Stationarity and existence of moments for some processes in Transport Protocols

Teunis J. Ott Jason Swanson

August 27, 2008

Abstract

This note establishes stationarity of a number of stochastic processes of interest in the study of Transport Protocols, and determines which moments of the stationary distributions are finite. For many of the processes studied in this note stationarity had been established before, but for one class the result is new. For that class, it was counterintuitive that stationarity was hard to prove. This note also explains why that class offered such stiff resistance.

The stationarity is proven using Liapunov functions, without first proving tightness by proving boundedness of moments. A different method then uses these results to prove existence of certain moments.

1 Introduction

The past 12 years has seen a large number of papers modeling aspects of TCP (Transport Control Protocol) and TCP–like protocols. We mention [1, 2, 9] as examples of earlier papers.

An excellent overview of the literature in this field up to about 2004 can be found in [10], and [8] gives a large (but almost by definition insufficient) number of references after 2004.

The paper [3] proposes a class of "TCP–like" Internet Transport Protocols and uses a class of stochastic processes to analyze the performance of these protocols, in particular the evolution over time of the window size. Over the past few years many results have been proven for that class of processes, such as (in most cases) stationarity of the processes, also in many cases weak convergence, if the drop probability goes to zero, of rescaled versions of the window size processes to a limit process, and in that case convergence of stationary distributions to the limiting stationary distribution, etc.

This note in a sense completes the picture. Stationarity of processes had been proven in many, but not all, cases. Several different techniques were used to prove stationarity, often involving proving boundedness of certain moments and thus tightness. Those methods did not always work.

This note proves stationarity in all cases. It uses a single (different) method (Liapunov functions, which bound expected return times to some compact set). For certain of those stationary distributions this note also finds out exactly which moments are finite.

There is at least one minor but intriguing question left open. For one class of processes existence of stationary distributions has been proven, but uniqueness of such a stationary distribution has not been established yet. This, and other, open issues are discussed in Section 8.

Some of the results in this note were presented at the 2006 MAMA workshop, and appeared in abbreviated form in [6].

2 The processes involved

The class of stochastic processes introduced in [3] is defined by:

Let $(U_n)_{n=0}^{\infty}$ be independent, identically distributed random variables, each distributed uniformly $[0,1]$. Let p be a probability, $0 < p < 1$. Define the i.i.d. random variables $\chi_{p,n}$ by

$$\chi_{p,n} = \begin{cases} \text{success} & \text{if } U_n \geq p \\ \text{failure} & \text{if } U_n < p \end{cases} \tag{2.1}$$

Further, let the discrete time, continuous state space process $W^*_{\alpha,p,C,n}$ $(n = 0, 1, 2, \cdots,$ $0 < W^*_{\alpha,p,C,n} < \infty, \quad 0 < p < 1)$ be defined by

$$W^*_{\alpha,p,C,n+1} = \begin{cases} W^*_{\alpha,p,C,n} + c_1 (W^*_{\alpha,p,C,n})^{\alpha} & \text{if } \chi_{p,n} = \text{success}, \\ \max\left(W^*_{\alpha,p,C,n} - c_2 (W^*_{\alpha,p,C,n})^{\beta}, C\right) & \text{if } \chi_{p,n} = \text{failure}, \end{cases} \tag{2.2}$$

where $\alpha < \beta \leq 1,\ c_1 > 0,\ c_2 > 0,\ C > 0$.

The special case with $\beta = 1,\ \alpha = -1,\ c_1 = 1, c_2 = \frac{1}{2}$ and (for example) $C = 1$ models "classical TCP".

The special case with $\beta = 1,\ \alpha = 0$ models Tom Kelly's "Scalable TCP", see [13, 14].

[3] shows that the more general case, even the case $0 < \alpha < \beta \leq 1$, is of interest in the study of transport protocols.

[7] proves that for all values $\alpha < \beta \leq 1,\ c_1 > 0,\ c_2 > 0,\ C > 0,\ 0 < p < 1$ (and $0 < c_2 < 1$ if $\beta = 1$) the process $W^*_{\alpha,p,C,n}$ has a unique stationary distribution. The uniqueness of that stationary distribution (independent of $W^*_{\alpha,p,C,0}$) is derived from the fact that eventually $W^*_{\alpha,p,C,n} = C$ for some (possibly large) n.

[4] mainly studies the case $\alpha < \beta = 1,\ c_1 > 0,\ 0 < c_2 < 1$. In that case we write $1 - c_2 = b$. In that case we can drop the "max (..., C)" in (2.2) (or choose $C = 0$). [4, 5] also draw some conclusions, from the case "$C = 0$", for the case "$C > 0$".

The process of main interest in that paper therefore is defined by

$$W_{\alpha,p,n+1} = \begin{cases} W_{\alpha,p,n} + c_1(W_{\alpha,p,n})^\alpha & \text{if } \chi_{p,n} = \text{success}, \\ W_{\alpha,p,n} - (1-b)W_{\alpha,p,n} = bW_{\alpha,p,n} & \text{if } \chi_{p,n} = \text{failure}. \end{cases} \tag{2.3}$$

but it also draws some conclusions for the process $(W^{**}_{\alpha,p,C,n})_{n=0}^\infty$ defined by

$$W^{**}_{\alpha,p,C,n+1} = \begin{cases} W^{**}_{\alpha,p,C,n} + c_1(W^{**}_{\alpha,p,C,n})^\alpha & \text{if } \chi_{p,n} = \text{success}, \\ \max(bW^{**}_{\alpha,p,C,n}, C) & \text{if } \chi_{p,n} = \text{failure}. \end{cases} \tag{2.4}$$

By an abuse of notation we will often denote $W^{**}_{\alpha,p,C,n}$ as $W^{**}_{p,n}$ and $W_{\alpha,p,n}$ as $W_{p,n}$, etc, and when the parameter n is dropped we assume the random variable has the stationary distribution.

[7] proves existence and uniqueness of the stationary distribution of the process $(W^{*}_{\alpha,p,C,n})_{n=0}^\infty$ in the case $C > 0$. In addition, [4] proves existence and uniqueness of stationary distributions in case ($0 \leq \alpha < \beta = 1$ **and** $C = 0$, $0 < c_2 < 1$). The results in the latter paper can also be used (using the tightness proven in that paper) to prove existence, but not uniqueness, of a stationary distribution in the case ($\alpha < 0$, $\beta = 1$, $C = 0$, $0 < c_2 < 1$ **and** $0 < p < 1$ sufficiently small), see [5] for details.

The latter extra requirement, that p is sufficiently small, is counterintuitive: It "should" be easier to prove stationarity for p close to one (therefore fewer successes, therefore (?) $W_{\alpha,p,n}$ more likely to be small (?)) than for p close to zero. More on this topic later in this paper.

[3, 4, 7] study rescaled versions of the processes described above. The cases $\beta = 1$ and $\beta < 1$ require different rescalings. [7] uses methods from [11] to obtain weak limit results for the rescaled processes (weak convergence of processes to a limit process) if $p \downarrow 0$, in the cases $\beta = 1$ as well as $\beta < 1$. [4] concentrates on the case $\beta = 1$ and obtains weak convergence results for stationary distributions of the rescaled processes, including

stochastic dominance results and rate of convergence results. Both papers leave a gap in the case $\beta = 1, C = 0, \alpha < 0$.

This note fills the the gap in the analysis of stationarity: By an alternative method (using Liapunov functions, and presented in the Sections 3–5 of this note) it proves stationarity in all cases $(\beta < 1, C > 0)$ as well as $(\beta = 1, C \geq 0)$. However, the greater generality comes at a cost: the proofs using Liapunov functions do not prove existence, let alone boundedness, of the various moments of interest. This approach was presented at the 2006 MAMA workshop, see [6].

The approach in [4] (basically) is to first prove, in the case $\beta = 1, C = 0$, tightness by proving finiteness and boundedness of certain moments of the rescaled processes, and then observe that the boundedness of moments implies existence of a stationary distribution. The approach in the sections 3–5 of this note proves existence of stationary distributions in all cases $(\beta < 1, C > 0)$ and $(\beta = 1, C \geq 0)$ without obtaining insight in existence of moments.

In the Sections 6 and 7 of this note stochastic dominance results from [4] are combined with the proven stationarity to obtain necessary and sufficient conditions for existence of moments of the stationary distribution of W_p in the case $\beta = 1, C = 0$.

The results in this note are formulated for the original processes $W^*_{\alpha,p,C,n}$ (etc), not for the rescaled processes. Translation is easy.

The method of proof in the Sections 3–5 of this note is to find a compact set $[v_1, v_2]$ and to prove that the "expected first return time" from $W^*_{\alpha,p,C,n}$ (etc) leaving that set to returning to that set is bounded. One of the results used is Theorem 12.3.4 on page 296 of [12].

In the case $C > 0$ we will choose $v_1 = C < v_2 < \infty$. In the case $(\beta = 1,\ C = 0,\ 0 < c_2 < 1)$ we will choose $0 < v_1 < v_2 < \infty$. In the old "holdout situation" $\alpha < 0$ this is necessary to make the proof work.

3 The case $\beta < 1, \;\; C > 0$

Throughout this section we have $\beta < 1$, $C > 0$ and we choose $v_1 = C$ and v_2 "large", to be described later. Among others we require that v_2 is large enough to ensure that $w - c_2 w^\beta$ is increasing in w for $w \geq v_2$, that $C + c_1 C^\alpha < v_2$, and that $w - c_2 w^\beta > C$ for all $w \geq v_2$.

We will find a function (Liapunov function) $V : [C, \infty) \to [0, \infty)$ with the following properties:

$$V(w) = 0 \;\text{ for }\; C \leq w \leq v_2,$$

$$V(w) \geq 1 + (1-p)V(w + c_1 w^\alpha) + pV(w - c_2 w^\beta) \;\text{ for }\; w > v_2, \tag{3.1}$$

and such that there is an upper bound $B < \infty$ with the property that

$$V(w + c_1 w^\alpha) < B \;\text{ for all }\; C \leq w \leq v_2. \tag{3.2}$$

(3.1) shows that $V(w)$ is an upper bound for the Expected First Passage Time from $W^*_{\alpha,p,C,n} = w > v_2$ to $W^*_{\alpha,p,C,n+T} \leq v_2$. This is Theorem 11.3.4 on page 265 of [12]. (3.2) then shows that the process $(W^*_{\alpha,p,C,n})_{n=0}^\infty$ has (at least one) stationary distribution. That is Theorem 12.3.4 on page 296 of [12].

Once we have the results above, it is obvious that the expected first passage time from "anywhere" to $W^*_{\alpha,p,C,n} = C$ is finite. This then proves the uniqueness of the stationary distribution.

We find $w_2 > v_2$ such, that $w_2 - c_2 w_2^\beta = v_2$. In fact, see below, one could say we choose w_2 large enough and then define $v_2 = w_2 - c_2 w_2^\beta$. Then we choose the function V of the form

$$V(w) = \nu + \mu w^{1-\beta} \text{ for } w > v_2. \tag{3.3}$$

Clearly, it now is sufficient to choose μ and ν such, that

$$\mu w^{1-\beta} \geq 1 + (1-p)\mu(w + c_1 w^{\alpha})^{1-\beta} + p\mu(1 - c_2 w^{\beta})^{1-\beta} \text{ for } w \geq w_2 \tag{3.4}$$

and

$$\nu \geq \frac{1}{p} \sup_{v_2 < w < w_2} \left(1 + (1-p)\mu(w + c_1 w^{\alpha})^{1-\beta} - \mu w^{1-\beta}\right). \tag{3.5}$$

For (3.4) to hold we choose

$$\mu > \frac{1}{p(1-\beta)c_2}. \tag{3.6}$$

By a simple Binomial expansion we see that for (3.4) to hold, w_2 must be chosen at least equal to, or larger than, w_{min}, where (roughly)

$$w_{min} \sim \left(\frac{(1-p)(1-\beta)c_1}{p(1-\beta)c_2 - \frac{1}{\mu}}\right)^{\frac{1}{\beta-\alpha}}. \tag{3.7}$$

The approach used has the disadvantage that it requires a special choice of v_2, namely, v_2 quite large. With more work smaller choices of v_2 can be obtained, but that extra work thus far has not led to sufficiently interesting results.

It now is obvious that there is a $B < \infty$ for which (3.2) holds.

4 The case $\beta = 1, C > 0$

Throughout this section we have $\beta = 1$, therefore $0 < c_2 < 1$, and we still have $C > 0$. We also have $b = 1 - c_2$. The development in this section parallels that in the previous section. We choose $v_1 = C$, v_2 sufficiently large, w_2 as in the previous section. In this section that means that $v_2 = bw_2$. We choose the function V to be of the form

$$V(w) = \nu + \mu \log w \quad \text{for} \quad w > v_2. \tag{4.1}$$

For (3.1) to hold we now need that

$$\mu \log(w) \geq 1 + (1-p)\mu \log(w + c_1 w^\alpha) + p\mu \log(bw) \quad \text{for} \quad w \geq w_2 \tag{4.2}$$

and that

$$\nu \geq \frac{1}{p} \sup_{v_2 < w < w_2} \left(1 + (1-p)\mu \log(w + c_1 w^\alpha) - \mu \log(w)\right). \tag{4.3}$$

For (4.2) to hold we choose

$$\mu > \frac{1}{p|\log(b)|} \tag{4.4}$$

and a simple expansion shows that for (4.2) to hold we must choose $w_2 > w_{min}$, where (roughly)

$$w_{min} \sim \left(\frac{(1-p)c_1}{p|\log(b)| - \frac{1}{\mu}}\right)^{\frac{1}{1-\alpha}}. \tag{4.5}$$

5 The case $\beta = 1$, $C = 0$

In this section we study the case $\beta = 1$ with $C = 0$. We also have $0 < c_2 < 1$ and $b = 1 - c_2$. In this case different approaches are necessary for the cases $\alpha < 0$, $\alpha = 0$, and $0 < \alpha < 1$. The only interesting situation, however, is $\alpha < 0$ because it illustrates why the original approach could not be extended to values of p close to one.

Thus, in most of this section we have $\alpha < 0$. In that case, $w + c_1 w^\alpha$ is minimal for $w = (c_1|\alpha|)^{\frac{1}{1+|\alpha|}} = w^*$. It is decreasing in w for $0 < w < w^*$ and increasing in w for $w > w^*$. It goes to infinity both for $w \downarrow 0$ and $w \uparrow \infty$.

We now choose v_1 and v_2 such, that $0 < v_1 < w^* < v_2 < \infty$, such that also $bv_2 > v_1$ and $v_1 + c_1 v_1^\alpha > v_2$, and such that also $v_1 < 1 < v_2$, and with some additional constraints, see below.

Thus, we insure that in order for the process $W_{\alpha,p,n}$ to move from the set (v_2, ∞) to the set $(0, v_1)$ it must pass through the set $[v_1, v_2]$, and in order for the process to pass from the set $(0, v_1)$ to the set $[v_1, v_2]$ it must first jump over the set $[v_1, v_2]$ into the set (v_2, ∞) and then, as in Section 4, drift down to the set $[v_1, v_2]$.

We choose the function $V(.)$ as

$$V(w) = \begin{cases} \nu_l + \mu_l |\log(w)| & \text{for } 0 < w < v_1, \\ 0 & \text{for } v_1 \le w \le v_2, \\ \nu_u + \mu_u \log(w) & \text{for } v_2 < w < \infty. \end{cases} \tag{5.1}$$

(l and u stand for "lower" and "upper"). ν_u and μ_u are chosen as in Section 4, and an additional lower bound for v_2 is obtained as in that same section. The critical inequality now becomes

$$\nu_l + \mu_l |\log(w)| \ge 1 + (1-p)\left(\nu_u + \mu_u \log(w + c_1 w^\alpha)\right) + p\left(\nu_l + \mu_l |\log(bw)|\right) \tag{5.2}$$

for all $0 < w < v_1$. This can be rewritten as

$$\begin{aligned} (1-p)(\nu_l - \nu_u) - p\mu_l |\log(b)| + (1-p)\left(\mu_l - \mu_u|\alpha|\right) |\log(w)| \ge \\ 1 + (1-p)\mu_u \left(\log(c_1) + \log(1 + \frac{w^{1+|\alpha|}}{c_1}) \right) \end{aligned} \tag{5.3}$$

for all $0 < w < v_1$. For given ν_u and μ_u, it is easy to choose ν_l, μ_l and w_1 such, that this holds. For example, we can take $\mu_l = \mu_u|\alpha|$, etc.

In the situation of this section we can leave the set $[v_1, v_2]$ in two ways: by jumping up past v_2 and by jumping down past v_1. It remains easy to prove that the expected return time remains bounded. That would not have been the case had we chosen $v_1 = 0$.

The approach in [4] in the case ($\alpha < 0$, $\beta = 1$, $C = 0$) proves existence of a stationary distribution only in the case p sufficiently small, but in that case also proves that every such stationary distribution has a finite first moment. The approach in the Sections 3–5 of this note works for all $0 < p < 1$ but does not prove existence of moments.

The cases $\alpha = 0$ and $0 < \alpha < 1$ in the situation of this section are easy to handle. In that situation the process $W_{\alpha,p,n}$ can not "jump over" the set $[v_1, v_2]$ (as long as v_1 is reasonably small and v_2 is reasonably large) and the analysis of $V(.)$ is split into two independent subproblems: one for $w > v_2$, one for $w < v_1$. The subproblem for $w > v_2$ remains as before.

In the case $0 < \alpha < 1$ we obtain for the subproblem $w < v_1$ that

$$V(w) = \nu_l + \mu_l \log\left(|\log(w)|\right) \tag{5.4}$$

satisfies if $v_1 > 0$ is sufficiently small.

In the case $\alpha = 0$ we can even choose the function $V(.)$ bounded on $0 \le w < v_1$.

6 Moments of $W_{\alpha,p}$ if $\beta = 1, C = 0$: Preliminaries

In this section and the next we study the case ($\beta = 1, C = 0$) and prove for the (or, strictly speaking, "any") stationary distribution of $W_{\alpha,p}$ that if $\alpha < 0$ then (still $0 < c_2 < 1$ and $b = 1 - c_2$)

$$E[W_{\alpha,p}^{\nu}] \begin{cases} = \infty & \text{if } \nu \ge \left|\frac{\log(p)}{\alpha\log(b)}\right|, \\ < \infty & \text{if } -\left|\frac{\log(p)}{\log(b)}\right| < \nu < \left|\frac{\log(p)}{\alpha\log(b)}\right|, \\ = \infty & \text{if } \nu \le -\left|\frac{\log(p)}{\log(b)}\right|, \end{cases} \tag{6.1}$$

while if $0 \le \alpha < 1$ then

$$E[W_{\alpha,p}^{\nu}] \begin{cases} < \infty & \text{if } \nu > -\left|\frac{\log(p)}{\log(b)}\right|, \\ = \infty & \text{if } \nu \le -\left|\frac{\log(p)}{\log(b)}\right|. \end{cases} \tag{6.2}$$

The proof utilizes a number of auxiliary processes defined in [4]:

The process $V_{\alpha,p,n} = \frac{p}{(1-\alpha)c_1}(W_{\alpha,p,n})^{1-\alpha}$, which also satisfies

$$V_{\alpha,p,n+1} = \begin{cases} V_{\alpha,p,n} + p + R_{\alpha,p}(V_{\alpha,p,n}) & \text{if } \chi_{p,n} = \text{success}, \\ cV_{\alpha,p,n} & \text{if } \chi_{p,n} = \text{failure}, \end{cases} \tag{6.3}$$

where $c = b^{1-\alpha}$ and the function $R_{\alpha,p}(.)$ is defined by

$$R_{\alpha,p}(v) = v\left(1 + \frac{p}{(1-\alpha)v}\right)^{1-\alpha} - v - p. \tag{6.4}$$

It is easily seen, and proven in [4], that for $\alpha < 0$ $R_{\alpha,p}(.)$ is a positive, decreasing function (in fact a completely monotone function).

A special role is played by the process $(X_{p,n})$, which evolves as

$$X_{p,n+1} = \begin{cases} X_{p,n} + p & \text{if } \chi_{p,n} = \text{success}, \\ cX_{p,n} & \text{if } \chi_{p,n} = \text{failure}. \end{cases} \tag{6.5}$$

($X_{p,n}$ still depends on α through $c = b^{1-\alpha}$).

We will also use the process $\Lambda_{\alpha,p,n}$ defined by

$$\Lambda_{\alpha,p,n+1} = \begin{cases} \Lambda_{\alpha,p,n} + p + R_{\alpha,p}(X_{p,n}) & \text{if } \chi_{p,n} = \text{success}, \\ c\Lambda_{\alpha,p,n} & \text{if } \chi_{p,n} = \text{failure}. \end{cases} \tag{6.6}$$

Often, the processes $V_{\alpha,p,n}$, $X_{p,n}$ and $\Lambda_{\alpha,p,n}$ are given well-chosen (possibly random) initial values $V_{\alpha,p,0}$, $X_{p,0}$ and $\Lambda_{\alpha,p,0}$ and then evolve as in (6.3), (6.5), (6.6).

Results from [4] we will use are that in that case

(1) If $0 \leq \alpha < 1$ and for some n $V_{\alpha,p,n} \leq X_{p,n}$ then $V_{\alpha,p,m} \leq X_{p,m}$ for all $m \geq n$, while

(2) If $\alpha \leq 0$ and for some n $X_{p,n} \leq V_{\alpha,p,n} \leq \Lambda_{\alpha,p,n}$ then $X_{p,m} \leq V_{\alpha,p,m} \leq \Lambda_{\alpha,p,m}$ for all $m \geq n$.

The flow of the logic is that the Sections 3–5 of this paper prove stationarity of the processes $W_{\alpha,p,n}$ and thereby also of $V_{\alpha,p,n}$. The required stochastic dominance relationships were proven in [4], and the combined results are utilized in this section and the next to prove (6.1) and (6.2).

Stationarity and uniqueness of the stationary distribution of the process $X_{p,n}$ was established in [4], which also gives that stationary distribution. The exact shape of that stationary distribution will not be needed here.

It is clear from (6.6) that if the process $\Lambda_{\alpha,p,n}$ has a stationary distribution then the process $(X_{p,n}, \Lambda_{\alpha,p,n})_{n=0}^{\infty}$ has a unique joint stationary distribution. As by–product of this section and the next we will see that this is indeed the case.

The random variables $W_{\alpha,p}$, $V_{\alpha,p}$, X_p and $\Lambda_{\alpha,p}$ will always have the stationary distribution of the corresponding processes. The results above show that if $0 < \alpha < 1$ then $V_{\alpha,p}$ and X_p have a joint distribution with $P\{V_{\alpha,p} \leq X_p\} = 1$, while if $\alpha < 0$ X_p, $V_{\alpha,p}$ and $\Lambda_{\alpha,p}$ have a joint distribution with $P\{X_p \leq V_{\alpha,p} \leq \Lambda_{\alpha,p}\} = 1$. The possibility that $P\{\Lambda_{\alpha,p} = \infty\} = 1$ can not yet be excluded, but a by–product of the arguments in the next section will show that in fact also the process $(\Lambda_{\alpha,p,n})$ is stationary for all $0 < p < 1$.

A third result from [4] we will use is

(3) There exists a random variable Z (which is elaborated on in [4]) which has a joint distribution with X_p with $P\{X_p < Z\} = 1$ and for which $E[Z^s] < \infty$ for all $-\infty < s < +\infty$.

In addition to the processes above, the proofs utilize the stochastic process $(L_{\alpha,p,n})_{n=0}^{\infty}$ defined by

$$L_{\alpha,p,n+1} = \begin{cases} c_1 L_{\alpha,p,n}^{\alpha} & \text{if } \chi_{p,n} = \text{success}, \\ b L_{\alpha,p,n} & \text{if } \chi_{p,n} = \text{failure}, \end{cases} \tag{6.7}$$

The process $\log(L_{\alpha,p,n})$ obeys

$$\log(L_{\alpha,p,n+1}) = \begin{cases} \alpha \log(L_{\alpha,p,n}) + \log(c_1) & \text{if } \chi_{p,n} = \text{success}, \\ \log(L_{\alpha,p,n}) + \log(b) & \text{if } \chi_{p,n} = \text{failure}. \end{cases} \tag{6.8}$$

It is easily seen that the process $L_{\alpha,p,n}$ is stationary if and only if $|\alpha| < 1$. In that case the stationary distribution is given by

$$E[L^s_{\alpha,p}] = E[exp\{s \log(L_{\alpha,p})\}] = c_1^{\frac{s}{1-\alpha}} \prod_{k=0}^{\infty} \left(\frac{1-p}{1-pb^{\alpha^k s}} \right), \tag{6.9}$$

which if $pb^s \geq 1$, i.e. $s \leq -\left|\frac{\log(p)}{\log(b)}\right|$, is interpreted as $+\infty$.

If $\alpha = 0$ (6.9) reduces to

$$E[L^s_{0,p}] = c_1^s \frac{1-p}{1-pb^s} \quad (\text{for } \ s > -\left|\frac{\log(p)}{\log(b)}\right|). \tag{6.10}$$

It is clear that if $0 \leq \alpha < 1$ and for some n $L_{\alpha,p,n} \leq W_{\alpha,p,n}$ then $L_{\alpha,p,m} \leq W_{\alpha,p,m}$ for all $m \geq n$ and therefore the stationary random variable $L_{\alpha,p}$ is stochastically less than the stationary random variable $W_{\alpha,p}$.

From (2.3) we see that for the stationary random variable $W_{\alpha,p}$, for all $-\infty < s < +\infty$,

$$E[W^s_{\alpha,p}] = pb^s E[W^s_{\alpha,p}] + (1-p)E[(W_{\alpha,p} + c_1 W^{\alpha}_{\alpha,p})^s], \tag{6.11}$$

possibly $\infty = \infty$. This leads to a number of simple observations:

Observation 1: If

$$E[(W_{\alpha,p} + c_1 W^{\alpha}_{\alpha,p})^s] = \infty \tag{6.12}$$

then

$$E[W^s_{\alpha,p}] = \infty$$

Observation 2: If

$$E[W^s_{\alpha,p}] < \infty \tag{6.13}$$

then

$$E[(W_{\alpha,p} + c_1 W_{\alpha,p}^{\alpha})^s] = \frac{1 - pb^s}{1 - p} E[W_{\alpha,p}^s] < \infty. \tag{6.14}$$

Since clearly $0 < E[W_{\alpha,p}^s] \leq \infty$ (etc), observation 2 leads to

Observation 3: If

$$pb^s \geq 1, \quad \text{i.e.} \ \ s \leq -\left|\frac{\log{(p)}}{\log{(b)}}\right| \tag{6.15}$$

then

$$E[W_{\alpha,p}^s] = \infty, \quad \text{for every value of} \ \ \alpha. \tag{6.16}$$

We also have **Observation 4**: If $s \geq 0$ and $\alpha < 0$ then

$$E[(W_{\alpha,p} + c_1 W_{\alpha,p}^{\alpha})^s] < \infty \tag{6.17}$$

if and only if

$$\text{both} \quad E[W_{\alpha,p}^s] < \infty \quad \text{and} \quad E[W_{\alpha,p}^{\alpha s}] < \infty. \tag{6.18}$$

Namely ($s \geq 0$):

$$W_{\alpha,p}^s < (W_{\alpha,p} + c_1 W_{\alpha,p}^{\alpha})^s, \quad (c_1 W_{\alpha,p}^{\alpha})^s < (W_{\alpha,p} + c_1 W_{\alpha,p}^{\alpha})^s, \tag{6.19}$$

and (Minkowski's inequality)

$$\left(E[(W_{\alpha,p} + c_1 W_{\alpha,p}^{\alpha})^s]\right)^{\frac{1}{s}} \leq \left(E[(W_{\alpha,p})^s]\right)^{\frac{1}{s}} + c_1 \left(E[(W_{\alpha,p})^{\alpha s}]\right)^{\frac{1}{s}}. \tag{6.20}$$

7 Existence of moments of $W_{\alpha,p}$: the Proofs

In this section we still have $\beta = 1, C = 0$.

For $0 \leq \alpha < 1$ we saw that $L_{\alpha,p}$ is stochastically smaller than $W_{\alpha,p}$, and from [4] we saw that (still $0 \leq \alpha < 1$) $V_{\alpha,p} = \frac{p}{(1-\alpha)c_1}(W_{\alpha,p})^{1-\alpha}$ is stochastically less than X_p, which is stochastically less than Z, for which $E[Z^s]$ is finite for all s. For $s > -\left|\frac{\log{(p)}}{\log{(b)}}\right|$ this gives finite lower and upper bounds for $E[W_{\alpha,p}^s]$. For $s \leq -\left|\frac{\log(p)}{\log(b)}\right|$ Observation 3 already told us that $E[W_{\alpha,p}^s] = \infty$. This completes the proof of (6.2).

Clearly, $X_{p,n}$ is the same as $V_{0,p,n} = \frac{p}{c_1} W_{0,p,n}$ with b replaced by $c = b^{1-\alpha}$. Hence, by the result above, $E[X_p^s] < \infty$ if and only if $s > -\left|\frac{\log(p)}{(1-\alpha)\log(b)}\right|$.

If $\alpha < 0$ we know from [4] that $V_{\alpha,p}$ is stochastically larger than X_p. With Observation 3 this shows that indeed $E[W_{\alpha,p}^s]$ is infinite for $s \leq -\left|\frac{\log(p)}{\log(b)}\right|$, finite for $-\left|\frac{\log(p)}{\log(b)}\right| < s \leq 0$.

In addition (still $\alpha < 0$), if $s > 0$ then

$$E[(W_{\alpha,p} + c_1 W_{\alpha,p}^{\alpha})^s] = c_1^s E[W_{\alpha,p}^{s\alpha}(1 + \frac{W_{\alpha,p}^{1-\alpha}}{c_1})^s] > c_1^s E[W_{\alpha,p}^{s\alpha}]. \tag{7.1}$$

Hence, by Observation 4, indeed $E[W_{\alpha,p}^s] = \infty$ when $s \geq \left|\frac{\log(p)}{\alpha\log(b)}\right|$.

Left over is the case $\alpha < 0,\ \ 0 < s < \left|\frac{\log(p)}{\alpha\log(b)}\right|$.

(6.5) and (6.6) show that

$$\Lambda_{\alpha,p,n+1} - X_{p,n+1} = \begin{cases} (\Lambda_{\alpha,p,n} - X_{p,n}) + R_{\alpha,p}(X_{p,n}) & \text{if } \chi_{p,n} = \text{success}, \\ c(\Lambda_{\alpha,p,n} - X_{p,n}) & \text{if } \chi_{p,n} = \text{failure}. \end{cases} \tag{7.2}$$

Hence,

$$E[(\Lambda_{\alpha,p,n+1} - X_{p,n+1})^{\nu}] = (1-p)E[((\Lambda_{\alpha,p,n} - X_{p,n}) + R_{\alpha,p}(X_{p,n}))^{\nu}] + pc^{\nu}E[(\Lambda_{\alpha,p,n} - X_{p,n})^{\nu}]. \tag{7.3}$$

In the remainder of this section we will use this to obtain crude, but finite, upper bounds for $E[(\Lambda_{\alpha,p,n} - X_{p,n})^{\nu}]$ in the case $\alpha < 0,\ \nu > 0$, as long as $E[(R_{\alpha,p}(X_p))^{\nu}] < \infty$.

Observation 5: Let G and H be any two random variables, and let $\nu \geq 0$. Let $\mu > 0$. Then

$$E[|G+H|^{\nu}] = E[|G+H|^{\nu}\chi(|H| \leq \mu|G|)] + E[|G+H|^{\nu}\chi(|H| > \mu|G|)] \leq$$
$$(1+\mu)^{\nu}E[|G|^{\nu}\chi(|H| \leq \mu|G|)] + \left(1 + \frac{1}{\mu}\right)^{\nu} E[|H|^{\nu}\chi(|H| > \mu|G|)] \leq$$

$$(1+\mu)^\nu E[|G|^\nu] + \left(\frac{1+\mu}{\mu}\right)^\nu E[|H|^\nu]. \tag{7.4}$$

We apply this result to (7.3) with $\nu \geq 0$ and obtain

$$E[(\Lambda_{\alpha,p,n+1} - X_{p,n+1})^\nu] \leq$$

$$((1-p)(1+\mu)^\nu + pc^\nu)\, E[(\Lambda_{\alpha,p,n} - X_{p,n})^\nu] + (1-p)\left(\frac{1+\mu}{\mu}\right)^\nu E[(R_{\alpha,p}(X_{p,n})^\nu]. \tag{7.5}$$

We now choose $X_{p,n}$ to have the stationary distribution, so that $E[(R_{\alpha,p}(X_{p,n}))^\nu] = E[(R_{\alpha,p}(X_p))^\nu]$ for all n. We choose $\nu > 0$ and choose $\mu > 0$ such, that $(1-p)(1+\mu)^\nu + pc^\nu < 1$ (that is clearly possible). If now for some n

$$E[(\Lambda_{\alpha,p,n} - X_{p,n})^\nu] \leq \frac{(1-p)\left(\frac{1+\mu}{\mu}\right)^\nu}{1-((1-p)(1+\mu)^\nu + pc^\nu)} E[(R_{\alpha,p}(X_p))^\nu] \tag{7.6}$$

then (7.6) holds for all $m \geq n$. Hence, if $E[(R_{\alpha,p}(X_p))^\nu] < \infty$ and $E[(\Lambda_{\alpha,p,0} - X_{p,0})^\nu] < \infty$ then

$$\limsup_{n\to\infty} E[(\Lambda_{\alpha,p,n} - X_{p,n})^\nu] \leq \frac{(1-p)\left(\frac{1+\mu}{\mu}\right)^\nu}{1-((1-p)(1+\mu)^\nu + pc^\nu)} E[(R_{\alpha,p}(X_p))^\nu] < \infty. \tag{7.7}$$

It is easily seen, and proven in [4], that (still $\alpha < 0$) $R_{\alpha,p}(v) = O(v^{-|\alpha|})$ for $v \downarrow 0$, and $R_{\alpha,p}(v) \to 0$ for $v \to \infty$. Hence $E[(R_{\alpha,p}(X_p))^\nu] < \infty$, and therefore also $E[(\Lambda_{\alpha,p} - X_p)^\nu] < \infty$, as long as $0 \leq \nu < \left|\frac{\log(p)}{|\alpha|(1+|\alpha|)\log(b)}\right|$. The remainder of the proof of (6.1) is left as an exercise for the reader.

As by–product of the proof above we see that $E[(\Lambda_{\alpha,p})^\nu] < \infty$ as long as $0 \leq \nu < \left|\frac{\log(p)}{|\alpha|(1+|\alpha|)\log(b)}\right|$. This shows that indeed also the process $(\Lambda_{\alpha,p,n})$ has a stationary distribution.

This section gives very coarse upper bounds to moments of interest, little more than proofs of finiteness. Once finiteness has been established the arguments in [4, 5] can be used to obtain sharper bounds.

It must be noted that $E[W^s_{\alpha,p}]$ can be infinite for certain values of s only because the process $(W_{\alpha,p,n})$ can, sufficiently often, achieve very small values. If $\alpha < 0$ then after the next upward jump $W_{\alpha,p,n}$ has a very large value. It is to be expected that with any $C > 0$ the process $(W^{**}_{\alpha,p,C,n})$ in (2.4) has a stationary distribution with $E[(W^{**}_{\alpha,p,C,n})^s] < \infty$ for all $-\infty < s < +\infty$.

8 Remaining Issues

Possibly the most intriguing remaining question is whether if $\beta = 1, 0 < c_2 < 1, \alpha < 0, C = 0$ the stationary distributions of the processes $W_{\alpha,p,n}$ and its rescaled version $V_{\alpha,p,n}$ are unique. In all other cases $(C > 0)$ and $(\beta = 1, 0 < c_2 < 1, 0 \leq \alpha < 1, C = 0)$ the uniqueness has been proven. It "seems obvious" that there is uniqueness also in the remaining case, but despite considerable effort, this has not yet been proven. If there is non–uniqueness the numbers α, c_1, c_2, p must have truly weird number–theoretic properties to allow multiple stationary distributions on sets that can not be reached from one another.

[7] proves in all cases with $C > 0$ weak convergence, for $p \downarrow 0$, of appropriately rescaled window size processes to a limit process. If also $(0 \leq \alpha < \beta = 1)$ this still holds if $C = 0$, as can be proven using the material in [5]. If $\alpha < 0 < \beta = 1, C = 0$ this additional step runs into difficulties. In that case there is an obvious "conjectured" limit process. In all cases weak convergence of the stationary distributions of the rescaled processes to the stationary distribution of the (proven or conjectured) limit process has been proven.

In the case $\beta = 1$ convergence of the moments of the stationary distributions of the rescaled processes to the moments of the stationary distribution of the (proven or conjectured) limit process has been proven, with rates of convergence. This problem has not been studied yet for the case $\beta < 1$.

References

[1] Ott, T.J., Kemperman, J.H.B., and Mathis, M. (1996)
The stationary behavior of ideal TCP congestion avoidance.
Available at www.teunisott.com/Papers.

[2] Mathis, M., Semke, J., Mahdavi, J. and Ott, T.J. (1997)
The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm.
Computer Communications Review **27**(3), pp 67 - 82 (July 1997).

[3] Ott, T.J. (2005)
Transport Protocols in the TCP Paradigm and their Performance.
Telecommunication Systems **30:**4 pp 351-385, Springer 2005.
Also available at www.teunisott.com/Papers.

[4] Ott, T.J. (2006)
Rate of convergence for the "Square Root Formula" in the Internet Transmission Control Protocol.
Adv. Appl. Prob. **38:**4 pp 1132 - 1154 (Dec 2006).
Also available at www.teunisott.com/Papers.

[5] Ott, T.J. (2006)
Rate of convergence for the "Square Root Formula" in TCP, extended version.
(Extended version of the paper above). www.teunisott.com/Papers.

[6] Ott, T.J. and Swanson, J. (2006)
Stationarity of some processes in transport protocols.
ACM Sigmetrics Performance Evaluation Review **34:**3 pp 30 - 32 (Dec 2006).
Based on material presented at the MAMA 2006 Workshop, St Malo, France, June 2006.

[7] Ott, T.J. and Swanson, J. (2007)
Asymptotic Behavior of a Generalized TCP Congestion Avoidance Algorithm.
J Appl Prob. **44,** pp 618 - 635 (Sept 2007).
Also available at www.teunisott.com/Papers.

[8] Ott, T.J. and Kemperman, J.H.B.K. (2008)
The Transient Behavior of processes in the TCP Paradigm.
Probability in the Engineering and Informational Sciences **22**(3) pp 431 - 471(July 2008).
Also available at www.teunisott.com/Papers.

[9] Padhye, J., Firoiu, V., Towsley, D., and Kurose, J. (2000)
Modeling of TCP RENO Performance: A simple model and its empirical validation.
IEEE/ACM Transactions on Networking **8**(2), pp 133–145 (2000).

[10] Budhiraja, A., Hernandez-Campos, F., Kulkarni, V.G. and Smith, F.D. (2004)
Stochastic Differential Equationfor TCP Window Size: Analysis and experimental validation.
Probability in the Engineering and Informational Sciences **18** pp 111-140 (2004).

[11] Kurtz, T.G. and Protter, P. (1991)
Weak limit theorems for stochastic integrals and stochastic differential equations.
Ann. Probab. **19**(3) pp 1035–1070 (1991).

[12] Meyn, S.P. and Tweedie, R.L. (1993)
Markov Chains and Stochastic Stability
Springer, 1993.

[13] Kelly, C.T. (2003)

Scalable TCP: Improving performance in highspeed wide area networks.

ACM SigComm Computer Communications Review **32**(2) pp 83-91 (2003).

[14] Kelly, C.T. (2004)

Engineering Flow Controls in the Internet.

PhD thesis, Cambridge Univ, England. Available at

http://www-lce.eng.cam.ac.uk/∼ctk21/papers.