

Stationarity of some processes in Transport Protocols

Teunis J. Ott *

Jason Swanson †

ABSTRACT

This note establishes stationarity of a number of stochastic processes of interest in the study of Transport Protocols. 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. After the 2006 MAMA workshop simple conditions for existence of such moments were obtained and were added to this note.

1. INTRODUCTION

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. That class of stochastic processes 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

$$\chi_{p,n} = \begin{cases} \text{success} & \text{if } U_n \geq p \\ \text{failure} & \text{if } U_n < p \end{cases} \quad (1.1)$$

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

$$W_{\alpha,p,C,n+1}^* = \begin{cases} W_{\alpha,p,C,n}^* + c_1(W_{\alpha,p,C,n}^*)^\alpha & \text{if } \chi_{p,n} = \text{success}, \\ \max(W_{\alpha,p,C,n}^* - c_2(W_{\alpha,p,C,n}^*)^\beta, C) & \text{if } \chi_{p,n} = \text{failure}, \end{cases} \quad (1.2)$$

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 [10, 11].

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

In [6] it is proven 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 .

*Ott Associates, Chester, NJ.

†Univ Wisconsin, Madison, Wisc.

The paper [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 “ $\max(\dots, C)$ ” in (1.2) (or choose $C = 0$). [4] also draws some conclusions, from the case “ $C = 0$ ”, for the case “ $C > 0$ ”.

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

$$W_{\alpha,p,n+1} = \begin{cases} W_{\alpha,p,n} + c_1(W_{\alpha,p,n})^\alpha & \text{if } \chi_{p,n} = \text{success}, \\ bW_{\alpha,p,n} & \text{if } \chi_{p,n} = \text{failure}. \end{cases} \quad (1.3)$$

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

$$W_{\alpha,p,C,n+1}^* = \begin{cases} W_{\alpha,p,C,n}^* + c_1(W_{\alpha,p,C,n}^*)^\alpha & \text{if } \chi_{p,n} = \text{success}, \\ \max(bW_{\alpha,p,C,n}^*, C) & \text{if } \chi_{p,n} = \text{failure}. \end{cases} \quad (1.4)$$

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.

As mentioned before, the paper [6] proves existence and uniqueness of the stationary distribution of the process $(W_{p,n}^*)_{n=0}^{\infty}$ in the case $C > 0$. In addition, the paper [4] proves existence and uniqueness of stationary distributions in case $(0 \leq \alpha < \beta = 1, C = 0, 0 < c_2 < 1)$. 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 $(\alpha < 0, \beta = 1, C = 0, 0 < c_2 < 1 \text{ and } 0 < p < 1 \text{ sufficiently small})$.

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

The papers [3], [4], [6] study rescaled versions of the processes described above. The paper [6] proves weak convergence, for $p \downarrow 0$, of these rescaled processes to interesting limit processes, and in many cases proves weak convergence of the stationary distributions of the rescaled processes to the stationary distributions of the limit processes. In the situation with $\beta = 1$ [4] even gives rate of convergence results and stochastic dominance results for that weak convergence of stationary distributions. If $\beta = 1$, it seems easiest to first study the case $C = 0$ and then apply the results to the case $C > 0$.

This note fills the the gap in the analysis of stationarity: By an alternative method it proves stationarity in all cases. However, the greater generality comes at a cost: for the time being, the results in this note do not give tightness of the family of stationary distributions for the rescaled processes for $p \downarrow 0$.

The results in this note are formulated for the original processes

$W_{p,C,n}^*$ and $W_{p,n}$, not for the rescaled processes.

The method of proof in this paper is to find a compact set $[v_1, v_2]$ and to prove that the “expected first return time” from W leaving that set to returning to that set is bounded. One of the results used is Theorem 12.3.4 on page 296 of [9].

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.

2. 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) $V : [C, \infty) \rightarrow [0, \infty)$ with the following properties:

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

$$V(w) \geq 1 + (1-p)V(w + c_1 w^\alpha) + pV(w - c_2 w^\beta) \text{ for } w > v_2, \quad (2.1)$$

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

$$V(w + c_1 w^\alpha) < B \text{ for all } C \leq w \leq v_2. \quad (2.2)$$

(2.1) shows that $V(w)$ is an upper bound for the Expected First Passage Time from $W = w > v_2$ to $W \leq v_2$. This is Theorem 11.3.4 on page 265 of [9]. (2.2) then shows that the process $(W_{p,C,n}^*)_{n=0}^\infty$ has (at least one) stationary distribution. That is Theorem 12.3.4 on page 296 of [9].

Once we have the results above, it is obvious that the expected first passage time from “anywhere” to $W = C$ 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

$$V(w) = \nu + \mu w^{1-\beta} \text{ for } w > v_2. \quad (2.3)$$

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

$$\mu w^{1-\beta} \geq 1 + (1-p)\mu(w + c_1 w^\alpha)^{1-\beta} + p\mu(1 - c_2 w^\beta)^{1-\beta} \quad (2.4)$$

for $w \geq w_2$ and

$$\nu \geq \frac{1}{p} \sup_{v_2 < w < w_2} \left(1 + (1-p)\mu(w + c_1 w^\alpha)^{1-\beta} - \mu w^{1-\beta} \right). \quad (2.5)$$

For (2.4) to hold we choose

$$\mu > \frac{1}{p(1-\beta)c_2}. \quad (2.6)$$

By a simple Binomial expansion we see that for (2.4) to hold, w_2 must be chosen at least equal to, or larger than, w_{min} , where (roughly)

$$w_{min} \sim \left(\frac{(1-p)(1-\beta)c_1}{p(1-\beta)c_2 - \frac{1}{\mu}} \right)^{\frac{1}{\beta-\alpha}}. \quad (2.7)$$

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 (2.2) holds.

3. 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

$$V(w) = \nu + \mu \log w \text{ for } w > v_2. \quad (3.1)$$

For (2.1) to hold we now need that

$$\mu \log(w) \geq 1 + (1-p)\mu \log(w + c_1 w^\alpha) + p\mu \log(bw) \quad (3.2)$$

for $w \geq w_2$ and that

$$\nu \geq \frac{1}{p} \sup_{v_2 < w < w_2} (1 + (1-p)\mu \log(w + c_1 w^\alpha) - \mu \log(w)). \quad (3.3)$$

For (3.2) to hold we choose

$$\mu > \frac{1}{p|\log(b)|} \quad (3.4)$$

and a simple expansion shows that for (3.2) to hold we must choose $w_2 > w_{min}$, where (roughly)

$$w_{min} \sim \left(\frac{(1-p)c_1}{p|\log(b)| - \frac{1}{\mu}} \right)^{\frac{1}{1-\alpha}}. \quad (3.5)$$

4. 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 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 3, drift down to the set $[v_1, v_2]$.

We choose the function $V(\cdot)$ as

$$V(w) = \begin{cases} \nu_l + \mu_l |\log(w)| & \text{for } 0 < w < v_1, \\ 0 & \text{for } v_1 \leq w \leq v_2, \\ \nu_u + \mu_u \log(w) & \text{for } v_2 < w < \infty. \end{cases} \quad (4.1)$$

(l and u stand for “lower” and “upper”). ν_u and μ_u are chosen as in Section 3, and an additional lower bound for v_2 is obtained as in that same section. The critical inequality now becomes

$$\nu_l + \mu_l |\log(w)| \geq$$

$$1 + (1-p)(\nu_u + \mu_u \log(w + c_1 w^\alpha)) + p(\nu_l + \mu_l |\log(bw)|) \quad (4.2)$$

for all $0 < w < w_1$. This can be re-written as

$$(1-p)(\nu_l - \nu_u) - p\mu_l |\log(b)| + (1-p)(\mu_l - \mu_u|\alpha|) |\log(w)| \geq 1 + (1-p)\mu_u \left(\log(c_1) + \log\left(1 + \frac{w^{1+|\alpha|}}{c_1}\right) \right) \quad (4.3)$$

for all $0 < w < w_1$. 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 first four sections of this note works for all $0 < p < 1$ but does not prove existence of moments.

After the 2006 MAMA workshop it was found exactly for what values of ν, α and p (etc) $E[W_{\alpha,p}^\nu]$ is finite for the stationary distribution of $W_{\alpha,p,n}$. The result is stated in the next section. The proof can be found in [7].

The cases $\alpha = 0$ and $0 < \alpha < 1$ in the situation of this section are easy to handle. In that situation the process W 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 $V(\cdot)$ is split into two independent sub-problems: 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

$$V(w) = \nu_l + \mu_l \log(|\log(w)|) \quad (4.4)$$

satisfies for w close to zero.

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

In the case $0 \leq \alpha < 1$ the next two sections of this note will show that $E[W_{\alpha,p}^\nu]$ is infinite for $\nu \leq -\left\lfloor \frac{\log(p)}{\log(b)} \right\rfloor$, finite for all larger values of ν .

The possible non-uniqueness of the stationary distributions (in the situation $\alpha < 0, \beta = 1, C = 0$) is an intriguing question. Do there exist combinations of $\alpha, b = 1 - c_2, c_1$ and p for which the process $W_{p,n}$ has multiple stationary distributions (on sets not reachable from one another)? If such combinations exist, they must have weird number-theoretic properties.

5. EXISTENCE OF MOMENTS OF $W_{\alpha,p}$

For the stationary distribution of $W_{\alpha,p}$, if $\alpha < 0$ then (still $0 < c_2 < 1$ and $b = 1 - c_2$):

$$E[W_{\alpha,p}^\nu] \begin{cases} = \infty & \text{if } \nu \geq \left\lfloor \frac{\log(p)}{\alpha \log(b)} \right\rfloor, \\ < \infty & \text{if } -\left\lfloor \frac{\log(p)}{\log(b)} \right\rfloor < \nu < \left\lfloor \frac{\log(p)}{\alpha \log(b)} \right\rfloor, \\ = \infty & \text{if } \nu \leq -\left\lfloor \frac{\log(p)}{\log(b)} \right\rfloor, \end{cases} \quad (5.1)$$

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

$$E[W_{\alpha,p}^\nu] \begin{cases} < \infty & \text{if } \nu > -\left\lfloor \frac{\log(p)}{\log(b)} \right\rfloor, \\ = \infty & \text{if } \nu \leq -\left\lfloor \frac{\log(p)}{\log(b)} \right\rfloor. \end{cases} \quad (5.2)$$

The proof of this statement can be found in [7].

We see that if p is close enough to 1 then $W_{\alpha,p,n}$ gets close enough to zero often enough to make $E[W_{\alpha,p}^\nu]$ infinite if ν is close

enough to $-\infty$. As long as $\alpha \geq 0$ this has no impact on the existence of $E[W_{\alpha,p}^\nu]$ for $\nu > 0$. However, if $\alpha < 0$, $W_{\alpha,p,n}$ being close to zero often implies that also $W_{\alpha,p,n}$ is close to $+\infty$ often enough to cause $E[W_{\alpha,p}^\nu]$ to be infinite for ν close enough to $+\infty$.

This observation, while mathematically of some interest, almost certainly has no consequences for protocol analysis: In real Transport Protocols the congestion window (apart from during time-outs) is bounded away from zero.

6. REFERENCES

- [1] Ott, T.J., Kemperman, J.H.B. and Mathis, M. (2004) The stationary behavior of ideal TCP congestion avoidance, revisited. Draft.
- [2] 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.
- [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 TCP. To appear in *Adv. Appl. Prob.* 38.4 (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) Asymptotic Behavior of a Generalized TCP Congestion Avoidance Algorithm. (Submitted) Available at www.teunisott.com/Papers.
- [7] Ott, T.J. and Swanson, J. (2006) Stationarity of some processes in Transport Protocols, extended version. (An extended version of this paper). Available at www.teunisott.com/Papers.
- [8] 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).
- [9] Meyn, S.P. and Tweedie, R.L. (1993) *Markov Chains and Stochastic Stability* Springer, 1993.
- [10] Kelly, C.T. (2003) Scalable TCP: Improving performance in highspeed wide area networks. *ACM SigComm Computer Communications Review* 32(2) pp 83-91 (2003).
- [11] Kelly, C.T. (2004) Engineering Flow Controls in the Internet. PhD thesis, Cambridge Univ, England. Available at <http://www-ice.eng.cam.ac.uk/~ctk21/papers>.