

This paper consists of three main parts. The first part (see Section 2) is devoted to analyzing the macroscopic behavior of the Markov chain $\{W_n\}$ as $\alpha \downarrow 0$. We identify the appropriate space scale and time scale that prevail when the loss rate $\alpha \downarrow 0$. This leads to a limit theorem which states that the dynamics of an appropriate scaled version of $\{W_n\}$ behave as the dynamics of a much more tractable Markov process $\{V_t, t \geq 0\}$. The way space is scaled leads us to the insights that are necessary to obtain a formula for the exponent b mentioned above. In fact, we shall show in Theorem 1 that the throughput ω is inversely proportional to $f(\alpha) = (\alpha/c(\alpha))^{1/(1+h-\kappa)}$. Thus, if we can choose $c(\alpha)$ to be a constant, which is a possibility if $\kappa < 1$, we shall have $b = 1/(1+h-\kappa)$. Thus, our proposed extension of the square root law is to approximate the throughput $E[W_\infty]$ with $E[V]/f(\alpha)$, with V a random variable which distribution is invariant for the process $\{V_t, t \geq 0\}$.

It would be convenient to have more information about the value of the constant $E[V]$. For the AIMD model considered in [5], it was shown that this constant equals 1.3098. Remarkably, this is in agreement with existing simulation results; see [5] for a discussion. This provides additional motivation to analyze the invariant distribution of $\{V_t, t \geq 0\}$, which is also interesting in its own right and forms the topic of Section 3. We derive an infinite product representation for the Mellin transform of the invariant distribution, which can be used to obtain the first moment of that distribution.

As argued and demonstrated in [3], the process $\{V_t : t \geq 0\}$ can also be used directly as a continuous time model for TCP like mechanisms. From that point of view, Section 2 can be seen as an additional motivation for the models considered in [3]. The conditions in that paper allow for more general decrease mechanisms than multiplicative decrease, but do not include clumps of losses. In addition, the performance formulae given in that paper focus on the distribution rather than moments. Therefore, the results in Section 3 of the present paper seem to complement the ones mentioned in [3].

Unfortunately, all these results do not guarantee that the invariant distribution of (a scaled version of) the Markov process $\{W_n\}$ converges to the invariant distribution of $\{V_t\}$. This is an infamous and often neglected problem in applied probability, and usually must be verified on a case by case basis. In Section 4, we exploit ideas recently developed by Ott and Swanson [20] to deal with this issue. As in [20], we first prove a compactness result by invoking a general result from Ethier and Kurtz [6]. The interchange of limits (n and/or $t \rightarrow \infty$ and $\alpha \downarrow 0$) can then be justified by using the results from Sections 2 and 3.

The models and techniques in this paper are closely related to the papers Dumas *et al.* [5], Guillemin *et al.* [11] and Ott and Swanson [20]. The first two papers deal with the special

AIMD case of our model, where in addition $h = 1$ is assumed. Our techniques in Section 2 are related to [5]. The approach in Section 3 differs from these papers, as we directly analyze the continuous time model, rather than looking at the embedded chain first. The analysis of Section 4 is more complicated than the corresponding result in [5], since it must be shown that the family of invariant distributions (indexed by α) is in some sense bounded away from 0, uniformly in α . This is only simple to verify in the AIMD case, see Section 4 for more details.

Similar problems to the one in Section 4 are encountered in [20]. The class of models considered there constitutes an extension of the basic AIMD scheme as well. The differences between our study and that work are that our model operates on the lattice (to reflect the discrete nature of packets), and that our model operates on a different time scale. The time scale considered in [20] coincides with the number of transmitted packets, which is completely different from the conventional time scale considered here and in [5, 11]. This difference allows us to treat more general cases involving dependencies in packet losses, as reflected by the possibility that h can be different from 1, and the fact that losses can occur in clumps. An experimental paper considering a model that is similar to [20] is [4].

The rest of this paper is organized as follows. More details of the model and the convergence result for a scaled version of the Markov chain $\{W_n\}$ to a limiting Markov process $\{V_t\}$ are given in Section 2. Section 3 is devoted to an analysis of the invariant distribution of this limiting Markov process. The convergence of the corresponding invariant distributions is shown in Section 4. Concluding remarks can be found in Section 5.

2 Convergence to a limiting Markov process

In this section, we provide the asymptotic behavior of the process $\{W_n\}$ scaled both in space and time. We exclude the case where both $h = 0$ and $\kappa = 1$, corresponding to scalable TCP with loss probability independent of the window size. This setup is considered later separately, see Remark 2.5.

We shall scale the space by

$$f(\alpha) = \left(\frac{\alpha}{c(\alpha)} \right)^{\frac{1}{1+h-\kappa}} \quad (3)$$

and the time by

$$\sigma(\alpha) = c(\alpha)^{\frac{h}{1+h-\kappa}} \alpha^{\frac{1-\kappa}{1+h-\kappa}}. \quad (4)$$

Note that $f(\alpha)^{1-\kappa} c(\alpha) = \sigma(\alpha) = \alpha / f(\alpha)^h$. We consider the scaled process

$$W_t^\alpha = f(\alpha) W_{\lfloor t/\sigma(\alpha) \rfloor}, t \geq 0.$$

The convergence result for this scaled process is summarized in the following theorem.

Theorem 1. *Assume that $\sigma(\alpha) \rightarrow 0$ as $\alpha \rightarrow 0$. When $\kappa = 0$ we assume that $c(\alpha)$ is constant and integer-valued. When $\kappa > 0$, we assume that $f(\alpha) = o(\sigma(\alpha))$. If $\lim_{\alpha \rightarrow 0} W_0^\alpha = v$, then the process W^α converges in distribution on the space $D[0, \infty)$ to the process $\{V_t, t \geq 0\}$ with $V(0) = v$ and generator*

$$\Omega u(x) = x^\kappa u'(x) + x^h (E[u(\delta^B x)] - u(x))$$

for any C^1 function u .

This theorem shows that, when the loss rate is small, losses occur at time scale $1/\sigma(\alpha)$, resulting in window sizes of the order $1/f(\alpha)$. To get a better feeling for the magnitude of these quantities, we provide several examples in Section 2.1. The proof of Theorem 1 is provided in Section 2.2.

Remark 2.1. As will be evident from the proof of Theorem 1, the condition that $f(\alpha) = o(\sigma(\alpha))$ is required only to handle the round-up corrections between losses. Since these are not present when $\kappa = 0$, we do not assume the condition there.

Remark 2.2. Observe that $f(\alpha)/\sigma(\alpha) = \alpha^{\kappa/(1-\kappa+h)}c(\alpha)^{-(1+h)/(1-\kappa+h)}$. Hence the extra condition required for truncation is equivalent to $\alpha^{\kappa}/c(\alpha)^{1+h} \rightarrow 0$, which is satisfied, for example, if $c(\alpha)$ is a constant, since $\kappa > 0$. Essentially this condition requires that $c(\alpha)$ does not go to 0 fast enough.

Remark 2.3. While Theorem 1 is stated for deterministic initial values so that $W_0^\alpha \rightarrow v$, it can be easily extended to the case where W_0^α converges weakly to a random variable.

2.1 Special cases

Let us consider some special choices of h and κ . We shall also take $c(\alpha) = c \in \mathbb{N}$ when $\kappa < 1$. Finally, we take $B_n \equiv 1$.

We first consider two examples involving AIMD, illustrating the impact of h .

Corollary 2 (AIMD case with loss rate independent of window size: $h = \kappa = 0$). *Let us consider the following model for the evolution of the window size*

$$W_{n+1} = \begin{cases} W_n + c, & \text{with probability } e^{-\alpha}, \\ \lceil \delta W_n \rceil, & \text{with probability } 1 - e^{-\alpha}. \end{cases}$$

Consider the scaled process $W_t^\alpha = \frac{\alpha}{c} W_{\lfloor t/\alpha \rfloor}$. If $\lim_{\alpha \rightarrow 0} W_0^\alpha = v$, then the process W^α converges in distribution on the space $D[0, \infty)$ to the process V with $V(0) = v$ and generator

$$\Omega u(x) = u'(x) + (E[u(\delta x)] - u(x))$$

for any C^1 function u .

In this case, the throughput is inversely proportional to the loss rate. A related model falling in this qualitative picture is considered by Altman *et al.* [1].

Our second example coincides with the case considered in Dumas *et al.* [5].

Corollary 3 (AIMD case with loss rate linear in window size: $h = 1, \kappa = 0$). *Let us consider the following model for the evolution of the window size*

$$W_{n+1} = \begin{cases} W_n + c, & \text{with probability } e^{-\alpha W_n}, \\ \lceil \delta W_n \rceil, & \text{with probability } 1 - e^{-\alpha W_n}. \end{cases}$$

Consider the scaled process $W_t^\alpha = \sqrt{\frac{\alpha}{c}} W_{\lfloor t/\sqrt{c\alpha} \rfloor}$. If $\lim_{\alpha \rightarrow 0} W_0^\alpha = v$, then the process W^α converges in distribution on the space $D[0, \infty)$ to the process $\{V_t, t \geq 0\}$ with $V(0) = v$ and generator

$$\Omega u(x) = u'(x) + x(E[u(\delta x)] - u(x))$$

for any C^1 function u .

In this case, the throughput is inversely proportional to the square root of the loss rate. As expected, h has a major impact on the order of magnitude of the throughput.

Next, we consider a case involving MIMD. Consider $c(\alpha) \rightarrow 0$, but not too fast so that $\alpha/c(\alpha)^{1+h} \rightarrow 0$. Let us consider the following model for the evolution of the window size:

$$W_{n+1} = \begin{cases} [(1 + c(\alpha))W_n], & \text{with probability } e^{-\alpha W_n}, \\ [\delta W_n], & \text{with probability } 1 - e^{-\alpha W_n}. \end{cases}$$

Corollary 4 (MIMD case with independent packet drop: $h = 1, \kappa = 1$). *Consider the scaled process $W_t^\alpha = \frac{\alpha}{c(\alpha)} W_{\lfloor t/c(\alpha) \rfloor}$. If $\lim_{\alpha \rightarrow 0} W_0^\alpha = v$, then the process W^α converges in distribution on the space $D[0, \infty)$ to the process $\{V_t, t \geq 0\}$ with $V(0) = v$ and generator*

$$\Omega u(x) = xu'(x) + x(E[u(\delta x)] - u(x))$$

for any C^1 function u .

We see that the throughput is not inversely proportional to α as suggested in [14], but is inversely proportional to $\alpha/c(\alpha)$. The reason is that we let the multiplicative increase factor $1 + c(\alpha) \downarrow 1$, which is necessary for our scaling to work (losses occur at a time scale of $1/c(\alpha)$). What happens if $c(\alpha)$ is kept constant is discussed next.

Remark 2.4. When $\kappa < 1$, a constant value of $c(\alpha)$ satisfies the condition of Theorem 1. However, we need to assume that $c(\alpha) \rightarrow 0$ if $\kappa = 1$. If we keep $c(\alpha)$ fixed, it can be shown that $\{\alpha W_n\}$ converges to the limiting process $\{\bar{W}_n\}$ governed by the evolution

$$\bar{W}_{n+1} = \begin{cases} (1 + c)\bar{W}_n, & \text{with probability } e^{-\bar{W}_n}, \\ \delta \bar{W}_n, & \text{with probability } 1 - e^{-\bar{W}_n}. \end{cases}$$

This process is self-similar in the sense that the evolution of $\{\bar{W}_n\}$ is statistically identical to $\{a\bar{W}_n\}$ for any $a > 0$. This result shows that the throughput is indeed inversely proportional to the loss rate as noted in [14]. It also shows that losses do not become rare as the loss rate goes to 0. In addition, the limiting process is not much more tractable than the original process. A peculiar phenomenon is that the process $\{\bar{W}_n\}$ can have uncountably many different invariant distributions if $\delta(1 + c) = 1$, while the original process $\{W_n\}$ has a unique invariant distribution, cf. the discussion in Section 4. This is an unusual phenomenon in stochastic network theory.

Remark 2.5. If $\kappa = 1$ and $h = 0$ then the process $\{W_n\}$ may be transient for some positive values of α . This can be seen as follows. Ignoring the round-offs except for the one at zero (so that the process lives on $[1, \infty)$), we define $L_n = \log W_n$. We see that $L_{n+1} = L_n + \log(1 + a)$ with probability $e^{-\alpha}$ and $L_{n+1} = (L_n - \log \delta)^+$ with probability $1 - e^{-\alpha}$, which is a simple version of Lindley's recursion; this transformation has been observed before in [2]. Thus, the phase transition does not occur at $\alpha = 0$ but at $\alpha = \alpha^*$, at which the drift of Lindley's recursion switches from negative to positive. One may apply standard heavy traffic limit theorems to obtain a scaling limit of the process L_n around $\alpha \approx \alpha^*$.

2.2 Proof of Theorem 1

Proof of Theorem 1. First we consider the case $\kappa > 0$, where round-ups between losses are made. Let P^α be the transition function for $f(\alpha)W_n$ and

$$A^\alpha = \frac{1}{\sigma(\alpha)}[P^\alpha - I].$$

Then, we have

$$\begin{aligned} A^\alpha u(x) &= \frac{u\left(x + \sigma(\alpha)\frac{f(\alpha)}{\sigma(\alpha)}\left[\frac{\sigma(\alpha)}{f(\alpha)}x^\kappa\right]\right) - u(x)}{\sigma(\alpha)}e^{-\sigma(\alpha)x^h} \\ &\quad + \left(E\left[u\left(f(\alpha)\left[\frac{\delta^B x}{f(\alpha)}\right]\right)\right] - u(x)\right)\frac{1 - e^{-\sigma(\alpha)x^h}}{\sigma(\alpha)}. \end{aligned}$$

We can bound the difference $|A^\alpha u(x) - \Omega u(x)|$ by $\Delta_1(x) + \Delta_2(x)$, where

$$\Delta_1(x) = \left| \frac{u\left(x + \sigma(\alpha)\frac{f(\alpha)}{\sigma(\alpha)}\left[\frac{\sigma(\alpha)}{f(\alpha)}x^\kappa\right]\right) - u(x)}{\sigma(\alpha)} - x^\kappa u'(x) \right| + x^\kappa |u'(x)| |1 - e^{-\sigma(\alpha)x^h}|$$

and

$$\begin{aligned} \Delta_2(x) &= \left(E\left[u\left(f(\alpha)\left[\frac{\delta^B x}{f(\alpha)}\right]\right)\right] + |u(x)| \right) \left| \frac{1 - e^{-\sigma(\alpha)x^h}}{\sigma(\alpha)} - x^h \right| \\ &\quad + x^h E\left[\left|u\left(f(\alpha)\left[\frac{\delta^B x}{f(\alpha)}\right]\right) - u(\delta^B x)\right|\right]. \end{aligned}$$

We check that $\Delta_1(x)$ and $\Delta_2(x)$ converge to 0 uniformly in $u \in \mathcal{C}_c^2$, the class of twice continuously differentiable functions with compact support. First observe that

$$0 \leq \frac{f(\alpha)}{\sigma(\alpha)}\left[\frac{\sigma(\alpha)}{f(\alpha)}x^\kappa\right] - x^\kappa < \frac{f(\alpha)}{\sigma(\alpha)}, \quad (5)$$

since $f(\alpha) = o(\sigma(\alpha))$, and hence converges uniformly in $x \in [0, \infty)$. We use a Taylor expansion to estimate $\Delta_1(x)$ and $\Delta_2(x)$. Since, for small enough α ,

$$\sigma(\alpha)\frac{f(\alpha)}{\sigma(\alpha)}\left[\frac{\sigma(\alpha)}{f(\alpha)}x^\kappa\right] \leq \sigma(\alpha)x^\kappa + f(\alpha) < x + 1,$$

we have, for small enough α ,

$$\begin{aligned} \Delta_1(x) &\leq |u'(x)| \left| \frac{f(\alpha)}{\sigma(\alpha)}\left[\frac{\sigma(\alpha)}{f(\alpha)}x^\kappa\right] - x^\kappa \right| + \frac{\sigma(\alpha)}{2} \left[\frac{f(\alpha)}{\sigma(\alpha)}\left[\frac{\sigma(\alpha)}{f(\alpha)}x^\kappa\right] \right]^2 \sup_{\xi \in [x, 2x+1]} |u''(\xi)| \\ &\quad + \sigma(\alpha)x^{1+h-\kappa}|u'(x)| \\ &\leq \left| \frac{f(\alpha)}{\sigma(\alpha)} + \sigma(\alpha)x^{1+\kappa-h} \right| |u'(x)| + \sigma(\alpha)(x^{2\kappa} + 1) \sup_{\xi \in [x, 2x+1]} |u''(x)|, \end{aligned}$$

which goes to 0 uniformly in $u \in \mathcal{C}_c^2$. Also, since $f(\alpha) \left\lceil \frac{\delta^B x}{f(\alpha)} \right\rceil \leq x + 1$, we have,

$$\Delta_2(x) \leq \sigma(\alpha)x^{2h} \sup_{\xi \leq x} |u(\xi)| + f(\alpha)x^h \sup_{\xi \leq x+1} |u'(\xi)|,$$

which goes to 0 as well uniformly in $u \in \mathcal{C}_c^2$. Thus, $A^\alpha u$ converges to Ωu uniformly in $u \in \mathcal{C}_c^2$. Observe finally that \mathcal{C}_c^2 is strongly separating, that is, given any $\delta > 0$ and $x \geq 0$, there exists an $f \in \mathcal{C}_c^2$, such that

$$\inf\{|f(x) - f(y)| : |x - y| > \delta\} > 0.$$

This enables us to apply Corollary 4.8.9 of [6] with $G_n = \mathbb{R}_+$ in condition (h), which yields the required result.

The case $\kappa = 0$, modeled by (2) is similar and in fact, easier, since there is no adjustment due to round-up. So we do not need a bound like in (5). We skip the details to avoid repetition. \square

3 Invariant distribution of the limiting process

The limiting Markov process $\{V_t, t \geq 0\}$, obtained in the last section, has generator

$$(\Omega u)(x) = x^\kappa u'(x) + x^h (E[u(x\delta^B)] - u(x)), \quad (6)$$

with $h \geq 0$, $\kappa \in [0, 1]$, and $\kappa < 1$ if $h = 0$. The goal of this section is to investigate the invariant distribution of this Markov process.

We use geometric Lyapounov functions to derive a bound to be satisfied by any invariant distribution, which implies that all moments of such distributions are finite. The property of having all moments finite allows us to characterize and derive an expression for the Mellin transform of the invariant distribution, establishing existence and uniqueness.

The convergence to the invariant distribution is non-trivial, since periodicity can occur (if B is lattice and $\kappa = 1$), and since downward jumps of arbitrary size can cause the process to reach zero arbitrary closely. Fortunately, from an applied point of view, it is not necessary to establish this convergence. The results derived in this section are sufficient to show convergence of the limiting distributions of W^α to the invariant distribution associated with (6); this will be demonstrated in Section 4.

3.1 Qualitative properties

It is useful to realize what the sample paths of V_t look like. The downward jumps occur at rate $\lambda(V_t) = V_t^h$, and in between downward jumps, the sample paths of V_t increase deterministically according to the differential equation $\frac{d}{dt} V_t = V_t^\kappa$. It is useful to note that the evolution of V_t till the first drop after time 0 is given by

$$V_{t,\kappa}^*(V_0) \equiv V_t^* = \begin{cases} V_0 e^t, & \text{if } \kappa = 1, \\ (V_0^{1-\kappa} + (1-\kappa)t)^{\frac{1}{1-\kappa}}, & \text{if } 0 \leq \kappa < 1. \end{cases} \quad (7)$$

Further observe that for a given κ and starting point V_0 , V_t is the largest when there is no downward jump in between, that is, the largest value is V_t^* , which is increasing in V_0 .

Our next concern is to derive a bound which must be satisfied by all invariant distributions. In the discussion that follows, let V be a random variable with such a distribution, that is, which satisfies $E[(\Omega u)(V)] = 0$ for all continuously differentiable functions u .

Proposition 5. *For any $s > 0$, the function $u(x) = x^s$ is a geometric Lyapounov function for V_t , that is, there exist constants $\epsilon > 0$, $\bar{x} < \infty$ and $0 < \eta < 1$ (all dependent on s) such that*

$$E[V_\epsilon^s \mid V_0 = x] \leq x^s \eta, x \geq \bar{x}. \quad (8)$$

Consequently, for any choice of V_0 which makes $\{V_t\}$ stationary, $E[V_0^s] < \infty$.

Proof. We first show that the first result implies the second by using similar techniques as in [9] and [18]. Let V_0 be such that V_t is stationary. As observed earlier, for a given $\kappa \in [0, 1]$, on the set $[V_0 < \bar{x}]$, the largest value of V_t is $V_{t,\kappa}^*(\bar{x})$, given in (9). Thus we obtain

$$\begin{aligned} E[V_0^s] &= E[V_\epsilon^s] \\ &= E[V_\epsilon^s I(V_0 < \bar{x})] + E[V_\epsilon^s I(V_0 \geq \bar{x})] \\ &\leq V_{\epsilon,\kappa}^*(\bar{x})^s + \int_{\bar{x}}^{\infty} E[V_\epsilon^s \mid V_0 = x] dP(V_0 \leq x) \\ &\leq V_{\epsilon,\kappa}^*(\bar{x})^s + \int_{\bar{x}}^{\infty} \eta x^s dP(V_0 \leq x) \\ &\leq V_{\epsilon,\kappa}^*(\bar{x})^s + \eta E[V_0^s], \end{aligned}$$

which implies

$$E[V_0^s] \leq \frac{V_{\epsilon,\kappa}^*(\bar{x})^s}{1 - \eta} < \infty.$$

To prove the first assertion, fix t and let J_t be the number of downward jumps between 0 and t . We first consider the case where $\kappa = 1$, and hence, $h > 0$. Observe that

$$E[V_t^s \mid V_0 = x] \leq (xe^t)^s P(J_t = 0 \mid V_0 = x) + E[V_t^s I(J_t \geq 1) \mid V_0 = x].$$

In presence of at least one downward jump, V_t is the largest when there is only one batch of drops, and hence, $V_t \leq \delta^B x e^t$. Furthermore, $P(J_t = 0 \mid V_0 = x) \leq e^{-tx^h}$. Consequently,

$$E[V_t^s \mid V_0 = x] \leq x^s \left(e^{-t(x^h - s)} + E[\delta^{Bs}] e^{st} \right).$$

Since $E[\delta^{Bs}] < 1$ for any $s > 0$ and $h > 0$, (8) follows by first choosing $t = \epsilon$ small enough and then \bar{x} large enough.

Next consider the case $\kappa < 1$. Now, in absence of downward jumps, we have

$$V_t = (V_0^{1-\kappa} + (1-\kappa)t)^{\frac{1}{1-\kappa}} = V_0 \left(1 + (1-\kappa) \frac{t}{V_0^{1-\kappa}} \right)^{\frac{1}{1-\kappa}} \leq V_0 e^{t/V_0^{1-\kappa}}. \quad (9)$$

Also, as in the case $\kappa = 1$, in presence of downward jumps, V_t is the largest when there is only one batch of drops. Furthermore, this is maximized when the downward jump happens at time 0 and the bound is obtained by replacing V_0 with δV_0 in (7). Consequently, arguing

as in the case $\kappa = 1$, we have, for any $M > 0$,

$$\begin{aligned}
& E[V_t^s \mid V_0 = x] \\
& \leq x^s e^{\frac{st}{1-\kappa}} e^{-t} + E \left[\{(\delta^B x)^{1-\kappa} + (1-\kappa)t\}^{\frac{s}{1-\kappa}} \right] \\
& \leq x^s e^{\frac{st}{1-\kappa}} e^{-t} + x^s E \left[\delta^{sB} e^{\frac{st}{(\delta^B x)^{1-\kappa}}}; B \leq M \right] + E \left[\{(\delta^B x)^{1-\kappa} + (1-\kappa)t\}^{\frac{s}{1-\kappa}}; B > M \right] \\
& \leq x^s e^{\frac{st}{1-\kappa}} e^{-t} + x^s e^{\frac{st}{(\delta^M x)^{1-\kappa}}} E[\delta^{sB}] + \{(\delta^B x)^{1-\kappa} + (1-\kappa)t\}^{\frac{s}{1-\kappa}} P(B > M) \\
& \leq x^s \left[e^{\frac{st}{1-\kappa}} e^{-t} + \{E[\delta^{sB}] + \delta^{sM} P(B > M)\} e^{\frac{st}{(\delta^M x)^{1-\kappa}}} \right].
\end{aligned}$$

Consequently, (8) follows by first choosing M large enough, then $t/x^{1-\kappa}$ small enough and then again choosing t large enough. \square

From Proposition 5, we have that all positive moments of invariant distributions are finite, so that the Mellin transform $v(s) = E[V^{s-1}]$ is well-defined for all real $s \geq 1$. We did not formally prove that an invariant distribution exists; we only showed that any invariant distribution must have all moments finite. We use this property in the next section to derive an expression for the Mellin transform of an invariant distribution, which also establishes existence and uniqueness.

3.2 Quantitative properties

Let V be a random variable with an invariant distribution. As was noted earlier, we necessarily have $E[(\Omega u)(V)] = 0$. Of special interest is the expected value of V , since $E[V]/f(\alpha)$ gives an approximation of the expected value of the window size. It therefore makes sense to investigate the Mellin transform $v(s) = E[V^{s-1}]$ of V . This is well defined for all values of $s > 1$.

Therefore, we take the class of test functions $u(x) = x^s$, $s \geq 0$. The equation $E[(\Omega u)(V)] = 0$ then simplifies to

$$sv(s + \kappa) = v(s + h + 1)(1 - E[\delta^{Bs}]). \quad (10)$$

Clearly $v \equiv 0$ is a solution, which corresponds to the identically 0 random variable. Also, there is a unique solution of v satisfying (10), and $v(1) = 1$ (corresponding to $V > 0$ a.s.) as is stated in the following theorem.

Theorem 6. *Suppose $\kappa < 1$. The unique solution of (10) satisfied by the Mellin transform of the invariant distribution of the limiting Markov process, with $v(1) = 1$, is given by, for $y > \kappa$,*

$$v(y) = (le^{-\gamma})^{\frac{y-1}{l}} \frac{1-\phi(y-\kappa)}{y-\kappa} \prod_{k=1}^{\infty} e^{\frac{y-1}{kl}} \frac{1-\phi(y-\kappa+kl)}{y-\kappa+kl}, \quad (11)$$

where $l = h + 1 - \kappa$, $\gamma = 0.5771\dots$ is Euler's constant and $\phi(s) = E[\delta^{sB}]$.

When $\kappa = 1$, we need to further assume that $E[B] < \infty$. Then the unique solution of (10) with $v(1) = 1$ is given by, for $y > 1$,

$$v(y) = (le^{-\gamma})^{\frac{y-1}{l}} \frac{1-\phi(y-1)}{y-1} \prod_{k=1}^{\infty} e^{\frac{y-1}{kl}} \frac{1-\phi(y-1+kl)}{y-1+kl}. \quad (12)$$

The steady-state average of the limiting scaled window size is given by $v(2)$, which is an infinite product that may be calculated numerically.

Remark 3.1. Before we prove Theorem 6, it is useful to note that the expression on the right side of (11), even though it looks complicated, is the transform of an infinite product of independent random variables. The first factor is the Mellin transform of the constant $(le^{-\gamma})^{1/l}$. To study the other factors, note that a simple application of Fubini's theorem gives

$$\frac{1 - \phi(s)}{s} = E \left[\frac{1 - \delta^{sB}}{s} \right] = \int_0^\infty e^{-su} P(-B \log \delta > u) du$$

and thus we have, with k a non-negative integer,

$$\frac{\frac{1 - \phi(y - \kappa + kl)}{y - \kappa + kl}}{\frac{1 - \phi(1 - \kappa + kl)}{1 - \kappa + kl}} = \frac{\int_0^\infty e^{-(y-1)u} e^{-(1-\kappa+kl)u} P(-B \log \delta > u) du}{\int_0^\infty e^{-(1-\kappa+kl)u} P(-B \log \delta > u) du} = E[e^{-(y-1)X_k}],$$

where X_k has density proportional to $\exp(-(1 - \kappa + kl)u)P(-B \log \delta > u)$. Note that this distribution appears after exponential tilting of the integrated tail distribution of $-B \log \delta$. Also, the representation works when $k = 0$ and $\kappa = 1$. Now take an independent sequence $\{Y_n\}$, where Y_k is distributed as $X_k - 1/(kl)$, whenever $k > 0$ and Y_0 is distributed as X_0 . Then the right side of (11) corresponds to the random variable

$$\sqrt[l]{le^{-\gamma}} \exp \left(- \sum_{k=0}^{\infty} Y_k \right).$$

Since this random variable has a Mellin transform given as above, which is continuous at 1 and has value 1 at 1, the random variable cannot have any mass at 0 and hence $\sum_{k=0}^{\infty} Y_k$ is finite with probability 1.

Proof of Theorem 6. The idea of the proof is similar to the analysis of the embedded Markov chain in [11]. We transform the functional equation for v in terms of the functional equation for the Gamma function and then apply a characterization theorem of the Gamma function known as the Bohr-Mollerup theorem. The proof is slightly more complicated here, as it involves a more tedious transformation of v , which is given by

$$h(s) = l^{1-s} \frac{v(ls + \kappa)}{v(l + \kappa)}. \tag{13}$$

It is easy to check that for any non-zero solution v of (10), we have

$$h(s + 1) = \frac{s}{1 - \phi(ls)} h(s). \tag{14}$$

Further note that $h(1) = 1$ and $\log h(s) = (1 - s) \log l + \log v(ls + \kappa) - \log v(l + \kappa)$ is convex. Now consider the infinite product $P(s)$ defined by

$$P(s) = \prod_{k=0}^{\infty} \frac{1 - \phi(l + lk)}{1 - \phi(ls + lk)}.$$

Using arguments similar to those in [17], Section 2 (namely, exploiting the fact that $\phi(ls + lN)$ converges to 0 monotonically as $N \rightarrow \infty$), one can show that this infinite product is well-defined, and that $P(s)$ is log-convex. In addition, note that $P(s) = P(s+1)/(1 - \phi(s))$, since $1 - \phi(ls + lN) \rightarrow 1$ as $N \rightarrow \infty$.

If we multiply both sides of (14) with $P(s+1)$, we see that

$$P(s+1)h(s+1) = sP(s)h(s). \quad (15)$$

Since both h and P are log-convex, so is their product. We also have $P(1) = h(1) = 1$. Applying the Bohr-Mollerup theorem to the functional equation (15), we have $P(s)h(s) = \Gamma(s)$. Consequently,

$$h(s) = \frac{\Gamma(s)}{P(s)} = \frac{\Gamma(s+1)}{sP(s)} = \Gamma(s+1) \frac{1 - \phi(ls)}{s} \prod_{k=1}^{\infty} \frac{1 - \phi(ls + lk)}{1 - \phi(lk)}.$$

Clearly, h is defined for all $s > 0$. Using Weierstrass' formula for the Gamma function, we further have

$$h(s) = e^{-\gamma s} \frac{1 - \phi(ls)}{s} \prod_{k=1}^{\infty} e^{s/k} \frac{\frac{1 - \phi(ls + lk)}{s + k}}{\frac{1 - \phi(lk)}{k}}.$$

Then, from (13), we obtain back, for $s > 0$,

$$v(ls + \kappa) = v(l + \kappa)(le^{-\gamma})^s \frac{1 - \phi(ls)}{ls} \prod_{k=1}^{\infty} e^{s/k} \frac{\frac{1 - \phi(ls + lk)}{ls + lk}}{\frac{1 - \phi(lk)}{lk}}.$$

Substituting $y = ls + \kappa$, we obtain for $y > \kappa$,

$$v(y) = v(l + \kappa)(le^{-\gamma})^{(y-\kappa)/l} \frac{1 - \phi(y - \kappa)}{y - \kappa} \prod_{k=1}^{\infty} e^{(y-\kappa)/(kl)} \frac{\frac{1 - \phi(y - \kappa + lk)}{y - \kappa + lk}}{\frac{1 - \phi(lk)}{kl}}.$$

Since, $v(1) = 1$, we can solve $v(l + \kappa)$. Plugging in its value, we obtain (11). \square

4 Convergence of the invariant distributions

For any $\alpha > 0$, the Markov chain $\{W_n\}$ has a unique invariant distribution, this can easily be shown by noting that state 1 can be reached from any other state, and by applying a standard Foster-type drift criterion. Since it is possible to stay in state 1 with positive probability, it follows that this Markov chain has a unique weak limit, which we denote by W . Given the results in Sections 2 and 3, it is tempting to conclude that $W^\alpha = f(\alpha)W$ converges in distribution to the unique invariant distribution (in $(0, \infty)$) of the Markov process $\{V_t\}$ studied in Section 3. That this is indeed the case is the main contribution of this section.

Theorem 7. *Suppose $E[B^2] < \infty$. Then $W^\alpha \xrightarrow{d} V$, with $V \in (0, \infty)$ the unique random variable whose distribution is invariant for the Markov chain with generator (6).*

The theorem implies that, when V is also the weak limit of the Markov chain specified by (6), it is allowed to interchange the steady-state limit, and the limit $\alpha \downarrow 0$. Interestingly, the

limiting process can be periodic as noted in Section 3, while the original process is not. So the distribution of V is not always the limiting distribution of the Markov process $\{V_t\}$.

To prove Theorem 7 requires some work. Recently, this problem received attention by Gamarnik and Zeevi [9] in the context of generalized Jackson networks in heavy traffic. In the context of TCP models there are two related studies. Dumas *et al.* [5] (which precedes [9]) treats a special AIMD case of our model, with $\kappa = 0$, $h = 1$, and $B = 1$. The property $\kappa = 0$ makes the problem tractable since one can “escape” from 0 to x in linear time. This is not the case in our model where 0 can be a trap in the limiting Markov chain. Therefore, the major problem is to show that W^α is in some sense bounded away from 0 uniformly in α (we make this precise later on). Similar problems were encountered in the related study of Ott and Swanson [20]. These authors came up with several elegant ideas of which some are exploited here. We believe that the analysis in this section has its own merits. The main additional technical problem we encounter here is that the downward jumps can be bigger than 1. To adequately handle this we need to control the size of the jumps by requiring $E[B^2] < \infty$; we refer to the body of the proof for more details.

The main technical results required to prove Theorem 7 are given in the following lemma’s.

Lemma 8. *For $\alpha < 1$, and any $\epsilon > 0$, there exists a constant δ_ϵ (not dependent on α) such that*

$$P(W^\alpha > \delta_\epsilon) > 1 - \epsilon. \quad (16)$$

Lemma 9. *Let α^* be small enough such that $\inf_{\alpha < \alpha^*} 1 - E[\delta^B] - f(\alpha) > 0$. For any $\epsilon > 0$, there exists a constant M^ϵ (not dependent on α) such that*

$$P(W^\alpha < M^\epsilon) > 1 - \epsilon, \quad (17)$$

for $\alpha \in (0, \alpha^*)$.

Proof of Theorem 7. Based on Lemma 8 and 9, we prove convergence of the stationary distributions. Let $\alpha \downarrow 0$ through a sequence and let $\{W^\alpha\}$ be the associated sequence of random variables. Lemma 8 and 9 show that $\{W^\alpha\}$ is relatively compact in $(0, \infty)$. It therefore suffices to show that each converging subsequence has the same distribution of V , with V the unique invariant distribution of the Markov chain with generator (6). Let V_0 be the limit point of a converging subsequence $\{W^{\alpha'}\}$ and let $W_0^{\alpha'} = W^{\alpha'}$. From Theorem 1 we deduce that $(W_0^{\alpha'}, W_t^{\alpha'}) \xrightarrow{d} (V_0, V_t)$ as $\alpha' \downarrow 0$. Since the choice of $W_0^{\alpha'}$ makes the process $\{W_t^{\alpha'}\}$ stationary, we see that $W_t^{\alpha'} \stackrel{d}{=} W_0^{\alpha'}$. Consequently, $V_0 \stackrel{d}{=} V_t$ for any $t > 0$. This implies that the distribution of V_0 is invariant for the Markov chain with generator (6). Since there is a unique such invariant distribution, the statement follows. \square

4.1 Proof of Lemma 8

In this section we prove Lemma 8. The proof of this lemma is similar to the proof of Theorem 2.3 in [20], although some of the details are different (in particular, working with the Lyapounov function $u(x) = 1/x$ would require the assumption $E[\delta^{-B}] < \infty$). As in [20], the main goal is to transform our problem into a case that can be handled by Corollary 4.9.8 of [6]. The first main idea is the following. Let $\{N(t), t \geq 0\}$ be a Poisson process with rate 1 independent of the Markov chain $\{W_n\}$, and consider the Markov chain

$$\tilde{W}_t^\alpha = f(\alpha)W_{N(t/\sigma(\alpha))}, t > 0.$$

It is clear that the steady-state distribution of $\{\tilde{W}_t^\alpha\}$ is identical to the distribution of W^α . A straightforward computation shows that this continuous-time Markov process has the following generator:

$$\tilde{A}_\alpha u(x) = e^{-\sigma(\alpha)x^h} \frac{u(f(\alpha)\lceil(\frac{x+\sigma(\alpha)x^\kappa}{f(\alpha)}\rceil) - u(x)}{\sigma(\alpha)} + \frac{1 - e^{\sigma(\alpha)x^h}}{\sigma(\alpha)} (E[u(f(\alpha)\lceil\frac{x\delta^B}{f(\alpha)}\rceil)] - u(x)). \quad (18)$$

Consider the function u given by $u(x) = (\log x)^2$ if $x < 1$ and $u(x) = 0$ if $x \geq 1$. This function is non-increasing and convex. Applying these properties, and exploiting that $1 - e^{-y} \leq y$, we see that, for $x < 1/(1 + \sigma(\alpha))$,

$$\begin{aligned} \tilde{A}_\alpha u(x) &\leq e^{-\sigma(\alpha)x^h} \frac{u(x + \sigma(\alpha)x^\kappa) - u(x)}{\sigma(\alpha)} + \frac{1 - e^{-\sigma(\alpha)x^h}}{\sigma(\alpha)} (E[u(x\delta^B)] - u(x)) \\ &\leq e^{-\sigma(\alpha)x^h} x^\kappa u'(x + \sigma(\alpha)x^\kappa) + x^h (E[u(x\delta^B)] - u(x)) \\ &= e^{-\sigma(\alpha)x^h} x^{\kappa-1} \frac{2x}{x + \sigma(\alpha)x^\kappa} \log(x + \sigma(\alpha)x^\kappa) + x^h E[B^2] (\log \delta)^2 - 2E[B] \log(1/\delta) x^h. \end{aligned}$$

Since $\sigma(\alpha)$ is bounded in α , since $E[B^2]$ is assumed to be finite, and most importantly since $h > \kappa - 1$, it is clear that, for every $m < \infty$, there exists an $x_m > 0$ independent of α , such that, for all $x < x_m$,

$$\sup_{\alpha \in (0,1), x \in (0, x_m)} \tilde{A}_\alpha u(x) \leq -m. \quad (19)$$

Since $u(x) = 0$ if $x > 1$, it follows that

$$\sup_{\alpha \in (0,1), x \in (0, \infty)} \tilde{A}_\alpha u(x) =: M < \infty. \quad (20)$$

Applying Corollary 4.9.8 of [6], we conclude that

$$P(W^\alpha > x_m) \geq \frac{m}{m + M}. \quad (21)$$

As this estimate holds for any $m < \infty$, the statement of the lemma follows.

4.2 Proof of Lemma 9

We use the same idea as in the previous subsection, but now take $u(x) = x$. It easily follows from (18) that

$$\tilde{A}_\alpha u(x) \leq \frac{f(\alpha)}{\sigma(\alpha)} + x^\kappa - x^{h+1} (1 - E[\delta^B] - f(\alpha)).$$

Let α^* be such that $(1 - E[\delta^B] - f(\alpha)) > \xi$ for a suitable $\xi > 0$ whenever $\alpha < \alpha^*$; which is possible since $f(\alpha)$ is monotone). Since $\frac{f(\alpha)}{\sigma(\alpha)}$ is bounded as $\alpha \downarrow 0$, for every $m < \infty$, it is possible to find a constant x^m (independent of $\alpha < \alpha^*$) such that $\tilde{A}_\alpha u(x) \leq -m$ for every $x > x^m$. Observe that

$$\sup_{\alpha \in (0, \alpha^*), x \in (0, \infty)} \tilde{A}_\alpha u(x) =: M_2 < \infty. \quad (22)$$

By applying Corollary 4.9.8 of [6], we conclude that

$$P(W^\alpha > x_m) \leq \frac{M_2}{m + M_2}, \alpha < \alpha^*. \quad (23)$$

Since this estimate holds for any $m < \infty$, the statement of the Lemma follows.

5 Concluding remarks

The theory developed in this paper ignores the fact that the loss rate is also a function of the throughput. Several studies in the AIMD context, see for example [10] and references therein, combine the square root formula with fixed point approximations for the loss rate in terms of the throughput to obtain approximations for the throughput and loss rate. Such fixed point approximations may be combined with the type of approximations developed in this paper to determine which scheme provides optimal throughput.

References

- [1] E. Altman, K. Avrachenkov, C. Barakat and R. Núñez Queija. (2002). State-dependent M/G/1 type queueing analysis for congestion control in data networks. *Computer Networks* **39**, 789–808.
- [2] E. Altman, K. Avrachenkov, C. Barakat, A.A. Kherani and B. Prabhu (2005). Analysis of MIMD congestion control algorithm for high speed networks. *Computer Networks* **48**, 972–989.
- [3] E. Altman, K. Avrachenkov, C. Barakat, A.A. Kherani and B. Prabhu (2005). Performance Analysis and Stochastic Stability of Congestion Control Protocols. In: *Proceedings of IEEE Infocom*, Miami, USA.
- [4] D. Bansal and H. Balakrishnan (2001). Binomial congestion control algorithms. In: *Proceedings of IEEE Infocom*, Anchorage, USA.
- [5] V. Dumas, F. Guillemin and P. Robert (2002). A Markovian analysis of additive-increase multiplicative-decrease algorithms. *Advances in Applied Probability* **34**, 85–111.
- [6] S. Ethier and T. Kurtz. (1986). *Markov Processes: Characterization and Convergence*. Wiley, New York.
- [7] S. Floyd (1991). Connections with multiple congested gateways in packet-switched networks part 1: One-way traffic. *Computer Communications Review* **21** 5 23–30.
- [8] S. Floyd. (2002). Highspeed TCP for Large Congestion Windows.
<http://www.icir.org/floyd/papers/draft-floyd-tcp-highspeed-01.txt>
- [9] D. Gamarnik and A. Zeevi. (2006). Validity of heavy traffic steady-state approximation in generalized Jackson networks. *Annals of Applied Probability* **16**, 56–90.
- [10] R.J. Gibbens, S.K. Sargood, C. Van Eijl, F.P. Kelly, H. Azmoodeh, R.N. Macfadyen and N.W. Macfadyen (2000). Fixed-point models for the end-to-end performance analysis of IP networks. In: *13th ITC Specialist Seminar: IP Traffic Measurement, Modeling and Management*, Monterey, California
- [11] F. Guillemin, P. Robert and B. Zwart (2004). AIMD algorithms and exponential functionals. *Annals of Applied Probability* **14**, 90–117.
- [12] V. Jacobson. (1988). Congestion avoidance and control. In: *Proceedings of SIGCOMM '88*.
- [13] F. P. Kelly, A.K. Maulloo and D.K.H. Tan (1998). Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* **49**, 237–252.
- [14] T. Kelly (2003). Scalable TCP: Improving performance in Highspeed wide area networks. *Computer Communications Review* **32**, 83–91.
- [15] L. Massoulié (2006). Structural properties of proportional fairness: stability and insensitivity To appear in *Annals of Applied Probability*.

- [16] M. Mathis, J. Semke, J. Mahdavi and T.J. Ott (1997). The macroscopic behavior of the TCP congestion avoidance algorithm. *Computer Communication Review* **27**.
- [17] K. Maulik and B. Zwart. (2006). Tail asymptotics for exponential functionals of Levy processes. *Stochastic Processes and their Applications* **116**, 156–177.
- [18] S. Meyn and R. Tweedie. (1994). *Markov Chains and Stochastic Stability*.
<http://decision.csl.uiuc.edu/~meyn/pages/book.html>
- [19] T.J. Ott (2005). Transport Protocols in the TCP Paradigm and their Performance. *Telecommunication Systems* **30**, 351–385.
- [20] T.J. Ott and J. Swanson (2006). Asymptotic behavior of a generalized TCP congestion avoidance algorithm. Preprint; see <http://front.math.ucdavis.edu/math.PR/0608476>.
- [21] J. Padhye, V. Firoiu, D. Towsley and J. Kurose (2000). Modeling TCP Reno Performance: A Simple Model and its Empirical Validation *IEEE/ACM Transactions on Networking* **8**, 133–145.
- [22] V. Paxson (1999). End-to-end internet packet dynamics, *IEEE/ACM Transactions on Networking* **7**, 277–292.