ASYMPTOTIC BEHAVIOR OF A GENERALIZED TCP CONGESTION AVOIDANCE ALGORITHM

By Teunis J. Ott

Winlab, Rutgers University

and

By Jason Swanson*

University of Wisconsin-Madison

August 18, 2006

The Transmission Control Protocol (TCP) is a Transport Protocol used in the Internet. In [34], a more general class of candidate Transport Protocols called "protocols in the TCP Paradigm" is introduced. The long run objective of studying this class is to find protocols with promising performance characteristics. This paper studies Markov chain models derived from protocols in the TCP Paradigm.

Protocols in the TCP Paradigm, as TCP, protect the network from congestion by decreasing the "Congestion Window" (the amount of data allowed to be sent but not yet acknowledged) when there is packet loss or packet marking, and increasing it when there is no loss. When loss of different packets are assumed to be independent events and the probability p of loss is assumed to be constant, the protocol gives rise to a Markov chain $\{W_n\}$, where W_n is the size of the congestion window after the transmission of the n-th packet.

For a wide class of such Markov chains, we prove weak convergence results, after appropriate rescaling of time and space, as $p \to 0$. The limiting processes are defined by stochastic differential equations. Depending on certain parameter values, the stochastic differential equation can define an Ornstein-Uhlenbeck process or can be driven by a Poisson process.

1. Introduction. The Congestion Avoidance algorithm of TCP is designed to prevent network congestion during the transmission of data over

*This work was supported in part by the VIGRE grant of University of Wisconsin-Madison.

AMS 2000 subject classifications: Primary 60F05; secondary 60G10, 60H10, 60J05

Keywords and phrases: Weak Convergence, Stochastic Differential Equations, Stationary Distributions, TCP/IP, Congestion Avoidance

a computer network. It does this by controlling the congestion window, i.e. the amount of data "transmitted but not yet acknowledged" by a sender. What follows is a simplified description of a more general class of Transport Protocols.

Under appropriate units, the congestion window W determines the maximum amount of data that a source can send without acknowledgement. The "TCP Paradigm" (see [34]) is a class of protocols that includes TCP (and other Transport Protocols). For each protocol in the TCP Paradigm there are two functions, $incr(\cdot)$ and $decr(\cdot)$. If, while the congestion window equals W, a packet is found to be lost (or marked, under ECN – see [19] and [44]), then the congestion window is reduced by $decr(W)$. However, the congestion window is never reduced below some fixed minimum value $\ell \geq 0$. If the packet is not lost, then the congestion window is increased by $incr(W)$. For protocols in the TCP Paradigm, $incr(W) = c_1 W^\alpha$ and $decr(W) = c_2 W^\beta$. In the special case of TCP, we have $c_1 = 1$, $\alpha = -1$, $c_2 = 1/2$, and $\beta = 1$. Another special case of interest is when $\alpha = 0$ and $\beta = 1$. This is the algorithm which Tom Kelly calls "Scalable TCP" in [22] and [23].

Let W_n denote the size of the congestion window after the transmission of the n-th packet, or, more accurately, after receipt of the n-th "good" acknowledgement. Let χ_n be the indicator function of the event that the n-th packet is lost, or, more accurately, that there is a loss between the $(n-1)$-th and n-th "good" acknowledgement. We shall assume that the χ_n's are independent and identically distributed. In particular, we are assuming that $p = P(\chi_n = 1)$ is a constant that does not change with time. Under these assumptions, we are led to the parameterized family of Markov processes

$$W_{p,n+1} = (W_{p,n} + c_1 W_{p,n}^\alpha (1 - \chi_{p,n+1}) - c_2 W_{p,n}^\beta \chi_{p,n+1}) \vee \ell. \tag{1.1}$$

The assumptions we place on the various parameters in the model are:

$$\{\chi_{p,n}\}_{n=1}^\infty \text{ is an iid sequence of } \{0,1\}\text{-valued random variables,} \tag{1.2}$$

$$p = P(\chi_{p,n} = 1), \tag{1.3}$$

$$c_1 > 0 \text{ and } c_2 > 0, \tag{1.4}$$

$$-\infty < \alpha < \beta \leq 1 \text{ and } \ell \geq 0, \tag{1.5}$$

$$\text{if } \beta = 1, \text{ then } c_2 < 1, \text{ and} \tag{1.6}$$

$$\text{if } \beta < 1, \text{ then } \ell > 0. \tag{1.7}$$

We will frequently drop the dependence on p from our notation and simply refer to the processes $\{\chi_n\}$ and $\{W_n\}$.

We are interested in studying the asymptotic behavior of $\{W_n\}$ as $p \to 0$. To this end, we define the continuous time process

$$Z_p(t) = p^\gamma W_{\lfloor tp^{-\nu} \rfloor}, \text{ where } \gamma = (\beta - \alpha)^{-1} \text{ and } \nu = (1-\alpha)\gamma. \tag{1.8}$$

In the case that $\beta = 1$, we will show that Z_p converges weakly as $p \to 0$ to the process Z defined by

$$Z(t) = Z(0) + c_1 \int_0^t Z(s)^\alpha \, ds - c_2 \int_0^t Z(s-) \, dN(s), \tag{1.9}$$

where N is a unit rate Poisson process, independent of $Z(0) = \lim Z_p(0)$. (Note that this is the conjecture given on page 362 of [34].) We will also show that, when $\ell > 0$, the stationary distributions of the discrete time Markov chains $\{p^\gamma W_n\}$ converge weakly to the unique stationary distribution of Z. Questions about the convergence of the stationary distributions when $\beta = 1$, as well as the rate of convergence, are addressed in [36] and [39] using techniques that differ from those used in this paper.

In the case that $\beta < 1$, we will show that Z_p converges to the process ζ defined by

$$\zeta(t) = \zeta(0) + \int_0^t (c_1 \zeta(s)^\alpha - c_2 \zeta(s)^\beta) \, ds, \tag{1.10}$$

where $\zeta(0) = \lim Z_p(0)$. With the exception of the initial condition, the process ζ is entirely deterministic. The convergence of Z_p to ζ is therefore a law of large numbers type of result. Hence, in the case $\beta < 1$, we can extend our analysis and study the fluctuations of Z_p around this central tendency. Unfortunately, it will not suffice to center Z_p by ζ. We must rather define

$$\zeta_p(t) = \zeta_p(0) + \int_0^t (c_1(1-p)\zeta_p(s)^\alpha - c_2 \zeta_p(s)^\beta) \, ds, \tag{1.11}$$

where $\zeta_p(0) \to \zeta(0)$, and consider the processes

$$\xi_p(t) = p^{-\tau}(Z_p(t) - \zeta_p(t)), \text{ where } \tau = (\nu - 1)/2. \tag{1.12}$$

We will show that ξ_p converges weakly as $p \to 0$ to the process ξ defined by

$$\begin{aligned} \xi(t) = \xi(0) &+ \int_0^t (c_1 \alpha \zeta(s)^{\alpha-1} - c_2 \beta \zeta(s)^{\beta-1}) \xi(s) \, ds \\ &- c_2 \int_0^t \zeta(s)^\beta \, dB(s), \end{aligned} \tag{1.13}$$

where B is a Brownian motion and $\xi(0) = \lim \xi_p(0)$.

A special case of this last result is worth mentioning. For each $p \in [0,1)$, define

$$c_p = (c_1(1-p)/c_2)^{\gamma}, \tag{1.14}$$

so that $\zeta_p(t) = c_p$ is an invariant solution to (1.11). Also, $\zeta(0) = \lim \zeta_p(0) = c_0$ is an invariant solution to (1.10). Hence, for an appropriate choice of $Z_p(0)$, ξ_p converges to the Ornstein-Uhlenbeck process defined by

$$d\xi = -\mu\xi dt + \sigma dW, \tag{1.15}$$

where $W = -B$,

$$\begin{aligned} \mu &= c_2\beta(c_1/c_2)^{\gamma(\beta-1)} - c_1\alpha(c_1/c_2)^{\gamma(\alpha-1)} \\ &= (\beta-\alpha)c_1^{-(1-\beta)/(\beta-\alpha)}c_2^{(1-\alpha)/(\beta-\alpha)}, \end{aligned}$$

and

$$\sigma = c_2(c_1/c_2)^{\gamma\beta} = c_1^{\beta/(\beta-\alpha)}c_2^{-\alpha/(\beta-\alpha)}.$$

(Note that this is the conjecture given on page 364 of [34].) We will also show that the stationary distributions of the discrete time Markov chains $\{p^{-\tau}(p^{\gamma}W_n - c_p)\}$ converge weakly to the unique stationary distribution of the above Ornstein-Uhlenbeck process.

It should be remarked that in this paper we use so-called "packet time." That is, the progress of time is expressed in number of data packets sent, or, more accurately, number of good acknowledgements received. Several papers analyzing TCP use "clock time," where the progress of time is expressed in number of Round Trip Times (RTTs) elapsed. If the congestion window is the only limit on the "flight size" (amount of data transmitted by the source for which no acknowledgement has yet been received), all packets contain one Maximum Segment Size (MSS) of data, and the congestion window is expressed in MSSs, then clock time t_C and packet time t_P are related by $dt_P = W\,dt_C$, where W denotes the size of the congestion window. Stationary distributions for "packet time" and "clock time" are related but are not the same. The relationship is given in [38].

2. Related work. When results like those in this paper are applied to "classical TCP", which has $\alpha = -1, \beta = 1$, they predict throughput for a (large) TCP flow in the order of $1/\sqrt{p}$. This is called the "Square Root Law" for TCP, and original papers in this area were often identified with the Square Root Law for TCP.

Work in this area started with [38], which among other things gives the stationary distribution of the limit process in the case $\beta = 1$, and the relationship between "packet time stationary distributions" and "clock time stationary distributions". The paper gives the stationary distribution of the limit process $W_{p,n}$ for $p \downarrow 0$ and assumes the weak convergence results which strictly speaking are not proven until [36] and this paper. That paper was presented at a workshop of the IFIP WG7.3 during Performance 1996 in Lausanne, Oct 1996 and also in a DIMACS workshop at Rutgers University in Nov 1996.

Another paper of historical interest is [33] which was presented in a workshop at ENS, Paris 2000. That paper first explicitly formulated the conjectures proven in this paper. In re–written form it appeared as [34].

In a non-distributional sense some of the results had been anticipated in [25].

The first papers identified with the "square root law" that made it into the open literature were [27], [41], [32], [42], [1] and [21].

Of these, [32] is the first to use the language of stochastic differential equations. It uses clock time, and assumes that the probability of a drop in a RTT is independent of the size of the congestion window, i.e. the drop–probability per packet is roughly inversely proportional to the size of the window.

An extensive bibliography and discussion of previous work is found in [16], which, among other things, studies the effect of a congestion window limited by a send window or receive window (through the advertised window).

The first papers to use "clock time" are [32] and [1]. Other papers that use clock time are [17] and [20].

Another paper of particular interest is [15], which uses stochastic differential equations, in clock time, to study joint evolution of RTT and congestion window size. The parameters of the two–dimensional stochastic differential equation are obtained from measurements in the Internet, not from postulating a particular behavior of sources and routers.

Other papers we want to mention are [4], where (as in [15]) the RTT depends on the flightsize, [2], which is an ambitious attempt to build an all–encompassing model where many flows keep each others' RTTs and drop probabilities in equilibrium, [3], which analyzes the performance of Scalable TCP ($\alpha = 0, \beta = 1$), [26], [7], [8], [6], [5], [10] and [9].

The papers [40], [28], [11], [13], [12], [29], [30], [14], [31] use "Square Root Law Results" and do analysis with, for example, drop probabilities that depend on the current size of the congestion window. That dependence is by assuming ECN and is evaluated by translating flightsize in queuelength in the router.

The conjectures proven in this paper are formulated in [34], which also obtains a number of other results, linked to "practicality" of control schemes, such as relaxation times, typical numbers of dropped or marked packets per RTT, etc.

An alternative proof of the stationarity of the processes $(W_{p,n})_{n=0}^{\infty}$ studied in this paper is given in [39].

For a more complete review of the literature, see [16].

Of possible future interest is [35] which makes a start with investigating the impact of delay of one RTT in the feedback on stability of the Internet, and [37] which studies the transient behavior of the limit process in the case $\beta = 1$, and thus, in so far the limit results apply, can be used to predict the amount of clock time it takes to ftp a very large file.

3. Main Results.

We first consider the case $\beta = 1$ and begin by cataloging some properties of the limit process Z.

LEMMA 3.1. *If $Z(0) > 0$ a.s., then the stochastic differential equation (1.9) has a unique solution Z. With probability one, $Z(t) > 0$ for all $t \geq 0$. Moreover, if $\tau = \inf\{t \geq 0 : Z(t) = c_0\}$, where c_0 is given by (1.14), then $\tau < \infty$ a.s.*

Proof. For each realization of the Poisson process, (1.9) can be solved deterministically and the solution is unique. Let

$$T = \inf\{t \geq 0 : Z(t) \notin (0, \infty)\}.$$

Since Z decreases only at the jump times of the Poisson process, and, with probability one, these jump times have no accumulation points, it follows that $T = \infty$ a.s.

To show that $\tau < \infty$ a.s., it will suffice to assume that $Z(0) = x > 0$ is deterministic. We first consider the case $x \leq c_0$. Suppose $\tau(\omega) = \infty$. Then $Z(t, \omega) < c_0$ for all $t \geq 0$. Find $u > r$ such that $u - r > \gamma c_2^{-1}$ and $N(u, \omega) = N(r, \omega)$. Then for all $t \in (r, u]$,

$$Z(t, \omega) = Z(r, \omega) + c_1 \int_r^t Z(s, \omega)^\alpha \, ds.$$

Since the solution to this integral equation is unique,

$$Z(t, \omega) = (c_1(1 - \alpha)(t - r) + Z(r, \omega)^{1-\alpha})^\gamma.$$

Therefore,

$$c_0 > Z(u, \omega) > (c_1(1 - \alpha)(u - r))^\gamma > c_0,$$

a contradiction. Hence, $\tau < \infty$ a.s.

We next consider the case $x > c_0$. Define

$$\sigma_1 = \inf\{t \geq 0 : Z(t) < c_0\} \text{ and } \sigma_2 = \inf\{t \geq \sigma_1 : Z(t) = c_0\},$$

so that $\tau \leq \sigma_2$, and it will suffice to show that $\sigma_2 < \infty$ a.s. Fix $L > x$ and define $\rho = \inf\{t \geq 0 : Z(t) \notin [c_0, L]\}$. Suppose $\rho(\omega) = \infty$. Then $Z(t, \omega) \in [c_0, L]$ for all $t \geq 0$. Let

$$K = \inf\{u^\alpha : c_0 \leq u \leq L\} > 0.$$

Find $u > r$ such that $u - r > (L - c_0)/(c_1 K)$ and $N(u, \omega) = N(r, \omega)$. Then

$$L \geq Z(u, \omega) = Z(r, \omega) + c_1 \int_r^u Z(s, \omega)^\alpha \, ds \geq c_0 + c_1(u - r)K > L,$$

a contradiction. Hence, $\rho < \infty$ a.s.

Now, observe that

$$Z(t \wedge \rho) = x + \int_0^{t \wedge \rho} (c_1 Z(s)^\alpha - c_2 Z(s)) \, ds - c_2 \int_0^{t \wedge \rho} Z(s-) \, dM(s),$$

where $M(t) = N(t) - t$ is the compensated Poisson process. If $s < t \wedge \rho$, then $Z(s) \geq c_0 = (c_1/c_2)^\gamma$. This implies that $c_1 Z(s)^\alpha - c_2 Z(s) \leq 0$. Since M is a martingale, $E[Z(t \wedge \rho)] \leq x$. Letting $t \to \infty$ gives $E[Z(\rho)] \leq x$. Hence, $P(Z(\rho) = L) \leq x/L$. Note that either $Z(\rho) = L$ or $Z(\rho) < c_0$. Therefore,

$$P(\sigma_1 = \infty) \leq P(Z(\rho) = L) \leq x/L.$$

Letting $L \to \infty$ shows $\sigma_1 < \infty$ a.s.

As in Theorem V.6.35 in [43], Z is a strong Markov process. Therefore,

$$P(\sigma_2 = \infty) = E[P^{Z(\sigma_1)}(\tau = \infty)].$$

But $Z(\sigma_1) < c_0$, and we have already shown that $P^x(\tau = \infty) = 0$ for all $x \leq c_0$. Hence, $\sigma_2 < \infty$ a.s. $\square$

We are now prepared to state our main results for the case $\beta = 1$. If μ_p and μ are Borel measures on a metric space S, then the notation $\mu_p \Rightarrow \mu$ will mean that μ_p converges weakly to μ as $p \to 0$, that is, $\int_S f \, d\mu_p \to \int_S f \, d\mu$ as $p \to 0$ for all bounded, continuous $f : S \to \mathbb{R}$. If X_p and X are S-valued random variables, then $X_p \Rightarrow X$ will mean that $PX_p^{-1} \Rightarrow PX^{-1}$. When X_p and X are processes, we will take our metric space to be $D_{\mathbb{R}^d}[0, \infty)$, the space of cadlag functions from $[0, \infty)$ to $\mathbb{R}^d$, with the Skorohod metric. See [18] for details.

THEOREM 3.2. *Suppose $\beta = 1$. Let the processes Z_p be given by (1.8) and suppose that $Z_p(0) \Rightarrow Z(0)$, where $Z(0) > 0$ a.s. Let Z be the unique solution to (1.9). Then $Z_p \Rightarrow Z$.*

THEOREM 3.3. *Suppose $\beta = 1$ and $\ell > 0$. Then the Markov chain $\{W_n\}$ has a unique stationary distribution. Moreover, the process Z given by (1.9) has a unique stationary distribution η on $(0, \infty)$. For each $p > 0$, let η_p be the stationary distribution for the Markov chain $\{p^\gamma W_n\}$. Then $\eta_p \Rightarrow \eta$.*

For some results on stationary distributions in the case $\beta = 1$ and $\ell = 0$, see [36] and [39].

For the case $\beta < 1$, we need some preliminary definitions. Assume that for all $p \in (0, 1)$, the processes $\{W_{p,n}\}$ are defined on the same probability space $(\Omega, \mathcal{F}, P)$. Define the σ-algebra

$$\mathcal{F}_0 = \sigma(W_{p,0} : 0 < p < 1) \vee \mathcal{N}, \tag{3.1}$$

where $\mathcal{N}$ denotes the collection of events $D \in \mathcal{F}$ with $P(D) = 0$.

THEOREM 3.4. *Suppose $\beta < 1$. Let the processes Z_p be given by (1.8). Suppose that $Z_p(0) \Rightarrow \zeta(0)$, where $\zeta(0) > 0$ a.s. Let ζ the unique solution to (1.10). Then $Z_p \Rightarrow \zeta$. Moreover, if $Z_p(0) \to \zeta(0)$ in probability, then $Z_p \to \zeta$ in probability.*

THEOREM 3.5. *Suppose $\beta < 1$. Let the processes Z_p be given by (1.8). For each $p \in (0, 1)$, let $\zeta_p(0)$ be a strictly positive random variable defined on $(\Omega, \mathcal{F}, P)$. Assume that $\zeta_p(0)$ is $\mathcal{F}_0$-measurable and $Z_p(0) - \zeta_p(0) \to 0$ in probability. Define ζ_p and ξ_p by (1.11) and (1.12), respectively.*

Suppose that there exists a pair of random variables $(\xi(0), \zeta(0))$, defined on $(\Omega, \mathcal{F}, P)$, such that $\zeta(0) > 0$ a.s. and $(\xi_p(0), \zeta_p(0)) \Rightarrow (\xi(0), \zeta(0))$. Let B be a standard Brownian motion independent of $(\xi(0), \zeta(0))$ and define the processes ζ and ξ by (1.10) and (1.13), respectively. Then $(\xi_p, \zeta_p) \Rightarrow (\xi, \zeta)$.

THEOREM 3.6. *Suppose $\beta < 1$. Then the Markov chain $\{W_n\}$ has a unique stationary distribution. For each $p > 0$, let η_p be the stationary distribution for the Markov chain $\{p^{-\tau}(p^\gamma W_n - c_p)\}$. Then $\eta_p \Rightarrow \eta$, where η is the stationary distribution of the Ornstein-Uhlenbeck process given by (1.15).*

4. General Definitions.

Define

$$\Lambda_n = (\ell - W_{n-1} - c_1 W_{n-1}^\alpha (1 - \chi_n) + c_2 W_{n-1}^\beta \chi_n) \vee 0,$$

so that

$$W_{n+1} = W_n + c_1 W_n^\alpha - (c_1 W_n^\alpha + c_2 W_n^\beta)\chi_{n+1} + \Lambda_{n+1}.$$

If we let $W(t) = W_{\lfloor t \rfloor}$, then we can rewrite this recursive relation as the integral equation

$$\begin{aligned} W(t) = W(0) &+ c_1 \int_0^t W(s-)^\alpha \, dm(s) \\ &- \int_0^t (c_1 W(s-)^\alpha + c_2 W(s-)^\beta) \, dS(s) + L(t), \end{aligned}$$

where

$$m(t) = \lfloor t \rfloor, \;\; S(t) = \sum_{j=1}^{\lfloor t \rfloor} \chi_j, \;\; \text{and } L(t) = \sum_{j=1}^{\lfloor t \rfloor} \Lambda_j.$$

Using (1.8), it is then easy to see that

$$\begin{aligned} Z_p(t) = Z_p(0) &+ c_1 \int_0^t Z_p(s-)^\alpha \, dm_p(s) \\ &- c_1 p \int_0^t Z_p(s-)^\alpha \, dS_p(s) - c_2 \int_0^t Z_p(s-)^\beta \, dS_p(s) + L_p(t), \end{aligned} \tag{4.1}$$

where

$$m_p(t) = p^\nu m(tp^{-\nu}), \;\; S_p(t) = p^{\nu-1} S(tp^{-\nu}), \;\; \text{and } L_p(t) = p^\gamma L(tp^{-\nu}).$$

Note that if we define the filtration

$$\mathcal{F}_t^p = \mathcal{F}_0 \vee \sigma(\chi_{p,j} : j \le \lfloor tp^{-\nu} \rfloor),$$

then m_p, S_p, and L_p are all $\{\mathcal{F}_t^p\}$-adapted.

Define the $\mathbb{R}^2$-valued cadlag $\{\mathcal{F}_t^p\}$-semimartingale

$$Y_p = (m_p, S_p)^T$$

and define the function $G_p : \mathbb{R}^2 \to \mathbb{R}$ by

$$G_p(x) = (c_1 x^\alpha, -c_1 p x^\alpha - c_2 x^\beta) 1_{\{x>0\}}.$$

Then (4.1) becomes

$$Z_p(t) = Z_p(0) + \int_0^t G_p(Z_p(s-)) \, dY_p(s) + L_p(t).$$

To show that Z_p converges as $p \to 0$, we will apply the theorems in [24]. This approach, however, comes with two technical difficulties. The first is the presence of the local time term L_p; the second is the fact that G_p may have a singularity at the origin. To deal with these issues, we introduce the process Z_p^ε, defined as the unique solution to

$$Z_p^\varepsilon(t) = Z_p(0) + \int_0^t G_p^\varepsilon(Z_p^\varepsilon(s-))\, dY_p(s), \tag{4.2}$$

where $G_p^\varepsilon = G_p(\varepsilon)1_{(-\infty,\varepsilon)} + G_p 1_{[\varepsilon,\infty)}$. To quantify the sense in which Z_p and Z_p^ε are close, we define the functional $h_\varepsilon : D_{\mathbb{R}^d}[0,\infty) \to [0,\infty]$ by

$$h_\varepsilon(x) = \inf\{t \geq 0 : |x(t)| \wedge |x(t-)| \leq \varepsilon\},$$

and the stopping times $\tau_p(\varepsilon) = h_\varepsilon(Z_p^\varepsilon)$, and we observe that

$$L_p = 0 \text{ and } Z_p = Z_p^\varepsilon \text{ on } [0, \tau_p(\varepsilon \vee p^\gamma \ell)). \tag{4.3}$$

By (3.5.2) in [18], if two cadlag functions x and y agree on the interval $[0,t)$, then $d(x,y) \leq e^{-t}$, where d is the metric on $D_{\mathbb{R}^d}[0,\infty)$.

5. Convergence of Z_p. In this section, we will prove Theorems 3.2 and 3.4 by applying the theorems in [24] to the processes Z_p^ε given by (4.2). We must therefore define the processes to which they converge in the cases $\beta = 1$ and $\beta < 1$.

Let $G(x) = (c_1 x^\alpha, -c_2 x^\beta)1_{\{x>0\}}$ and $G^\varepsilon = G(\varepsilon)1_{(-\infty,\varepsilon)} + G 1_{[\varepsilon,\infty)}$, and note that $G_p^\varepsilon \to G^\varepsilon$ uniformly on compacts as $p \to 0$. Let N be a unit rate Poisson process, define

$$Y(t) = (t, N(t))^T \text{ and } y(t) = (t,t)^T,$$

and let Z^ε and ζ^ε be the unique solutions to

$$Z^\varepsilon(t) = Z(0) + \int_0^t G^\varepsilon(Z^\varepsilon(s-))\, dY(s), \tag{5.1}$$

$$\zeta^\varepsilon(t) = \zeta(0) + \int_0^t G^\varepsilon(\zeta^\varepsilon(s-))\, dy(s), \tag{5.2}$$

where $Z(0)$ and N are independent. Note that if $\beta = 1$, then $Z^\varepsilon = Z$ on $[0, h_\varepsilon(Z^\varepsilon))$ and $h_\varepsilon(Z^\varepsilon) = h_\varepsilon(Z) \to \infty$ a.s. as $\varepsilon \to 0$. Hence, $d(Z^\varepsilon, Z) \leq \exp(-h_\varepsilon(Z)) \to 0$ a.s. That is, $Z^\varepsilon \to Z$ a.s. in $D_\mathbb{R}[0,\infty)$. Similarly, if $\beta < 1$, then $\zeta^\varepsilon = \zeta$ on $[0, h_\varepsilon(\zeta^\varepsilon))$, $h_\varepsilon(\zeta^\varepsilon) = h_\varepsilon(\zeta) \to \infty$ a.s., and $\zeta^\varepsilon \to \zeta$ a.s. in $D_\mathbb{R}[0,\infty)$.

We will show that $Z_p^\varepsilon \Rightarrow Z^\varepsilon$ and $\zeta_p^\varepsilon \Rightarrow \zeta^\varepsilon$. To pass from this to the conclusions of Theorems 3.2 and 3.4, we will need the following lemma, which is easily proved using the Prohorov metric. (See Section 3.1 in [18].)

LEMMA 5.1. *Let (S,d) be a complete and separable metric space. Let $\{X_p\}_{p>0}$ be a family of S-valued random variables and suppose, for each ε, there exists a family $\{X_p^\varepsilon\}_{p>0}$ such that*

$$\limsup_{p\to 0} E[d(X_p, X_p^\varepsilon)] \le \delta_\varepsilon,$$

where $\delta_\varepsilon \to 0$ as $\varepsilon \to 0$. Suppose also that for each ε, there exists Y^ε such that $X_p^\varepsilon \Rightarrow Y^\varepsilon$ as $p \to 0$. Then there exists X such that $X_p \Rightarrow X$ and $Y^\varepsilon \Rightarrow X$.

Proof of Theorem 3.2. Suppose $\beta = 1$, Z_p is given by (1.8), and $Z_p(0) \Rightarrow Z(0)$, where $Z(0) > 0$ a.s. Let Z be the solution to (1.9).

Let Z_p^ε and Z^ε be given by (4.2) and (5.1). We first show that $Z_p^\varepsilon \Rightarrow Z^\varepsilon$. Recall that $G_p^\varepsilon \to G^\varepsilon$ uniformly on compacts. Also observe that $S_p \Rightarrow N$ (see, for example, Problem 7.1 in [18]). Hence, since $Z_p(0)$ and Y_p are independent, $(Z_p(0), Y_p) \Rightarrow (Z(0), Y)$ in $D_{\mathbb{R}^3}[0,\infty)$. Hence, by Theorem 5.4 in [24], it will suffice to show that Y_p has a semimartingale decomposition $Y_p = M_p + A_p$ into a martingale part and a bounded variation part such that for each $t \ge 0$,

$$\sup_p E[[M_p]_t + T_t(A_p)] < \infty, \tag{5.3}$$

where $[M_p]_t$ is the quadratic variation process of M_p and $T_t(A_p)$ is the total variation of A_p on the interval $[0,t]$.

For this, define

$$\tilde{S}_p(t) = S_p(t) - m_p(t) = p^{\nu-1} \sum_{j=1}^{\lfloor tp^{-\nu}\rfloor} (\chi_j - p),$$

so that $\tilde{S}_p$ is an $\{\mathcal{F}_t^p\}$-martingale. Note that $T_t(m_p) = m_p(t)$ and

$$E[\tilde{S}_p]_t = p^{2\nu-2} \sum_{j=1}^{\lfloor tp^{-\nu}\rfloor} E|\chi_j - p|^2 = p^{2\nu-2}\lfloor tp^{-\nu}\rfloor p(1-p) \le tp^{\nu-1}. \tag{5.4}$$

Since $\beta = 1$ implies $\nu = 1$, this verifies (5.3) and shows that $Z_p^\varepsilon \Rightarrow Z^\varepsilon$.

By passing to a subsequence, we can assume there exists a $[0,\infty]$-valued random variable $\sigma(\varepsilon)$ such that $(Z_p^\varepsilon, h_\varepsilon(Z_p^\varepsilon)) \Rightarrow (Z^\varepsilon, \sigma(\varepsilon))$. By (4.3),

$$\begin{aligned}\limsup_{p\to 0} E[d(Z_p, Z_p^\varepsilon)] &\le \limsup_{p\to 0} E[\exp(-\tau_p(\varepsilon \vee p^\gamma \ell))]\\ &= \limsup_{p\to 0} E[\exp(-h_\varepsilon(Z_p^\varepsilon))]\\ &= E[\exp(-\sigma(\varepsilon))].\end{aligned}$$

We claim that $E[\exp(-\sigma(\varepsilon))] \le E[\exp(-h_\varepsilon(Z^\varepsilon))]$. To see this, let us assume by the Skorohod Representation Theorem (see Theorem 3.1.8 in [18]) that $(Z_p^\varepsilon, h_\varepsilon(Z_p^\varepsilon)) \to (Z^\varepsilon, \sigma(\varepsilon))$ a.s. Then $h_\varepsilon(Z^\varepsilon) \le \sigma(\varepsilon)$ a.s., which proves the claim.

Since $h_\varepsilon(Z^\varepsilon) = h_\varepsilon(Z) \to \infty$ a.s. as $\varepsilon \to 0$, we can apply Lemma 5.1 to conclude that $Z_p \Rightarrow Z$. □

Proof of Theorem 3.4. Suppose $\beta < 1$, Z_p is given by (1.8), and $Z_p(0) \Rightarrow \zeta(0)$, where $\zeta(0) > 0$ a.s. Let ζ be the solution to (1.10).

Note that $\beta < 1$ implies $\nu > 1$. Hence, (5.4) implies that (5.3) is satisfied and $\tilde{S}_p \to 0$ in probability. Therefore, $(Z_p(0), Y_p) \Rightarrow (Z(0), y)$ in $D_{\mathbb{R}^3}[0,\infty)$. By Theorem 5.4 in [24], $Z_p^\varepsilon \Rightarrow \zeta^\varepsilon$. By Corollary 5.6 in [24], if $Z_p(0) \to \zeta(0)$ in probability, then $Z_p^\varepsilon \to \zeta^\varepsilon$ in probability. By the same argument as above, this implies that Z_p converges to ζ in distribution or in probability, respectively. □

6. Fluctuations of Z_p.

In this section, we prove Theorem 3.5. Let us first recall the setting of that theorem. We have $\beta < 1$ and Z_p given by (1.8). Recall that the processes Z_p are all defined on the same probability space $(\Omega, \mathcal{F}, P)$. For each $p > 0$, $\zeta_p(0)$ is an $\mathcal{F}_0$-measurable random variable, where $\mathcal{F}_0$ is given by (3.1), such that $\zeta_p(0) > 0$ a.s. and $Z_p(0) - \zeta_p(0) \to 0$ in probability. The processes ζ_p and ξ_p are then given by (1.11) and (1.12).

To apply the theorems in [24], we wish to write ξ_p as the solution to a stochastic differential equation. By (1.11) and (4.1), we have

$$\begin{aligned}\xi_p(t) = \xi_p(0) &+ c_1(1-p)\int_0^t p^{-\tau}(Z_p(s-)^\alpha - \zeta_p(s)^\alpha)\,dm_p(s)\\ &- c_2\int_0^t p^{-\tau}(Z_p(s-)^\beta - \zeta_p(s)^\beta)\,dS_p(s)\\ &- c_2\int_0^t \zeta_p(s)^\beta\,dB_p(s) + R_p(t),\end{aligned} \tag{6.1}$$

where

$$B_p(t) = p^{-\tau}(S_p(t) - m_p(t)) = p^{(\nu-1)/2}\sum_{j=1}^{\lfloor tp^{-\nu}\rfloor}(\chi_j - p)$$

and

$$\begin{aligned}R_p(t) = p^{-\tau}&\int_0^t (c_1(1-p)\zeta_p(s)^\alpha - c_2\zeta_p(s)^\beta)\,d(m_p(s) - s)\\ &- c_1 p\int_0^t Z_p(s-)^\alpha\,dB_p(s) + p^{-\tau}L_p(t).\end{aligned} \tag{6.2}$$

Given a real number r, let us define the continuous function $F_r : (0,\infty)^2 \to \mathbb{R}$ by

$$F_r(x,y) = \frac{x^r - y^r}{x-y} 1_{\{x\neq y\}} + r y^{r-1} 1_{\{x=y\}}.$$

Using this, (6.1) becomes

$$\begin{aligned}\xi_p(t) &= \xi_p(0) + c_1(1-p)\int_0^t \xi_p(s-)\mathcal{D}_p^\alpha(s-)\, dm_p(s) \\ &\quad - c_2\int_0^t \xi_p(s-)\mathcal{D}_p^\beta(s-)\, dS_p(s) - c_2\int_0^t \zeta_p(s)^\beta\, dB_p(s) + R_p(t),\end{aligned} \tag{6.3}$$

where $\mathcal{D}_p^r = F_r(Z_p, \zeta_p)$.

Proof of Theorem 3.5. Suppose that there exists a pair of random variables $(\xi(0), \zeta(0))$, defined on $(\Omega, \mathcal{F}, P)$, such that $\zeta(0) > 0$ a.s. and $(\xi_p(0), \zeta_p(0)) \Rightarrow (\xi(0), \zeta(0))$. By the Skorohod representation theorem (see, for example, Theorem 2.1.8 in [18]), we can assume without loss of generality that $(\xi_p(0), \zeta_p(0)) \to (\xi(0), \zeta(0))$ a.s. Since the map that takes a point $x > 0$ to the unique solution to (1.11) with $\zeta_p(0) = x$ is continuous, $\zeta_p \to \zeta$ in probability and $(\xi_p(0), \zeta_p) \Rightarrow (\xi(0), \zeta)$. Also, since F_r is continuous, $\mathcal{D}_p^r \to r\zeta(\cdot)^{r-1}$ in probability.

Let

$$\begin{aligned}\mathcal{U}_p(t) &= \xi_p(0) - c_2\int_0^t \zeta_p(s)^\beta\, dB_p(s) + R_p(t), \text{ and} \\ \mathcal{Y}_p(t) &= c_1(1-p)\int_0^t \mathcal{D}_p^\alpha(s-)\, dm_p(s) - c_2\int_0^t \mathcal{D}_p^\beta(s-)\, dS_p(s),\end{aligned}$$

so that (6.3) becomes

$$\xi_p(t) = \mathcal{U}_p(t) + \int_0^t \xi_p(s-)\, d\mathcal{Y}_p(s). \tag{6.4}$$

We will apply the theorems in [24] to this integral equation.

We first show that $R_p \to 0$ in probability. By the Martingale Central Limit Theorem (Theorem 7.1.4 in [18]), $B_p \Rightarrow B$, where B is a standard Brownian motion; by Theorem 3.4, $Z_p \to \zeta$ in probability; and by (5.4), $\{B_p\}$ satisfies (5.3). Hence, by Theorem 2.2 in [24],

$$c_1 p\int_0^t Z_p(s-)^\alpha\, dB_p(s) \to 0$$

in probability. By (4.3), $p^{-\tau}L_p = 0$ on $[0, h_{p^\gamma\ell}(Z_p))$. Since $h_{p^\gamma\ell}(Z_p) \to \infty$ in probability, $p^{-\tau}L_p \to 0$ in probability.

For the final term in (6.2), note that $p^{-\tau}|m_p(t) - t| \le p^{\nu-\tau}$ and $\nu - \tau = (\nu+1)/2 > 0$. Hence, $p^{-\tau}(m_p(t) - t) \to 0$ uniformly. Let $f_p(s) = c_1(1-p)\zeta_p(s)^\alpha - c_2\zeta_p(s)^\beta$. Since $\zeta_p \to \zeta$ in probability, we can pass to a subsequence and assume that $\zeta_p \to \zeta$ uniformly on $[0,t]$, a.s. By (1.11), this implies that $\zeta_p' \to \zeta'$ uniformly on $[0,t]$. Hence, f_p and f_p' converge uniformly. Integrating by parts, we have

$$p^{-\tau}\int_0^t f_p(s)\,d(m_p(s)-s) = p^{-\tau}f_p(t)(m_p(t)-t) - p^{-\tau}\int_0^t (m_p(s)-s)f_p'(s)\,ds,$$

which goes to zero uniformly and completes the proof that $R_p \to 0$ in probability.

It now follows from Theorem 5.2 in [24] that $(\mathcal{U}_p, \mathcal{Y}_p, \zeta_p) \Rightarrow (\mathcal{U}, \mathcal{Y}, \zeta)$, where

$$\mathcal{U}(t) = \xi(0) - c_2\int_0^t \zeta(s)^\beta\,dB(s), \text{ and}$$
$$\mathcal{Y}(t) = c_1\int_0^t \alpha\zeta(s)^{\alpha-1}\,ds - c_2\int_0^t \beta\zeta(s)^{\beta-1}\,ds,$$

and B is a standard Brownian motion independent of $(\xi(0), \zeta(0))$. By Remark 2.5 in [24], we may apply Theorem 5.4 in [24] to (6.4) and conclude that $(\xi_p, \zeta_p) \Rightarrow (\xi, \zeta)$, where ξ is the unique solution to (1.13). □

7. Stationary Distributions. In this section, we prove Theorems 3.3 and 3.6. For this, we make time continuous in a slightly different manner than before. Let N be a unit rate Poisson process independent of $\{W_n\}$ and let $X(t) = W_{N(t)}$. Then X is a continuous time Markov chain on $E = [\ell, \infty)$ with generator

$$A\varphi(x) = p(\varphi(x - g(x)) - \varphi(x)) + (1-p)(\varphi(x + c_1x^\alpha) - \varphi(x)),$$

where $g(x) = (c_2x^\beta) \wedge (x - \ell)$. When $\beta = 1$, we will study the process

$$\hat{Z}_p(t) = p^\gamma X(tp^{-1}),$$

whereas when $\beta < 1$, we will consider

$$\hat{\xi}_p(t) = p^{-\tau}(p^\gamma X(tp^{-\nu}) - c_p),$$

where c_p is given by (1.14). It is easy to see that a probability measure is a stationary distribution for $\{p^\gamma W_n\}$ or $\{p^{-\tau}(p^\gamma W_n - c_p)\}$ if and only if it is a stationary distribution for $\hat{Z}_p$ or $\hat{\xi}_p$, respectively.

LEMMA 7.1. *If $\ell > 0$, then $\{W_n\}$ has a unique stationary distribution.*

Proof. It will suffice to show that X has a unique stationary distribution. Let $\varphi(x) = x$ so that

$$A\varphi(x) = -pg(x) + (1-p)c_1 x^{\alpha}.$$

Since $g(x) = c_2 x^{\beta}$ for x sufficiently large, $A\varphi$ is bounded above and $A\varphi(x) \to -\infty$ as $x \to \infty$. By Lemmas 4.9.5 and 4.9.7 in [18], the family of probability measures $\{\mu_t\}_{t\geq 1}$ defined by

$$\mu_t(\Gamma) = \frac{1}{t}\int_0^t P^x(X(s) \in \Gamma)\, ds$$

is relatively compact. By Theorem 4.9.3 in [18], any subsequential weak limit of $\{\mu_t\}$ is a stationary distribution for X.

To show that the stationary distribution is unique, it will suffice to show that for all $x \in E$,

$$\tau = \inf\{t \geq 0 : X(t) = \ell\} < \infty, \quad P^x\text{-a.s.}$$

(See, for example, Problem 4.36 in [18].) Let $x \in E$ be arbitrary and let $\varepsilon > 0$. Choose M such that $\mu_t([\ell, M]) \geq 1 - \varepsilon$ for all $t \geq 0$. Note that there exists $K > 0$ such that $P^y(\tau < \infty) \geq K$ for all $y \in [\ell, M]$.

Define the stopping times $\tau_0 = 0$ and

$$\tau_{j+1} = \inf\{t \geq \tau_j + 1 : X(t) \leq M\},$$

and note that $\tau_j \to \infty$ a.s. By the strong Markov property,

$$\begin{aligned} P(\tau = \infty, \tau_j < \infty) &= E[1_{\{\tau \geq \tau_j, \tau_j < \infty\}} P^{X(\tau_j)}(\tau = \infty)] \\ &\leq (1-K)P(\tau \geq \tau_j, \tau_j < \infty) \end{aligned}$$

Letting $j \to \infty$ shows that $P(\{\tau = \infty\} \cap D) = 0$, where D is the event that $\tau_j < \infty$ for all j. Note that

$$1_{D^c} \leq \liminf_{t\to\infty} \frac{1}{t}\int_0^t 1_{\{X(s) > M\}}\, ds.$$

Hence, by Fatou's Lemma, $P(D^c) \leq \liminf_{t\to\infty} \mu_t((M, \infty)) \leq \varepsilon$. Therefore, $P(\tau = \infty) = P(\{\tau = \infty\} \cap D^c) \leq \varepsilon$. Since ε was arbitrary, $\tau < \infty$ P^x-a.s. and the stationary distribution is unique. □

Proof of Theorem 3.3. In what follows, C and K will denote strictly positive, finite constants that do not depend on p and may change value from line to line.

Suppose $\beta = 1$, $\ell > 0$, and η_p is the stationary distribution for $\{p^\gamma W_n\}$. Then η_p is the stationary distribution for $\hat{Z}_p$, which is a continuous time Markov chain on $E_p = [p^\gamma \ell, \infty)$ with generator

$$\begin{aligned} A_p\varphi(x) &= \varphi(x - p^\gamma g(p^{-\gamma}x)) - \varphi(x) \\ &\qquad + p^{-1}(1-p)(\varphi(x + pc_1x^\alpha) - \varphi(x)). \end{aligned}$$

Let $\varphi(x) = x + x^{-1}$, so that

$$\begin{aligned} A_p\varphi(x) &= -p^\gamma g(p^{-\gamma}x) + (1-p)c_1x^\alpha \\ &\qquad + \frac{p^\gamma g(p^{-\gamma}x)}{x(x - p^\gamma g(p^{-\gamma}x))} - \frac{(1-p)c_1x^\alpha}{x(x + pc_1x^\alpha)}. \end{aligned}$$

Since $x \mapsto 1 + pc_1x^{\alpha-1}$ is decreasing,

$$1 + pc_1x^{\alpha-1} \leq 1 + pc_1(p^\gamma\ell)^{\alpha-1} = 1 + c_1\ell^{\alpha-1}$$

for all $x \in E_p$. Hence,

$$A_p\varphi(x) \leq -p^\gamma g(p^{-\gamma}x) + Cx^\alpha + \frac{p^\gamma g(p^{-\gamma}x)}{x(x - p^\gamma g(p^{-\gamma}x))} - Kx^{\alpha-2}$$

whenever $p < 1/2$.

If $x \geq p^\gamma\ell/(1-c_2)$, then $g(p^{-\gamma}x) = c_2p^{-\gamma}x$ and

$$A_p\varphi(x) \leq -Kx + Cx^\alpha + Cx^{-1} - Kx^{\alpha-2}.$$

If $x < p^\gamma\ell/(1-c_2)$, then $g(p^{-\gamma}x) = p^{-\gamma}x - \ell$ and

$$A_p\varphi(x) \leq Cx^\alpha + \frac{x - p^\gamma\ell}{xp^\gamma\ell} - Kx^{\alpha-2} \leq Cx^\alpha + (p^\gamma\ell)^{-1} - Kx^{\alpha-2}.$$

But in this case, $(p^\gamma\ell)^{-1} < Cx^{-1}$. It therefore follows that

$$A_p\varphi(x) \leq C - Kx - Kx^{\alpha-2}$$

for all $x \in E_p$.

Let $\varepsilon > 0$. Define

$$L = \sup_{p<1/2} \sup_{x\in E_p} A_p\varphi(x) < \infty$$

and let $m = L(1-\varepsilon)/\varepsilon$. Choose $M > 0$ such that $x \notin [M^{-1}, M]$ implies $A_p\varphi(x) < -m$ for all $p < 1/2$. By Corollary 4.9.8 in [18],

$$\eta_p([M^{-1}, M]) \geq \eta_p(\{x : A_p\varphi(x) \geq -m\}) \geq \frac{m}{L+m} = 1 - \varepsilon.$$

The family of measures $\{\eta_p\}$ is therefore relatively compact on $(0,\infty)$. By passing to a subsequence, we can assume that $\eta_p \Rightarrow \eta$ for some probability measure η on $(0,\infty)$.

Now let $p^\gamma W_0$ have distribution η_p and let Z_p be given by (1.8). By Theorem 3.2, $Z_p \Rightarrow Z$, where Z satisfies (1.9) with $PZ(0)^{-1} = \eta$. Fix $t_1 \leq \cdots \leq t_n$. Then

$$\begin{aligned}(Z_p(t_1), \ldots, Z_p(t_n)) &= p^\gamma(W_{\lfloor t_1 p^{-1}\rfloor}, \ldots, W_{\lfloor t_n p^{-1}\rfloor}) \\ &\overset{d}{=} p^\gamma(W_0, W_{\lfloor t_2 p^{-1}\rfloor - \lfloor t_1 p^{-1}\rfloor}, \ldots, W_{\lfloor t_n p^{-1}\rfloor - \lfloor t_1 p^{-1}\rfloor}) \\ &= (Z_p(0), Z_p(t_2 - t_1), \ldots, Z_p(t_n - t_1)) + \varepsilon,\end{aligned}$$

where $\varepsilon_j = Z_p(h_j) - Z_p(t_j - t_1)$ and $h_j = (\lfloor t_j p^{-1}\rfloor - \lfloor t_1 p^{-1}\rfloor)p$. Note that $h_j \to t_j - t_1$ as $p \to 0$ and, for fixed t, Z is almost surely continuous at t. Hence, $\varepsilon \to 0$ a.s., which gives

$$(Z_p(t_1), \ldots, Z_p(t_n)) \Rightarrow (Z(0), Z(t_2 - t_1), \ldots, Z(t_n - t_1)).$$

But

$$(Z_p(t_1), \ldots, Z_p(t_n)) \Rightarrow (Z(t_1), \ldots, Z(t_n)),$$

so Z is a stationary process, and η is a stationary distribution for Z. The uniqueness of η follows from Lemma 3.1. □

For the proof of Theorem 3.6, note that $\hat{\xi}_p$ is a continuous time Markov chain on $E_p = [p^{-\tau}(p^\gamma \ell - c_p), \infty)$ with generator

$$\begin{aligned}A_p\varphi(x) &= p^{-\nu+1}(\varphi(x - p^{\gamma-\tau}g(p^{\tau-\gamma}x + p^{-\gamma}c_p)) - \varphi(x)) \\ &\quad + p^{-\nu}(1-p)(\varphi(x + p^{\gamma-\tau}c_1(p^{\tau-\gamma}x + p^{-\gamma}c_p)^\alpha) - \varphi(x)).\end{aligned} \tag{7.1}$$

We will use the same argument as in the proof of Theorem 3.3, this time using the Lyapunov function $\varphi(x) = |x|^r$, where r is sufficiently large. Our key estimate on $A_p\varphi(x)$ is given in the following lemma and is valid as long as $|x|$ is not too large.

LEMMA 7.2. *Suppose $\beta < 1$. Let $\varphi(x) = |x|^r$, where $r \geq 2$, and let A_p be given by (7.1). Let $0 < \delta < M < \infty$ be arbitrary. Then there exists $p_0 > 0$ and strictly positive, finite constants C and K such that*

$$A_p\varphi(x) \leq C - K|x|^r$$

for all $p \leq p_0$ and all $x \in E_p$ satisfying $\delta \leq p^\tau x + c_p \leq M$.

Proof. For notational simplicity, let us define $y_p(x) = p^\tau x + c_p$ so that

$$\begin{aligned} A_p\varphi(x) = p^{-\nu+1}&(\varphi(x - p^{\gamma-\tau}g(p^{-\gamma}y_p)) - \varphi(x)) \\ &+ p^{-\nu}(1-p)(\varphi(x + p^{\gamma-\tau}c_1(p^{-\gamma}y_p)^\alpha) - \varphi(x)). \end{aligned}$$

Either $g(x) = c_2 x^\beta$ or $g(x) < c_2 x^\beta$. Note that there exists $x_0 > \ell$ such that $g(x) = c_2 x^\beta$ if and only if $x \geq x_0$. Hence, if $g(p^{-\gamma}y_p) < c_2(p^{-\gamma}y_p)^\beta$, then $p^{-\gamma}y_p < x_0$, which implies $x < p^{-\tau}(p^\gamma x_0 - c_p)$. If p is sufficiently small, this implies $x < 0$. Since φ is decreasing on $(-\infty, 0]$, it follows that

$$\begin{aligned} A_p\varphi(x) \leq p^{-\nu+1}&(\varphi(x - p^{\gamma-\tau-\gamma\beta}c_2 y_p^\beta) - \varphi(x)) \\ &+ p^{-\nu}(1-p)(\varphi(x + p^{\gamma-\tau-\gamma\alpha}c_1 y_p^\alpha) - \varphi(x)) \end{aligned}$$

for all $x \in E_p$.

Observe that

$$\begin{aligned} |\varphi(z) - \varphi(x) - \varphi'(x)(z-x)| &= \left| \int_x^z (z-u)\varphi''(u)\,du \right| \\ &\leq C|z-x|^2(|x|^{r-2} + |z|^{r-2}) \\ &\leq C|x|^{r-2}|z-x|^2 + C|z-x|^r. \end{aligned}$$

Hence,

$$\begin{aligned} A_p\varphi(x) \leq -\varphi'(x)p^{-\tau}&(p^{-\nu+1+\gamma-\gamma\beta}c_2 y_p^\beta - p^{-\nu+\gamma-\gamma\alpha}c_1(1-p)y_p^\alpha) \\ &+ C|x|^{r-2}(p^{-\nu+1+2\gamma-2\tau-2\gamma\beta}c_2^2 y_p^{2\beta} + p^{-\nu+2\gamma-2\tau-2\gamma\alpha}c_1^2 y_p^{2\alpha}) \\ &+ C(p^{-\nu+1+r\gamma-r\tau-r\gamma\beta}c_2^r y_p^{r\beta} + p^{-\nu+r\gamma-r\tau-r\gamma\alpha}c_1^r y_p^{r\alpha}). \end{aligned}$$

We can simplify these exponents by observing that

$$\begin{aligned} -\nu + \gamma - \gamma\alpha &= 0 \\ -\nu + 1 + \gamma - \gamma\beta &= 0 \\ -\nu + 2\gamma - 2\tau - 2\gamma\alpha &= 1 \\ -\nu + 1 + 2\gamma - 2\tau - 2\gamma\beta &= 0 \\ -\nu + 1 + r\gamma - r\tau - r\gamma\beta &= \tau(r-2) \\ -\nu + r\gamma - r\tau - r\gamma\alpha &= r - 1 + \tau(r-2). \end{aligned}$$

Thus,

$$\begin{aligned} A_p\varphi(x) \leq -\varphi'(x)p^{-\tau}&(c_2 y_p^\beta - c_1(1-p)y_p^\alpha) + C|x|^{r-2}(y_p^{2\beta} + p y_p^{2\alpha}) \\ &+ C(p^{\tau(r-2)}y_p^{r\beta} + p^{r-1+\tau(r-2)}y_p^{r\alpha}). \end{aligned}$$

Since $\varphi'(x)$ and $c_2 y_p^\beta - c_1(1-p)y_p^\alpha$ have the same sign, this gives

$$A_p\varphi(x) \le -r|x|^{r-1}p^{-\tau}|c_2 y_p^\beta - c_1(1-p)y_p^\alpha| + C|x|^{r-2}(y_p^{2\beta} + p y_p^{2\alpha}) + C(p^{\tau(r-2)}y_p^{r\beta} + p^{r-1+\tau(r-2)}y_p^{r\alpha}) \tag{7.2}$$

for all $x \in E_p$.

If $r \ge 2$ and $\delta \le y_p \le M$, then

$$A_p\varphi(x) \le -r|x|^{r-1}p^{-\tau}c_2 y_p^\alpha |y_p^{\beta-\alpha} - c_p^{\beta-\alpha}| + C|x|^{r-2} + C.$$

By the Mean Value Theorem,

$$\begin{aligned}\psi_p(x) &\le -K|x|^{r-1}p^{-\tau}|y_p - c_p| + C|x|^{r-2} + C\\ &= -K|x|^r + C|x|^{r-2} + C,\end{aligned}$$

which completes the proof. □

The following two lemmas provide the needed estimates on $A_p\varphi$ in the extreme regimes.

Lemma 7.3. *Suppose $\beta < 1$. Let $\varphi(x) = |x|^r$, where $r \ge 2$, and let A_p be given by (7.1). Then there exists $p_0 > 0$, $M < \infty$ and $K > 0$ such that*

$$A_p\varphi(x) \le -K|x|^{(r-1)\wedge(r-1+\beta)}$$

for all $p \le p_0$ and all $x \in E_p$ satisfying $p^\tau x + c_p > M$.

Proof. Let $p \le p_0$ and $y_p = p^\tau x + c_p > M$. If p_0 is sufficiently small and M is sufficiently large, then $x \ge Kp^{-\tau}$ and $y_p \le x$. By (7.2),

$$\begin{aligned}A_p\varphi(x) &\le -K|x|^{r-1}y_p^\beta + C|x|^{r-2}y_p^{2\beta} + C y_p^{r\beta}\\ &= -|x|^{r-1}y_p^\beta(K - C|x|^{-1}y_p^\beta - C|x|^{-r+1}y_p^{\beta(r-1)}).\end{aligned}$$

If $\beta \le 0$, then for p sufficiently small,

$$A_p\varphi(x) \le -|x|^{r-1}y_p^\beta(K - C|x|^{-1} - C|x|^{-r+1}) \le -K|x|^{r-1+\beta}.$$

If $\beta > 0$, then

$$A_p\varphi(x) \le -|x|^{r-1}y_p^\beta(K - C|x|^{\beta-1} - C|x|^{(\beta-1)(r-1)}),$$

so for p sufficiently small, $A_p\varphi(x) \le -K|x|^{r-1}y_p^\beta \le -K|x|^{r-1}$. □

LEMMA 7.4. *Suppose $\beta < 1$. Let $\varphi(x) = |x|^r$, where $r \geq 2$, and let A_p be given by (7.1). Then there exists $p_0 > 0$, $\delta > 0$ and $K > 0$ such that*

$$A_p\varphi(x) \leq -K|x|^{r\wedge(r-2\alpha/(1-\beta))}$$

for all $p \leq p_0$ and all $x \in E_p$ satisfying $p^\tau x + c_p < \delta$.

Proof. Let $p \leq p_0$ and $y_p = p^\tau x + c_p < \delta$. Note that since $x \in E_p$, $y_p \geq p^\gamma \ell$. If p_0 and δ are sufficiently small, then $x < 0$ and $Kp^{-\tau} \leq |x| \leq Cp^{-\tau}$. By (7.2), for δ sufficiently small,

$$\begin{aligned}
A_p\varphi(x) &\leq -|x|^r y_p^\alpha (K|y_p^{\beta-\alpha} - c_p^{\beta-\alpha}| - C(p^{2\tau}y_p^{2\beta-\alpha} + p^{2\tau+1}y_p^\alpha) \\
&\qquad - C(p^{\tau r+\tau(r-2)}y_p^{r\beta-\alpha} + p^{\tau r+r-1+\tau(r-2)}y_p^{r\alpha-\alpha})) \\
&\leq -|x|^r y_p^\alpha (K - C(p^{2\tau}y_p^{2\beta-\alpha} + p^{2\tau(r-1)}y_p^{r\beta-\alpha}) \\
&\qquad - C(p^{2\tau+1}y_p^\alpha + p^{(2\tau+1)(r-1)}y_p^{\alpha(r-1)})).
\end{aligned}$$

Let us first estimate the term $p^{2\tau}y_p^{2\beta-\alpha}$. If $2\beta-\alpha \geq 0$, then $p^{2\tau}y_p^{2\beta-\alpha} \leq Cp^{2\tau}$. If $2\beta - \alpha < 0$, then $p^{2\tau}y_p^{2\beta-\alpha} \leq Cp^{2\tau+\gamma(2\beta-\alpha)}$. Note that $2\tau + \gamma(2\beta - \alpha) = \gamma + 1$. Hence, for all values of α and β, there exists some $s > 0$ such that $p^{2\tau}y_p^{2\beta-\alpha} \leq p^s$.

Similarly, for the remaining terms in the above inequality, we observe that

$$\begin{aligned}
2\tau(r-1) + \gamma(r\beta - \alpha) &= (2\tau + \gamma\beta)(r-1) + 1 = \gamma(r-1) + 1 \\
2\tau + 1 + \gamma\alpha &= \gamma \\
(2\tau+1)(r-1) + \gamma\alpha(r-1) &= \gamma(r-1).
\end{aligned}$$

Therefore, if p_0 is sufficiently small, then $A_p\varphi(x) \leq -K|x|^r y_p^\alpha$. If $\alpha < 0$, then $A_p\varphi(x) \leq -K|x|^r$. If $\alpha \geq 0$, then

$$A_p\varphi(x) \leq -K|x|^r p^{\gamma\alpha} \leq -K|x|^{r-\gamma\alpha/\tau}.$$

Since $\gamma\alpha/\tau = 2\alpha/(1-\beta)$, this completes the proof. □

Proof of Theorem 3.6. Suppose $\beta < 1$ and η_p is the stationary distribution for $\{p^{-\tau}(p^\gamma W_n - c_p)\}$. Then η_p is the stationary distribution for $\hat{\xi}_p$. Let $\varphi(x) = |x|^r$, where $r \geq 2$. By Lemmas 7.2, 7.3, and 7.4, if r is sufficiently large, there exists $p_0 > 0$ and strictly positive, finite constants C and K such that

$$A_p\varphi(x) \leq C - K|x|^s$$

for some $s > 0$ and all $p \leq p_0$ and $x \in E_p$. As in the proof of Theorem 3.3, this implies that the family of measures $\{\eta_p\}$ is relatively compact on $\mathbb{R}$. By passing to a subsequence, we can assume that $\eta_p \Rightarrow \eta$ for some probability measure η on $\mathbb{R}$.

Let $p^{-\tau}(p^\gamma W_0 - c_p)$ have distribution η_p, let Z_p be given by (1.8), and let ξ_p be given by (1.12) with $\zeta_p \equiv c_p$. Note that $\xi_p(0)$ converges in distribution, so $p^\tau \xi_p(0) = Z_p(0) - \zeta_p(0) \to 0$ in probability. Hence, by Theorem 3.5, $\xi_p \Rightarrow \xi$, where ξ satisfies (1.15) with $P\xi(0)^{-1} = \eta$. As in the proof of Theorem 3.3, ξ is a stationary process, so η is the stationary distribution for ξ. □

Acknowledgements. The authors gratefully acknowledge Tom Kurtz, Ruth Williams, and Timo Seppäläinen for their support and helpful ideas. Thanks also to an anonymous referee for his or her helpful suggestions.

REFERENCES

[1] E. Altman, K. Avrachenkov, and C. Barakat. A stochastic model of TCP/IP with stationary random losses. *ACM SIGCOMM Computer Communications Review*, 30(4):231–242, 2000.
[2] E. Altman, K. Avrachenkov, and C. Barakat. TCP network calculus: the case of large delay-bandwidth product. *Proceedings IEEE Infocom 2002*, 2002.
[3] E. Altman, K. Avrachenkov, C. Barakat, A.A. Kherani, and B.J. Prabhu. Analysis of scalable TCP. *Proceedings IEEE HSNMC 2004*, pages 51–62, 2004.
[4] E. Altman, K. Avrachenkov, C. Barakat, and R. Nunez-Queija. TCP modeling in the presence of nonlinear window growth. *Proceedings of ITC-17*, 2001.
[5] E. Altman, K. Avrachenkov, A. Kherani, and B. Prabhu. Performance analysis and stochastic stability of congestion control protocols. *Proceedings of IEEE Infocom 2005*, 2005.
[6] E. Altman, K. Avrachenkov, and B. Prabhu. Fairness in MIMD congestion control algorithms. *Proceedings IEEE Infocom 2005*, 2005.
[7] E. Altman, C. Barakat, and V.M. Ramos. Analysis of AIMD protocols over paths with variable delay. *Proceedings IEEE Infocom 2004*, 2004.
[8] E. Altman, T. Jimenez, and D. Kofman. DPS queues with stationary ergodic service times and the performance of TCP in overload. *Proceedings IEEE Infocom 2004*, 2004.
[9] A. Baccelli, F.and Chaintreau, D. de Vleeschauwer, and D McDonald. A mean-field analysis of short lived interacting TCP flows. *Performance Evaluation*, 2004. Also Proceedings ACM Sigmetrics 2004.
[10] F. Baccelli, D. McDonald, and J. Reynier. A mean-field model for multiple TCP connections through a buffer implementing RED. *Performance Evaluation*, 49:77–97, 2002. Also Proceedings Performance 2002.
[11] J. Baras, A Misra, and T.J. Ott. The window distribution of multiple TCPs with random loss queues. *Proceedings of Globecomm'99*, 1999.
[12] J. Baras, A Misra, and T.J. Ott. Generalized TCP congestion avoidance and its effect on bandwidth sharing and variability. *Proceedings of Globecomm 2000*, 2000.
[13] J. Baras, A Misra, and T.J. Ott. Using drop-biasing to stabilize the occupancy of random drop queues with TCP traffic. *Proceedings of ICCS 2000*, 2000.
[14] J. Baras, A Misra, and T.J. Ott. Predicting bottleneck bandwidth sharing by generalized TCP flows. *Computer Networks*, 40(4):557–576, 2002.

[15] S. Bohacek. A stochastic model of TCP and fair video transmission. *Proceedings Infocom 2003*, 2003.

[16] A. Budhiraja, Hernandez-Campos F., V.G. Kukarni, and F.D. Smith. Stochastic differential equation for TCP window size: Analysis and experimental validation. *Probability in the Engineering and Informational Sciences*, 18:111–140, 2004.

[17] V. Dumas, F. Guillemin, and Ph. Robert. A Markovian analysis of additive-increase, multiplicative decrease (AIMD) algorithms. *Adv Appl Prob*, 34(1):85–111, 2002.

[18] Stewart N. Ethier and Thomas G. Kurtz. *Markov processes: Characterization and convergence*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, 1986.

[19] S. Floyd. TCP and explicit congestion notification. *ACM Computer Communications Review*, 21(5):8–23, 1994.

[20] F. Guillemin, Ph. Robert, and B. Zwart. AIMD algorithms and exponential functionals. *Ann Appl Prob*, 14(1):90–117, 2004.

[21] C. Hollot, V. Misra, D. Towsley, and W.B. Gong. A control theoretic analysis of RED. *Proceedings of Infocom 2001*, 2001.

[22] C. T. Kelly. Scalable TCP: Improving performance in high speed wide area networks. *ACM SIGCOMM Computer Communication Review*, 32(2):83–91, 2003.

[23] C. T. Kelly. *Engineering flow controls in the internet*. PhD thesis, Cambridge University, 2004.

[24] Thomas G. Kurtz and Philip Protter. Weak limit theorems for stochastic integrals and stochastic differential equations. *Ann. Probab.*, 19(3):1035–1070, 1991.

[25] T.V. Lakshman and U. Madhow. The performance of networks with high bandwidth-delay products and random loss. *IEEE/ACM Transactions on Networking*, 5(3):336–350, 1997.

[26] R. Marquez, E. Altman, and S. Sole-Alvarez. Modeling TCP and high speed TCP: a nonlinear extension to AIMD mechanisms. *Proceedings IEEE HSNMC 2004*, pages 132–143, 2004.

[27] Matthew Mathis, Jeffrey Semke, Jamshid Mahdavi, and Teunis J. Ott. The macroscopic behavior of the TCP congestion avoidance algorithm. *ACM SIGCOMM Computer Communication Review*, 27(3):67–82, 1997.

[28] A. Misra and T.J. Ott. The window distribution of idealized TCP congestion avoidance with variable packet loss. *Proceedings of IEEE Inforcom'99*, pages 1564–1572, 1999.

[29] A. Misra and T.J. Ott. Effect of exponential averaging and the variability of a RED queue. *Proceedings of ICC 2001*, 2001.

[30] A. Misra and T.J. Ott. Jointly coordinating ECN and TCP for rapid adaptation to varying bandwidth. *Proceedings of Milcom 2001*, 2001.

[31] A. Misra and T.J. Ott. Performance sensitivity and fairness of ECN-aware "modified TCP". *Journal of Performance Evaluation (PEVA)*, 53(3):255–272, 2003. Also published as Lecture Notes in Computer Science (LNCS), editors Gregori, Conti, Campbell, Omidyar, Zukerman.

[32] V. Misra, W.B. Gong, and D. Towsley. Stochastic differential equation modeling and analysis of TCP-windowsize behavior. *Proceedings of IFIP WG 7.3 Performance 1999*, 1999.

[33] Teunis J. Ott. ECN protocols and the TCP paradigm. Available at www.teunisott.com/Papers, 1999.

[34] Teunis J. Ott. Transport protocols in the TCP paradigm and their performance. *Telecommunication Systems*, 30(4):351–385, 2005.

[35] Teunis J. Ott. On the Ornstein–Uhlenbeck process with delayed feedback. Available

at www.teunisott.com/Papers, 2006.

[36] Teunis J. Ott. Rate of convergence for the “square root formula”. *Adv. in Appl. Probab.*, 38(4):1132–1154, 2006.

[37] Teunis J. Ott. The transient behavior of processes in the TCP paradigm. Work in Progress, 2007.

[38] Teunis J. Ott, J. H. B. Kemperman, and M. Mathis. The stationary behavior of idealized TCP congestion behavior. Available at `www.teunisott.com/Papers`, 1996.

[39] Teunis J. Ott and Jason Swanson. Stationarity of some processes in transport protocols. *SIGMETRICS Perform. Eval. Rev.*, 34(3):30–32, 2006. Also presented at the IFIP WG7.3 MAMA workshop, June 2006.

[40] T.J. Ott, T.V. Lakshman, and L.H. Wong. SRED: Stabilized RED. *Proceedings of IEEE Infocom'99*, pages 1346–1355, 1999.

[41] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling TCP throughput: A simple model and its empirical validation. *Proceedings ACM Sigcomm 1998*, 1998.

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

[43] Philip E. Protter. *Stochastic integration and differential equations. Second Edition.* Springer-Verlag, Berlin, 2004.

[44] K. K. Ramakrishnan, F. Floyd, and D. Black. The addition of explicit congestion control (ECN) to IP. *IETF RFC 3168*, 2001.

Winlab, Rutgers University
New Brunswick, NJ 07930
e-mail: ott@winlab.rutgers.edu

Mathematics Department
University of Wisconsin-Madison
480 Lincoln Dr.
Madison, WI 53706-1388
url: www.math.wisc.edu/∼swanson