

# STOCHASTIC DIFFERENTIAL EQUATION FOR TCP WINDOW SIZE

## *ANALYSIS AND EXPERIMENTAL VALIDATION*

**A. BUDHIRAJA**

*Department of Statistics  
University of North Carolina  
Chapel Hill, NC 27599*

**F. HERNÁNDEZ-CAMPOS**

*Department of Computer Science  
University of North Carolina  
Chapel Hill, NC 27599*

**V. G. KULKARNI**

*Department of Operations Research  
University of North Carolina  
Chapel Hill, NC 27599  
E-mail: vkulkarni@e-mail.unc.edu*

**F. D. SMITH**

*Department of Computer Science  
University of North Carolina  
Chapel Hill, NC 27599*

In this paper we develop a stochastic differential equation to describe the dynamic evolution of the congestion window size of a single TCP session over a network. The model takes into account recovery from packet losses with both fast recovery and time-outs, boundary behavior at zero and maximum window size, and slow-start after time-outs. We solve the differential equation to derive the distribution of the window size in steady state. We compare the model predictions with the output from the NS simulator.

## 1. INTRODUCTION

The dominant transport protocol for the Internet is TCP (transmission control protocol). The applications using this protocol range from simple e-mail exchanges to Web browsing and live audio or video streaming. Approximately 90% of the data carried by the Internet are governed by TCP. TCP provides an end-to-end (application-to-application) service that ensures the reliable, in-order, and unduplicated delivery of bytes flowing from a sender to a receiver. The mechanisms (sequence numbers explicitly acknowledged by the receiver, delivery timers, and retransmissions) used in TCP for providing this reliable data transport also provide a basis for end-to-end flow- and congestion-control mechanisms. Flow control is necessary to ensure that the sending application does not send data at a rate exceeding the rate at which the receiving application can process it. Congestion control is necessary to ensure that the sending application does not send data at a rate exceeding the currently available transmission capacity along the network path from the sender to the receiver.

Engineering experience has shown that the vast majority of lost data occurs because finite-capacity buffers overflow and arriving data are discarded at intermediate routing nodes in the Internet. TCP, therefore, treats the detection of lost data as a signal of congestion along the Internet path from the sender to the receiver. Buffer overflows mean that the arrival rate of in-bound data at a router destined for one of its out-bound links exceeds the transmission rate of that link. The overload may be a transient condition caused by bursty traffic demands or a persistent condition caused by poor network capacity planning. In response to lost data as an indication of congestion, the TCP at the sender reduces the maximum rate at which the application can send data. When there is no indication of congestion, TCP allows the sending rate to increase so it can accommodate the nominal rate used by the application.

Both flow and congestion controls are accomplished with a window-based transmission mechanism. A “window” is simply an amount of data the sender is allowed to have in transit in the network (sent but not yet acknowledged by the receiver) at any point in time and the window can be increased or decreased to govern the sender’s transmission rate. The multibyte data transmission unit created by TCP is called a segment and window size is often expressed as a number of maximum-size TCP segments. For flow control, the receiving TCP explicitly informs the sending TCP how much additional data beyond the segment being acknowledged it has buffers available to hold. This amount is the receiver-advertised window. A second window, the congestion window, is regulated in response to indications of congestion (or no congestion). At any time, the TCP sender is restricted to the minimum of the receiver-advertised window and the congestion window. The explicit increase and decrease algorithms used by TCP to adjust these windows are described in Section 3.

Because TCP is the dominant mechanism controlling the flow of data in the Internet, there has been intense research interest in producing good models of its dynamic window-adjustment behaviors (and therefore its control of transmission rates and application throughputs). These behaviors can be extremely complex de-

pending on the specific network conditions. A brief review of this research is given in Section 2. The practical applications for a model of TCP behavior have become more numerous in the recent past. TCP congestion-control mechanisms have been successful in allowing the Internet to handle an unanticipated scale of growth without suffering collapse in periods of congestion. It has, therefore, become the de facto standard for evaluating other congestion-control mechanisms. In particular, the other transport protocol for the Internet, UDP (user datagram protocol), does not provide any mechanisms for reliable data delivery or congestion control. It is up to applications that use UDP, notably those having multimedia data streams that require periodic real-time data delivery, to provide their own response to congestion conditions.

There is considerable concern in the Internet engineering community that UDP applications do not meet the de facto standard set by TCP for controlling congestion and that growth in their usage will lead to congestion collapse in the Internet. This has led to various proposals (e.g., [18]) for monitoring and policing UDP flows to ensure that they provide an effective level of congestion control. The basic notion is that routers can monitor their link queues to determine the rate of discarded data and can compare the data arrival rate from UDP flows with the result from a model of TCP throughput at the same loss rate. If the UDP flow persistently sends data at a higher rate, the router can subject it to a “penalty box,” in which a much higher fraction of its data is discarded. A related application for a TCP model is where a UDP application monitors its own rate of dropped data (with feedback from the receiver about those that do not arrive) and adjusts its sending rate to that determined by a model of TCP at the same loss rate (a so-called “TCP-friendly” behavior called “equation-based” congestion control [19], where the equation is an analytic TCP model). Recently, it was proposed in [22] to use a model of TCP performance at a certain loss rate to write contractual service agreements between an Internet service provider and its customers.

## 2. LITERATURE REVIEW

The large number of publications on modeling TCP behavior is a statement both about the difficulty of constructing a robust model of such complex behavior and the importance of TCP congestion control to the Internet. It is impossible for space reasons to comment on all of these models, but we do review the most prominent examples from several approaches used. All of the models described here consider the dynamic behavior of a single TCP connection as it would respond to loss events as indicators of link congestion. At a high level, the models are distinguished by three characteristics: the model of loss events (independent versus correlated), the aspects of TCP mechanisms assumed to have negligible influence on the results and not considered in the model, and whether the analysis computes expected values (e.g., the TCP window size or throughput) or produces a probability distribution of window sizes. An excellent discussion of many important issues in TCP model construction and validation can be found in [11].

The important distinctions between our model and the ones reviewed below are that it avoids certain common simplifying assumptions (e.g., an unbounded window

size or loss rates independent of window size) and it gives the probability distribution of window sizes.

## 2.1. Models with Independent Losses

An early derivation of a simple stochastic model of steady-state TCP behavior in its congestion-avoidance mode is presented in [32] and explained in more detail in Section 3. This model produced the well-known result that the mean TCP window size is  $C/\sqrt{p}$ , where  $p$  is the probability that an arbitrary segment will be independently dropped in the network and  $C$  is a constant term that reflects details of the TCP acknowledgment algorithm ( $C = 1.22$  if every segment is acknowledged and  $C = 0.87$  if at least every other segment is acknowledged). Given this mean window size estimate, throughput is bounded by  $(MSS/RTT) * (C/\sqrt{p})$ , where MSS is the maximum segment size and RTT is the mean round-trip time of the connection. This model is often called the “periodic-drop” model because the underlying analysis assumes that a constantly sending TCP source emits  $1/p$  segments between segment drops. An even earlier model derivation that is closely related can be found in [17,20].

The form of the model is to compute the mean window size based on its evolution as a series of discrete segment-by-segment events. The model only captures TCP window adjustments associated with TCP’s congestion-avoidance phase and only for segment drops that can be recovered without a time-out by reacting to three duplicate acknowledgments. Key elements of TCP behavior not considered in this model are the effects of time-out periods (in which throughput is zero), the TCP connection-establishment latency, the time TCP spends in its slow-start mode (also explained in Section 3), either following connection establishment or after a loss recovered by time-out, multiple losses in a single window of segments (that typically results in a time-out unless selective acknowledgment (SACK) is used), throughput limits imposed by the size of the receiver’s advertised window, and delayed acknowledgments. Because these elements are not considered, the duration (latency) of the connection cannot be derived for a given size of data transfer.

References [27,28] present an analysis of the mean throughput of a TCP connection that was developed concurrently with (and apparently independently of) [32,37]. Two cases are modeled. The first is one in which losses are not random but instead caused when the connection’s window grows to a point that causes a segment to be lost because the capacity of the buffer at the bottleneck link is exceeded (assuming only the one connection is using the link). This analysis considers both slow-start and congestion-avoidance phases. The second case considers both losses from overflow at the bottleneck buffer and random losses (from other causes), where the loss events are independent with fixed per-segment probability. Only congestion-avoidance with fast retransmission (no recovery with time-outs) is considered in this analysis. In both cases, connection establishment, receiver window limits, time-out intervals (zero throughput), and delayed acknowledgments are not modeled. Reference [29] extends the analysis in [27,28] to include asymmetric delays and congestion-induced loss on the “forward” (data) and “reverse” (acknowledgment) paths.

Reference [1] builds on the model and analysis from [27] to include loss models for i.i.d. losses and also a two-state (loss, no loss) first-order Markov model for the per-segment error process to include effects of correlated loss. The model also accounts for the effects of loss recovery by time-outs (including exponential back-off of the interval timer). Otherwise, it has the same assumptions and limitations as [27]. Reference [2] gives an earlier version of this model that considers only i.i.d. losses.

Reference [26] describes a Markov renewal-reward process model assuming independent segment losses. The model takes into account the explicit details of several versions of TCP (Tahoe, Reno, NewReno) along with factors such as time-out values and their granularity, receiver window size, and the number of duplicate acknowledgments used to infer loss. The detailed evolution of the window size for each TCP version is examined to determine the probability that multiple losses can be recovered without a time-out and the probability of recovery by time-outs. Reference [49] presents a modeling approach similar to [26] but using a two-state (loss, no loss) first-order Markov model for the per-segment error process to include effects of correlated loss.

References [38,39] cover the most widely cited and influential of the early TCP models. The new considerations in this model account for the effects of error recovery by retransmission after a time-out (given a mean time-out duration), the limits imposed by the maximum receiver's window (assumed to be constant at the value advertised at connection establishment), and the number of segments covered by an acknowledgment. The model of segment loss defines a probability,  $p$ , that a segment is lost given that no preceding segment from the window is lost (i.e., the probability that a given segment becomes the first loss within a window). Loss in one window is assumed independent of loss in any other window. To approximate the effects of correlated losses, if one segment from a window is lost, all following segments are considered lost also. Two assumptions also made explicit are that RTT and window size are independent [i.e., window size does not influence queuing delays (or that RTT consists only of propagation delays)], and that the time to transmit all segments in a window is less than the RTT.

Construction of the model in [39] proceeds from a detailed segment-by-segment analysis for the evolution of the TCP window as a stochastic process (Markov regenerative process with rewards). The result is an expression for the long-term expected value of throughput in segments/unit time,  $B$ , expressed as  $E(Y)/E(A)$ , where  $Y$  is the number of segments sent in intervals of time,  $A$ , between loss events of different types using the mean RTT as a measure of time. The model equation includes a term that expresses the  $C/\sqrt{p}$  relationship between mean window size and loss rate (e.g., the models of [32,37] are special cases of this model). They note also that for larger values of  $p$ , TCP window behavior is better approximated by  $C/p$  than by  $C/\sqrt{p}$ . The validity of this model is established with extensive empirical measurements of TCP connection throughputs along with the loss rates experienced.

In [22], the model of [39] is revised and extended. The primary motivation was to express the loss model in terms of independent random loss (Bernoulli) model with

intensity,  $p$ , which can be determined from conventional router-maintained counters. Using this model, the authors derive an expression for the probability that  $i$  segments will be lost from window of size  $W$  given independent losses, replacing the assumption that all segments following the first loss in a window are also lost (an approximation to correlated loss). From this, they also derive a new estimate of the probability of a time-out. A further refinement is to model the fact that the window size following loss events depends on the number of losses occurring in a window of size  $W$ , replacing the assumption in [39] that the window is reduced to  $W/2$ . Other than these two refinements, the model has the same assumptions and limitations as [39].

Reference [13] significantly extends the model in [38] to consider the effects of TCP connection establishment (including the possibility of segment loss), throughput during the initial slow-start phase (but not slow start following a time-out for a loss event), and the timing effects of delayed acknowledgments. Because the model is an extension of [38], the same assumptions and parameters are used. Another novel aspect of this model is that it is expressed in terms of expected values for connection latency given a size of data transfer (from which a throughput rate over the entire latency interval can also be calculated). Note that this allows the model to be accurate for transfers of any size, including those that are quite short (never encounter a loss or leave slow start). This model is validated through comparison with network simulator (NS) simulations and live Internet measurements.

The models presented in [44,45] are based on the same assumptions and some of the analysis and equations (such as approximations for correlated losses) from [13,38]. Their analysis and model differs from those earlier models because it is based on determining how many segments from an  $N$ -segment transfer are transmitted in each of the slow-start phases (both the initial one and following a loss event indicated by time-out) or in the congestion-avoidance phase as a function of the number of loss events. The approach in their analysis is somewhat similar to [26], in which the detailed evolution of the window size is examined to determine if multiple losses can be recovered without a time-out. They also develop a more refined model of the slow-start phase that accounts for effects of the delayed-acknowledgment timer. In [45], the analysis is refined to include specific details that reflect differences among TCP Tahoe, Reno, and SACK algorithms. They claim this model is more accurate for estimating total latency (and throughput) than [13] for short transfers but found little difference for transfers of more than 10,000 bytes. Results of empirical measurements (both new measurements and those in [39]) are presented to support this claim.

Reference [47] presents a generalization of the model in [39] to include two additional parameters,  $\alpha$  and  $\beta$ , where  $\alpha$  is the rate of increase in the window (measured in segments) when there is no loss in an RTT and  $\beta$  is the fraction of the current window that exists following a loss. In all other aspects, their model is identical to [39] and requires the same assumptions and has the same limitations related to elements of behavior not reflected in the model.

Reference [34] presents a model of steady-state TCP throughput. The new result is a model of window size behavior as a Poisson-Counter-driven stochastic differ-

ential equation where the loss events are treated as a Poisson stream of arrivals with rate  $\lambda$  at the source (instead of modeling segments as being sent out but lost with a given fixed probability) and the window size is approximated as a fluid (with continuous changes). Like [39], the model accounts for time-outs, receiver window size, and number of segments covered by an acknowledgment (but not connection establishment, initial slow start, and delays in returning acknowledgments). The result is an expression for the expected values of TCP window size and throughput as functions of the arrival rates of time-outs and duplicate-acknowledgment loss events, RTT, the mean time-out delay, and the maximum window size. There is no  $C/\sqrt{p}$  term in this model, but it was shown that the model includes those such as [32] as a special case. The model was also shown to produce results comparable to or better than [39] when compared to the same empirical measurements.

Stochastic differential equation models for TCP throughput have also been considered in [25]. The authors consider a network with multiple resources and users. The throughput rate for the  $r$ th user increases additively in the absence of loss signals and decreases multiplicatively at a rate proportional to the stream of feedback signals received from the various resources. The stream of feedback signals is modeled either via a Poisson process (with a stochastic rate function) or a diffusion process. The rate of the Poisson process (or the governing parameters in the diffusion), for a given resource, depends on the load on the resource. In that respect, the model is quite similar to the model studied in the current article; however, the authors do not take into account the finiteness of the receiver window size or time-outs. Furthermore, they do not obtain the asymptotic distribution of the window size or throughput rate.

Reference [36] used simulations to explore the implications of some assumptions common to a number of the analytical models (loss model, no lost acknowledgments, etc.). The authors examined the distribution of latencies for transfers of a given number of segments,  $N$ , and found that the shape of the distribution function was largely invariant for difference scenarios but was shifted or scaled by linear transformations. They suggested that this indicates that analytic models can be extended to compute distributions of latencies, not just first and second moments.

## 2.2. Models with Correlated Losses

The emphasis in [3] is on improving the realism of the loss models. In particular, the expressions derived include a term for the correlation function of the interloss times in addition to the usual parameters for the loss rate, RTT, number of segments per acknowledgment, the maximum receiver window, the mean time-out, and the window increase/decrease rates resulting from losses indicated by duplicate acknowledgments or time-outs. Loss processes can then be varied in the models by specifying the correlation function of interloss times. In terms of the elements of TCP behavior represented, this model is equivalent to [34,38]. It is also shown that for i.i.d. random losses, the  $C/\sqrt{p}$  result is a special case. An extensive set of empirical measurement is used to show that a variety of loss models are necessary to capture the

effects found on real Internet paths and that the proposed model produces good results for different patterns of loss actually observed. Other aspects of their work on TCP performance with correlated losses can be found in [4,5].

Reference [8] uses a two-state (loss, no loss) first-order Markov model for the per-segment error process to include effects of correlated loss. Five versions of TCP algorithms (from Tahoe to SACK) are modeled during steady-state data transfers. The TCP elements not considered are the initial connection-establishment phase and exponential back-off of the time-out value for multiple consecutive losses. The detailed analysis considers the number of segments sent as a “train” during each interval during the evolution of the TCP window (in both slow-start and congestion-avoidance phases) between loss-recovery events and produces an expression for the expected value of throughput.

### 2.3. Models with a Distribution of Window Sizes

The computation of the stationary distribution of the window size has not received as much attention in the literature as the expected window size. The early model in [37] (apparently never published) is the first use of a fluid approximation to the discrete TCP window-evolution process. It assumes that segment losses are independent events with probability  $p$ . In addition to the use of a fluid approximation, the other novel contribution is to give the stationary distribution function (actually the complementary c.d.f.) for the window size process. The mean of the window-evolution process is derived from this distribution with the final result being  $C/\sqrt{p}$  with a slightly different value of  $C = 1.31$ . The same elements of TCP behavior not considered in [32] are also not considered in this model.

Reference [33] presents a generalization of the model from [37] to consider segment loss probability that varies as a nondecreasing function of the TCP window size instead of being a constant probability. The effects of delayed acknowledgments are also modeled, but, otherwise, this model omits all the same elements of TCP behavior as [32,37], especially the limits imposed by the receiver’s window and the impact of time-outs. As in [37], the result is expressed in terms of computing the cumulative distribution function of the TCP window size in number of TCP segments. No throughput results are given, but they can presumably be computed from the distribution of window sizes given round-trip times. The results of computing window sizes with the model are shown to compare favorably with results from NS simulations.

The models described in [14,15] are closely related to [35,37] and mostly consider only those aspects of TCP behavior modeled in [32]. Specifically, they consider only idealized TCP steady-state behavior in the congestion-avoidance mode with independent segment losses of constant probability and no time-outs. They do, however, model the effects of the limitations imposed by the receiver’s window, and the increase/decrease ratios are parameterized as in [47]. The major contribution of these works is to provide convergence results for the distribution of window sizes

(and throughput) for small loss probabilities ( $p \rightarrow 0$ ). A subset of the limit results for certain special cases is given in [14].

In [15,33], the authors show that a time and space scaled version of the window-size process converges in distribution to a Markov process, which can then be analyzed by the standard methods to obtain the limiting distribution. The general nature of the loss function in [33] forces them to write the limiting distribution as an infinite sum of increasing dimensional multiple integrals, which can be evaluated numerically. The analysis in [15] is closer to our analysis. They derive the limiting distribution as an infinite sum of hypoexponential probability density functions (p.d.f.'s), whereas we get finite sums of simpler functions. Their approach seems to hide the interesting features of the limiting distribution, such as the atom at zero, and discontinuities in the p.d.f., that we have illustrated in this article.

Our model is described in more detail in Section 3. It takes into account recovery from packet losses with both fast recovery and time-outs, boundary behavior at zero and maximum window size, and slow start after time-outs. As far as we know, ours is the first model that accounts for all these aspects of TCP and also provides a distribution of window sizes.

## 2.4. Network-Level Models of TCP

In addition to research on models that yield the throughput of a single TCP connection for a particular packet loss process, there is also an active line of research on models that include more of the factors that influence TCP behavior directly in the model. The approaches that fall in this category generally consider multiple TCP connections that share network resources (link capacity, router buffers) and model their effects on each other's performance and on the utilization of the network resources. Many of these models also include queuing effects and, in particular, the behavior of the queue management algorithms at routers that ultimately determine how segments are dropped when link congestion occurs. Examples of models that consider elements of the entire network when modeling TCP connections are given in [6,7,9,10,21,23,24,30,31,35,41,43]. We do not discuss these further because our work presented here falls in the category of single TCP connection models.

## 3. TCP DYNAMICS

We consider a single application instance sending data over the Internet using one TCP connection. We briefly describe the features of TCP that are relevant to the models developed here. TCP implements window-based protocols that control the rate at which segments are transmitted by the sender to the receiver. For each connection, TCP maintains two variables: congestion window size and receiver window size. Both are specified in bytes, but because we assume that all segments have mean size  $B$ , we assume that the window sizes are in number of segments of size  $B$ . The minimum of the two window sizes is called the effective TCP window size, or just window size.

Let  $W_C(t)$  be the congestion window size (in number of segments) at time  $t$ , let  $W_R(t)$  be the receiver's window size at time  $t$ , and let RTT be the mean round-trip time (in seconds) for the TCP connection. The receiver's window size is the number of segments the receiver has space for and is typically a fixed quantity  $K$ . It rarely changes, due to the high speed at which the receiver's buffer can be emptied as compared to the speed at which the segments are received. Similarly, the round-trip time and the segment sizes may vary over the duration of the connection, but we will use their mean values in the analysis and define  $B$  as the mean segment size and RTT as the mean round-trip time over the connection duration. The TCP window size is given by

$$W(t) = \min\{W_C(t), W_R(t)\} = \min\{W_C(t), K\}.$$

The TCP congestion control has two phases: slow start and congestion-avoidance. In the congestion-avoidance phases, TCP behaves roughly as follows. At time  $t$ , the sender transmits  $W(t)$  number of segments back to back. As soon as the receiver receives a segment, an acknowledgment for that segment is sent back to the sender. (We ignore the batching of acknowledgments that is done by some TCP versions.) Thus, by time  $t + \text{RTT}$ , the sender receives  $W(t)$  acknowledgments, if there are no problems in transmission and delivery of the segments.

If the segment is error-free and is in the correct sequence, the receiver sends an acknowledgment for the highest sequence number received correctly. If a segment has unrecoverable errors or is out of sequence, the receiver continues to send the highest sequence number received correctly in its acknowledgment (sends a duplicate acknowledgment). A loss is indicated if the sender receives three duplicate acknowledgments or if the time-out period for a segment expires without an acknowledgment for its sequence number.

The congestion window  $W_C(t)$  changes dynamically at every instant an acknowledgment is received. In the absence of any loss indications, the window size increases by one every round-trip time. If an acknowledgment is received that is a third duplicate acknowledgment (indicating a segment loss), the window is reduced to  $(1 - \alpha)W(t) + 1$  (typically  $\alpha = 0.5$ ), and the TCP protocol continues in the congestion-avoidance phase. This operation is called fast recovery. We use  $(1 - \alpha)W(t) + 1$  in place of the customary  $(1 - \alpha)W(t)$  for two reasons: TCP's congestion window is never reduced below one segment and increasing the window slightly is an approximation to account for the segments sent during fast recovery before the window is reduced to its final value. If a time-out occurs, the window size is reduced to one. After time-out, the TCP protocol switches to the slow-start phase and rapidly increases the window size to one plus  $(1 - \alpha)$  fraction of the prevailing window size just before the time-out occurred. Again, there are many variations of TCP; we will analyze this one. Note that the response to a segment loss is determined by  $W(t)$ , not  $W_C(t)$ .

We will study the dynamics of the window size process  $\{W(t), t \geq 0\}$  in two steps. In the next section, we assume that there are no time-outs and develop the

probabilistic analysis of the window size process in steady state. In Section 4, we extend this analysis to include the time-outs.

#### 4. TCP MODEL WITH NO TIME-OUTS

As a first step toward the analysis of the complete TCP dynamics, we begin with the assumption that there are no time-outs. Thus, once the TCP gets out of the slow-start phase, it stays in the congestion-avoidance phase forever. Assume that each returning acknowledgment indicates a segment loss with probability  $\hat{p}$ , independent of anything else. Let  $X(t)$  be the number of acknowledgments received during  $(t, t + \text{RTT})$  that indicate segment losses. With the assumption of independence, we see that  $X(t)$  is a binomial  $(W(t), \hat{p})$  random variable. We assume that  $\hat{p}$  is sufficiently small so that we can ignore the possibility that  $X(t) \geq 2$ , and, thus,  $X(t)$  can be taken to be a Bernoulli random variable with

$$P(X(t) = 1) = \hat{p}W(t), \quad P(X(t) = 0) = 1 - \hat{p}W(t).$$

Thus, in effect, we assume that there is at most one segment loss per RTT. Now, as long as  $0 < W_c(t) < K$ , we have  $W(t) = W_c(t)$ , and the dynamics of the TCP protocol can be approximated as

$$W(t + \text{RTT}) = (W(t) + 1)(1 - X(t)) + (1 + (1 - \alpha)W(t))X(t).$$

Rearranging, we get

$$W(t + \text{RTT}) - W(t) = 1 - \alpha W(t)X(t). \quad (4.1)$$

Hence, we get

$$\begin{aligned} W(t + \text{RTT}) - W(t) &= \frac{1}{\text{RTT}} \text{RTT} - \alpha W(t)X(t) \\ &= (f)\text{RTT} - \alpha W(t)X(t), \end{aligned} \quad (4.2)$$

where

$$f = \frac{1}{\text{RTT}}.$$

Now, we approximate the distribution of  $X(t)$  by that of  $N(t + \text{RTT}) - N(t)$ , where  $\{N(t), t \geq 0\}$  is a Poisson process with a time-varying stochastic rate function  $pW(t)$ , where  $p = \hat{p}/\text{RTT}$ . Thus, Eq. (4.2) can be written as

$$W(t + \text{RTT}) - W(t) = (f)\text{RTT} - \alpha W(t)(N(t + \text{RTT}) - N(t)).$$

The above equation can be thought of as the finite-difference version of the following stochastic differential equation for  $\{W(t), t \geq 0\}$ :

$$dW(t) = f dt - \alpha W(t) dN(t), \quad 0 < W(t) < K. \quad (4.3)$$

Next, we modify the above equation to account for the boundary behavior at  $W(t) = K$ . Once  $W(t)$  reaches  $K$ ,  $W_C(t)$  may keep changing, but  $W(t)$  stays at  $K$  until a segment loss occurs. We incorporate this condition by modifying the stochastic differential equation (4.3) as follows:

$$dW(t) = f1_{\{W(t) < K\}} dt - \alpha W(t) dN(t), \quad 0 \leq W(t) \leq K. \quad (4.4)$$

The above stochastic differential implies that the  $W$  process increases continuously at a constant rate  $f$  as long as it is less than  $K$  and there are no events in the  $N$  process. When it hits the upper limit  $K$ , it stays at  $K$ . When an event occurs in the  $N$  process, the  $W$  process jumps down by a factor  $1 - \alpha$ . Figure 1 shows a typical sample path of the  $W$  process.

### 5. LIMITING BEHAVIOR OF $W$ WITH NO TIME-OUTS

In this section, we study the limiting behavior of the TCP window size, assuming that it evolves as described in the previous section; that is, we obtain the limiting distribution of the solution to Eq. (4.4) as  $t \rightarrow \infty$ . Let

$$F(x) = \lim_{t \rightarrow \infty} P(W(t) \leq x), \quad 0 \leq x \leq K,$$

be the limiting distribution of the  $W(\cdot)$  process. The sample path of the  $W$  process shows that the limiting distribution of  $W$  has a probability mass  $v_K$  at  $x = K$  and a probability density  $u(x)$  (defined as the left derivative of  $F$ ) over  $x \in [0, K)$ . Thus,




FIGURE 1. Sample paths of  $W$  and  $W_C$  processes in the absence of time-outs.

$$F(x) = \int_0^x u(y) dy, \quad 0 \leq x < K,$$

$$F(K) = \int_0^K u(y) dy + v_K = 1.$$

The following theorem gives  $v_K$  and  $u(\cdot)$ .

**THEOREM 5.1:** *The density  $u(x)$ ,  $0 < x < K$ , is given by*

$$u(x) = a_0 \sum_{r=0}^n \alpha_{n-r} \beta_r g_r(x), \quad x \in (K(1-\alpha)^{n+1}, K(1-\alpha)^n), n \geq 0, \quad (5.1)$$

where  $a_0$  is a normalizing constant,

$$g_r(x) = \exp\left\{\frac{-px^2}{2f(1-\alpha)^{2r}}\right\}, \quad x \geq 0, r \geq 0, \quad (5.2)$$

and the constants  $\{\alpha_i\}$  and  $\{\beta_i\}$  are

$$\beta_0 = 1, \quad \beta_{r+1} = \frac{(1-\alpha)^{2r}}{(1-\alpha)^{2r+2} - 1} \beta_r, \quad r \geq 0, \quad (5.3)$$

$$\alpha_0 = 1,$$

$$\alpha_1 = \frac{1}{g_0(K(1-\alpha))} (g_0(K(1-\alpha)) - (1 + \beta_1)g_1(K(1-\alpha))),$$

$$\begin{aligned} \alpha_{r+1} = & \frac{1}{g_0(K(1-\alpha)^r)} \sum_{j=0}^r \alpha_j (\beta_{r-j} g_{r-j}(K(1-\alpha)^{r+1}) \\ & - \beta_{r+1-j} g_{r+1-j}(K(1-\alpha)^{r+1})) \end{aligned} \quad (5.4)$$

for  $r > 0$ . The probability mass  $v_K$  is given by

$$v_K = \frac{fu(K-)}{pK}. \quad (5.5)$$

Finally, the normalizing constant  $a_0$  is chosen so that

$$\int_0^K u(y) dy + v_K = 1.$$

**PROOF:** Because  $W$  is an irreducible Markov process on  $[0, K]$ , it has a stationary distribution  $F$ . Let  $W(0)$  have distribution  $F$ . Then,  $W(t)$  has distribution  $F$  for all  $t \geq 0$ . Then, for  $0 < x < K$ ,

$$\begin{aligned}
 F(x) &= P(W(h) \leq x) \\
 &= \int_{[0, K]} P(W(h) \leq x | W(0) = y) dF(y) \\
 &= ph \int_{\left[0, \min\left(\frac{x}{1-\alpha}, K\right)\right]} y dF(y) + \int_{[0, x-fh]} (1 - phy) dF(y) + o(h) \\
 &= ph \int_{\left[0, \min\left(\frac{x}{1-\alpha}, K\right)\right]} y dF(y) + F(x - fh) - + ph \int_{[0, x-fh]} y dF(y) + o(h).
 \end{aligned}$$

Thus, we have that

$$\frac{F(x) - F(x - fh)}{h} = p \int_{\left[0, \min\left(\frac{x}{1-\alpha}, K\right)\right]} y dF(y) - p \int_{[0, x-fh]} y dF(y) + o(1).$$

Taking the limit as  $h \rightarrow 0$ , we have that

$$fu(x) = p \int_{\left[0, \min\left(\frac{x}{1-\alpha}, K\right)\right]} y dF(y) - p \int_{[0, x]} y dF(y), \quad x \in (0, K]. \quad (5.6)$$

In particular, we have that

$$fu(K-) = pKv_K,$$

which yields Eq. (5.5).

Consider now the case when  $x \in I_0 \equiv (K(1 - \alpha), K)$ . Then, from (5.6), it follows that

$$fu(x) = p \int_{[0, K]} y dF(y) - p \int_{[0, x]} y dF(y).$$

Differentiating, we have that

$$fu'(x) = -pxu(x),$$

and solving the above equation, we get

$$u(x) = a_0 g_0(x),$$

where  $a_0$  is some constant. Thus, we have shown (5.1) for the case  $n = 0$ . We will prove (5.1) for a general  $n$  via induction. Now, suppose that (5.1) holds for  $n = 0, 1, \dots, j$ . We begin by observing that from (5.6), it follows that for  $x \in [0, K(1 - \alpha))$ ,

$$u'(x) = p \frac{px}{f(1 - \alpha)^2} u\left(\frac{x}{1 - \alpha}\right) - \frac{px}{f} u(x). \quad (5.7)$$

Thus, from the induction hypothesis, it follows that for  $x \in (K(1 - \alpha)^{j+2}, K(1 - \alpha)^{j+1})$ ,

$$fu'(x) = p \frac{x}{(1 - \alpha)^2} \left\{ a_0 \sum_{r=0}^j \alpha_{j-r} \beta_r g_r \left( \frac{x}{1 - \alpha} \right) \right\} - pxu(x).$$

We can solve the above differential equation to get

$$\begin{aligned} u(x) &= \frac{pa_0}{(1 - \alpha)^2} e^{-(px^2/2f)} \sum_{r=0}^j \alpha_{j-r} \beta_r \int x g_r \left( \frac{x}{1 - \alpha} \right) e^{px^2/2f} dx + a_{j+1} e^{-(px^2/2f)} \\ &= \frac{a_0}{(1 - \alpha)^2} \sum_{r=0}^j \alpha_{j-r} \beta_r g_{r+1}(x) \left( \frac{(1 - \alpha)^{2r+2}}{(1 - \alpha)^{2r+2} - 1} \right) + a_{j+1} e^{-(px^2/2f)} \\ &= a_0 \sum_{r=0}^j \alpha_{j-r} \beta_{r+1} g_{r+1}(x) + a_{j+1} \beta_0 g_0(x), \end{aligned} \quad (5.8)$$

where  $a_{j+1}$  is some constant.

In order to determine  $a_{j+1}$ , we now consider the cases  $j = 0$  and  $j \geq 1$  separately.

*Case 1:  $j = 0$ :* In this case, we have from (5.6) that

$$u(K(1 - \alpha)^+) = u(K(1 - \alpha)^-) + \frac{pKv_K}{f}.$$

Using this observation, the induction hypothesis, and (5.8), we now have that

$$a_0 \beta_1 g_1(K(1 - \alpha)) + a_1 g_0(K(1 - \alpha)) = a_0 g_0(K(1 - \alpha)) - \frac{pKv_K}{f}.$$

Thus, solving for  $a_1$ , we get

$$\begin{aligned} a_1 &= \frac{a_0}{g_0(K(1 - \alpha))} (g_0(K(1 - \alpha)) - (\beta_1 + 1)g_1(K(1 - \alpha))) \\ &= a_0 \alpha_1. \end{aligned}$$

*Case 2:  $j \geq 1$ :* Next, note that from (5.6), it follows that  $u(\cdot)$  is continuous on  $(0, K(1 - \alpha))$ ; thus, if  $j \geq 1$ , then  $u(K(1 - \alpha)^{j+1}+) = u(K(1 - \alpha)^{j+1}-)$ . Using this observation, (5.8), and the induction hypothesis, we have that

$$\begin{aligned} &a_0 \sum_{r=0}^j \alpha_{j-r} \beta_{r+1} g_{r+1}(K(1 - \alpha)^{j+1}) + a_{j+1} \beta_0 g_0(K(1 - \alpha)^{j+1}) \\ &= a_0 \sum_{r=0}^j \alpha_{j-r} \beta_r g_r(K(1 - \alpha)^{j+1}). \end{aligned}$$




FIGURE 2. Typical graph of the density  $u$  and  $u^T$  of the window size.

Solving for  $a_{j+1}$  and rearranging the terms, we now have that

$$\begin{aligned}
 a_{j+1} &= \frac{a_0}{g_0(K(1-\alpha)^{j+1})} \\
 &\times \sum_{r=0}^j \alpha_r \{ \beta_{j-r} g_{j-r}(K(1-\alpha)^{j+1}) - \beta_{j+1-r} g_{j+1-r}(K(1-\alpha)^{j+1}) \} \\
 &= a_0 \alpha_{j+1}.
 \end{aligned}$$


Substituting the above observations in (5.8), we have (5.1) for the case  $n = j + 1$ . This proves the result. ■

The typical form of the limiting distribution is shown by the upper graph in Figure 2. Note that  $u(\cdot)$  has a jump discontinuity at  $x = K(1 - \alpha)$  and its derivative has a jump discontinuity at  $x = K(1 - \alpha)^2$ . The lower graph shows  $u^T$ , the limiting density of the window size in the presence of time-outs, which is studied in the next section.

### 6. TCP MODEL WITH TIME-OUTS

The above analysis does not account for the response of TCP for lost segments that will be detected by time-outs. Typically, such events are accompanied by intervals of zero throughput for a random amount of time. The reason is that the sender at some point reaches the limit of its current effective window, but the lost segment’s time-out interval (which depends on an estimated RTT and thus behaves like a random

variable) has not occurred to initiate recovery and start sending again. We make the assumption that a fixed fraction  $c$  of the events in the  $N$  process are time-outs, and the remaining  $1 - c$  fraction of the events are fast-recovery events. We assume that the effective window size jumps down to zero in the event of a time-out and it stays zero for an exponential duration with mean  $1/\mu$ , the expected duration of the time-out interval. We assume that after this zero window-size interval, the window size jumps instantaneously to  $(1 - \alpha)W^-$ , where  $W^-$  is the window size when the time-out occurred. In other words, we approximate the rapid exponential increase in the window size during the slow-start phase by an instantaneous jump at the end of the time-out period. The window size then begins to evolve according to Eq. (4.4) until the next loss event is encountered. We denote the window-size process in the presence of time-outs by  $\{W^T(t), t \geq 0\}$ , to distinguish it from the window size process  $\{W(t), t \geq 0\}$  without the time-outs. Figure 3 shows the sample paths of  $W^T(t)$  and  $W_C(t)$  and their relation to each other, taking into account the fast recovery, the time-out events and the boundary behavior at  $K$  and zero. It is clear the sample paths of the  $W^T$  process can be thought of as the sample paths of the  $W$  process, where periods of zero window size are inserted after a fraction  $c$  of the downward jumps in the  $W$  process in a completely random fashion. This observation implies that there is a very close relation between the limiting c.d.f. of the  $W$  process and that of the  $W^T$  process. This relation is made clear in Theorem 6.1.



**FIGURE 3.** Sample paths of  $W_C(t)$  and  $W(t)$  in the presence of time-outs.

The sample path of the  $W^T$  process shows that the limiting distribution of  $W^T$  has a probability mass  $v_0^T$  at  $x = 0$  and  $v_K^T$  at  $x = K$ , and a probability density  $u^T(x)$ ,  $0 < x < K$ . Let

$$F^T(x) = \lim_{t \rightarrow \infty} P(W^T(t) \leq x), \quad 0 \leq x \leq K,$$

be the limiting distribution of the  $W^T(\cdot)$  process.

Theorem 6.1 gives  $v_0^T$ ,  $v_K^T$ , and  $u^T(\cdot)$  in terms of the quantities  $v_K$  and  $u(\cdot)$  of Theorem 5.1 and

$$m = \int_0^K yu(y) dy + Kv_K.$$

THEOREM 6.1: *We have*

$$v_0^T = \frac{mpc}{\mu + mpc}, \quad (6.1)$$

$$u^T(x) = (1 - v_0^T)u(x), \quad 0 < x < K, \quad (6.2)$$

$$v_K^T = (1 - v_0^T)v_K. \quad (6.3)$$

PROOF: Let

$$F^T(x) = \lim_{t \rightarrow \infty} P(W^T(t) \leq x), \quad 0 \leq x \leq K.$$

Since the window size is zero when the time-out is in progress, we see that

$$\lim_{t \rightarrow \infty} P(W^T(t) \leq x \mid \text{Time-out is in progress at time } t) = 1, \quad 0 \leq x \leq K.$$

Also,

$$\lim_{t \rightarrow \infty} P(W^T(t) \leq x \mid \text{Time-out is not in progress at time } t) = F(x), \quad 0 \leq x \leq K,$$

where  $F(x)$  is from Theorem 5.1.

It is clear that  $v_0^T$  is the fraction of the time the  $W^T$  process is in the time-out period. Hence, conditioning on whether the time-out is in progress or not, we get

$$F^T(x) = v_0^T + (1 - v_0^T)F(x).$$

Equations (6.2) and (6.3) immediately follow from this. We thus need to derive Eq. (6.1) to complete the proof. Note that the rate at which the  $W^T$  process leaves state 0 in steady state is  $v_0^T\mu$ . The rate at which it enters state 0 is given by

$$\int_0^K y dF^T(y) pc = (1 - v_0^T) \int_0^K dF(y) pc = (1 - v_0^T)mpc.$$

Equating these two, we get Eq. (6.1). This completes the proof. ■

Figure 2 shows the typical shape of the density function  $u^T$ . It also shows the two probability masses  $v_0^T$  and  $v_K^T$ . The fact that  $u^T(0) = 0$  is the result of our approximation of the exponential growth of window size in the slow-start phase by an instantaneous jump. In the next section, we validate the analytical results of this section by simulation using NS. As expected, we find that the empirical p.d.f. of the window size near zero is indeed not zero. An extension of the analytical technique to better account for the slow-start growth is complicated, but it is under consideration.

## 7. VALIDATION BY SIMULATION

To evaluate the effectiveness of this model at representing TCP behavior, we compared the results from the analytic model with the results from simulations using the NS simulator [12]. This simulator is widely used in networking research and has a very detailed model of TCP error recovery and congestion-control mechanisms. It has been subjected to careful validations and is regarded as a standard tool for TCP research. The simulated environment we chose for evaluating the model emphasizes the role of link-level congestion as the dominant factor that causes segment losses. For our experiments, we constructed a model of an enterprise or campus network having a single wide-area link to an upstream Internet service provider (ISP). The high-level view of the simulations is that a single, continuous TCP flow sending at the maximum rate it can achieve shares a bottleneck link with a large number of TCP flows used for Web-like traffic. The motivation for choosing Web-like traffic to share the link was the assumption that this traffic would exhibit highly variable and bursty demands on the link (specifically, a self-similar segment arrival process). The empirical data used to generate our Web traffic [16] had heavy-tailed distributions for both “think” (OFF) times and response size (ON) times, which would lead to self-similar traffic from the aggregation of sources [46]. This segment arrival process entering the fixed capacity queue in the router serving the congested link provides a realistic simulation of the loss processes found in the Internet when router buffers overflow. Thus, the simulation does not assume a particular loss model (independent, correlated, etc.), but, instead, the loss process is a direct result of running the simulation with different levels of load from self-similar Web traffic.

We study the dynamics of the single continuously sending TCP flow when it shares this congested link with many other TCP flows carrying Web-like traffic. The simulated Web clients (browsers) are all located on the enterprise or campus network and all of their requests are satisfied by Web servers located somewhere on the Internet beyond the ISP link. The traffic model that drives the simulation is based on a contemporary model of Web browsing reported in [16]. This model is an application-level description of the critical elements that characterize how HTTP/1.0 protocol is used. It is based on empirical data and is intended for use in generating synthetic Web workloads. Because response sizes are much larger than requests in this traffic model, the load on the network link that carries traffic from the servers to the browsers will be much greater than on the link carrying traffic in the opposite direction that

is basically uncongested. The continuously sending TCP flow also sends its segments along the congested path carrying Web traffic from servers to clients.

The load generated by the Web-like flows at the bottleneck link was varied in different runs of the simulation from 40% to 98% of a link capacity that was fixed at 10,000,000 bits per second. The continuous TCP flow attempted to send at the maximum rate it could sustain, given the competing load from the Web traffic on the shared link. For example, at a 40% load for Web traffic, the continuous flow was able to send at a rate equivalent to nearly 60% of the link capacity. Note that because the continuous TCP flow can use all of link bandwidth the web traffic is not using, the link is congested in all cases. As the load from the Web traffic is increased, the continuously sending TCP flow will experience greater segment loss and lower throughput.

The version of TCP simulated is TCP Reno and the delayed-acknowledgments option was turned off. The continuous TCP flow had a maximum receiver's window size of 45 segments of 1460 bytes. The buffer size in the router was 60 of these maximum size segments.

Each simulation was run for 60 min of simulated time and data from the first 10 min was discarded to eliminate transient start-up effects. Scripts for running the NS simulations which also show various parameter settings can be found in [50]. In each simulation run, we recorded for the continuous TCP flow, the number of segments transmitted, the number of three-duplicate-acknowledgment loss events, the number of time-out loss events, the mean round-trip time, the mean throughput, and the effective size of the TCP window sampled every 10 ms. The frequencies of loss events experienced by the continuous TCP flow in the simulations were used as loss probabilities in the analytic TCP model for a comparison of results. The mean round-trip time (including transmission, propagation, and queuing delays) for the continuous TCP flow in the simulations was used as the mean RTT in the analytic model. The estimates of  $p$ ,  $c$ , and RTT (in milliseconds) are given in Table 1 for the seven different load conditions. They are also shown in Figures 5 and 6. Figure 5 shows the values of  $p1 = p(1 - c)$  and  $p2 = pc$ ; Figure 6 shows the plot of the mean RTT (in milliseconds) for the seven runs.

In the NS simulation of TCP, the effective size of the TCP window is not reset to its minimum size (one segment) until the end of a time-out period, whereas the

**TABLE 1.** Estimates of  $p$ ,  $c$ , and RTT for the Seven Runs

	Load						
	40%	50%	60%	70%	80%	90%	98%
$p$	0.0007	0.0010	0.0026	0.0042	0.0090	0.0135	0.0195
$c$	0.3675	0.3316	0.3023	0.2760	0.4338	0.5332	0.6350
RTT	74.7	76.1	77.8	79.5	81.0	0.0843	88.2

analytic model considers the effective window to be minimum (zero segments) for a random interval during the time-out period. The TCP window sizes sampled at 10-ms intervals during the simulations were adjusted to match this definition by setting the sample value to zero if the sample was taken during a time-out period. If the sample was not taken during a time-out period, the sample value was recorded without change.

## 8. COMPARISON OF ANALYTIC AND SIMULATION RESULTS

Figure 4 compares the throughput for the continuous TCP flow observed in the simulation with the throughput predicted by the analytic model (the line labeled “theory” in the plot) for several link loads generated by background Web traffic. The loss probabilities used in the analytic model for each load were those values observed in the corresponding simulation as given in Figure 5. For example, at a load of 70% Web traffic, the simulation results for the continuous TCP flow showed a per-packet probability of a loss indicated by three duplicate acknowledgments of 0.003 and per-packet probability of a loss indicated by a time-out of 0.0012. The mean RTT used in the analytic model for each load were those values observed in the corresponding simulation as given in Figure 6. For example, at a 70% load, the mean RTT for the continuous TCP flow was 79 ms.

The analytic model predicts slightly lower throughput only at the lowest loads. At all higher loads where the TCP connection experienced more loss and queuing delays, the simulation and analytic models produce essentially the same results. Thus, for a single TCP connection, the model presented here can be expected to give




FIGURE 4. Throughput comparison.




FIGURE 5. Loss rates in NS simulations.




FIGURE 6. Mean RTTs in NS simulations.



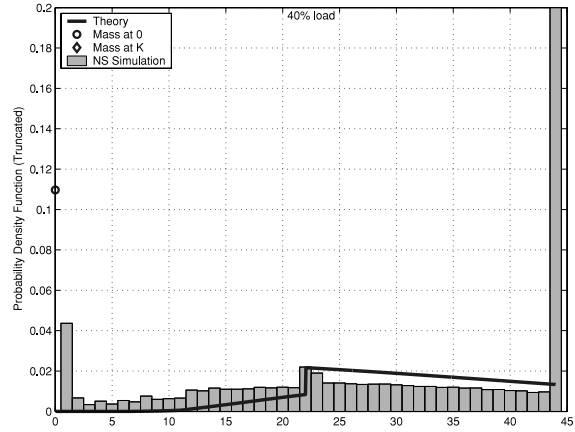
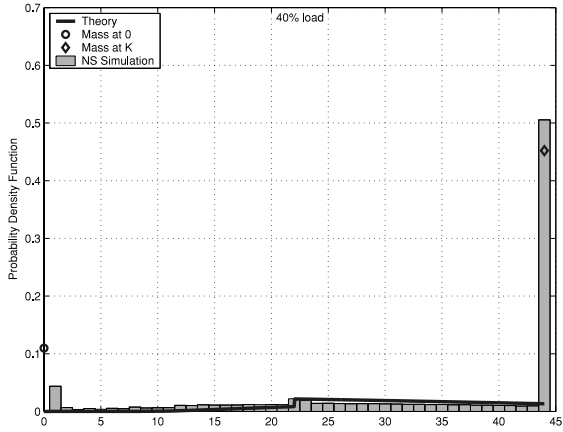
an excellent prediction of throughput for given probabilities of fast recovery and time-out events and a mean RTT.

Figures 7–9 compare the p.d.f. of effective TCP window size computed with the analytic model (“theory”) with histograms of window size observed in the NS simulations for different loads of background Web traffic. The plots on the left of each figure show a full range of probability density values on the y axis, whereas the plots on the right show more detail for densities in the range  $[0, 0.2]$ . All related plots are shown with a uniform scale for ease of comparison. For the analytic results, the computed probability mass at window sizes of zero and  $K$  are plotted with special symbols.

In general, the analytic model matches well with the distribution of effective window sizes from the simulation and also shows mass peaks at 0,  $K$ , and, at lower loads,  $K/2$  (resulting from halving the maximum window size). Note that the analytic model treats the minimum effective window size as zero and the simulated TCP correctly uses a minimum window size of one. The theoretical p.d.f. of window sizes differs from the simulation results mainly for small window sizes. A significant factor contributing to this, especially at the higher loads, is that the model approximates slow start after a time-out event by an instantaneous jump in the window size to the initial window for congestion avoidance and enters that phase. In the simulation, the window is advanced from one segment through some number of small values before reaching the congestion-avoidance phase.

## 9. SUMMARY AND CONCLUSIONS

We have described a new analytic model for the dynamic behavior of TCP windows as they respond to congestion indicated by data loss in the network. This model is based on stochastic differential equations and allows us to compute throughput as well as the complete p.d.f. for effective window sizes. The results computed from this model have been compared with NS simulations of TCP behavior in a realistic network environment and were found to produce comparable results.









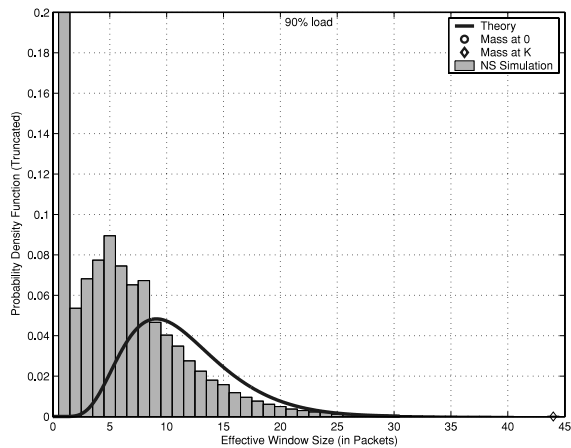



FIGURE 7. Comparison of effective window-size distributions at 40%, 50%, and 60% loads.





**FIGURE 8.** Comparison of effective window-size distributions at 70%, 80%, and 90% loads.

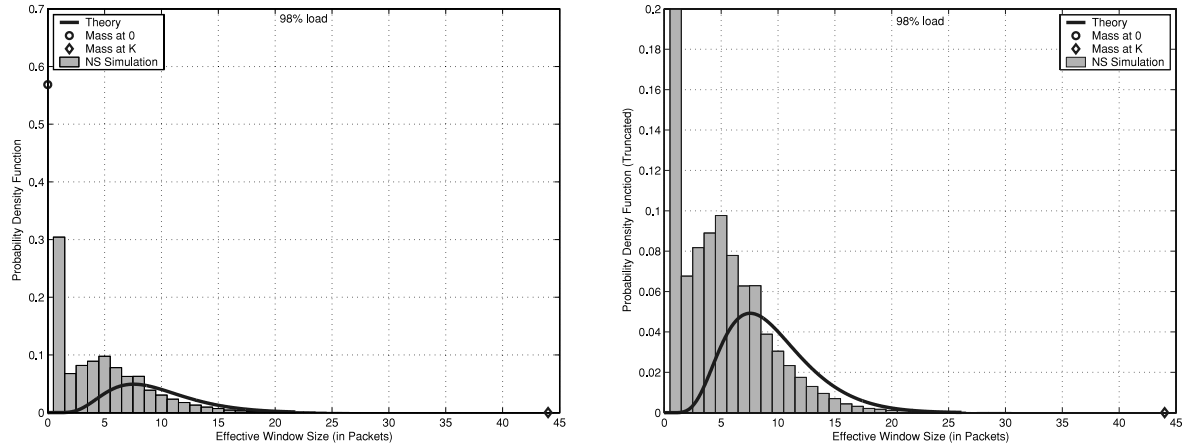


FIGURE 9. Comparison of effective window-size distributions at 98% loads.

## References

1. Abouzeid, A., Roy, S., & Azizoglu, M. (2000). Stochastic modeling of TCP over lossy links. *Proceedings of INFOCOM 2000*.
2. Abouzeid, A. Roy, S., & Azizoglu, M. (1999). Stochastic modeling of TCP over random loss channels. *Proceedings of the 6th International Conference on High Performance Computing*.
3. Altman, E., Avrachenkov, K., & Barakat, C. (2000). A stochastic model of TCP/IP with stationary random loss. *Proceedings of SIGCOMM 2000*.
4. Altman, E., Avrachenkov, K., & Barakat, C. (2000). Impact of bursty losses on TCP performance. *Performance Evaluation* 42(2–3): 129–147.
5. Altman, E., Avrachenkov, K., & Barakat, C. (2000). TCP in presence of bursty losses. *Proceedings of SIGMETRICS 2000*.
6. Altman, E., Avrachenkov, K., & Barakat, C. (2002). TCP network calculus: The case of large delay–bandwidth product. *Proceedings of INFOCOM 2002*.
7. Altman, E., Jimenez, T., & Nunez-Queija, R. (2002). Analysis of two competing TCP/IP connections. *Performance Evaluation* 49(1–4): 43–55.
8. Anjum, F. & Tassioulas, L. (1999). On the behavior of different TCP algorithms over a wireless channel with correlated packet losses. *Proceedings of SIGMETRICS 1999*.
9. Baccelli, F. & Hong, D. (2000). TCP is max-plus linear. *Proceedings of SIGCOMM 2000*.
10. Baccelli, F. & Hong, D. (2002). AIMD, fairness and fractal scaling of TCP traffic. *Proceedings of INFOCOM 2002*.
11. Barakat, C. (2001). TCP modeling and validation. *IEEE Network* 15(3): 38–47.
12. Breslau, L., Estrin, D., Fall, K., Floyd, S., Heidemann, J., Helmy, A., Huang, P., McCanne, S., Varadhan, K., Xu, Y., & Yu, H. (2000). Advances in network simulation. *IEEE Computer* 33(5): 59–67.
13. Cardwell, N., Savage, S., & Anderson, T. (2000). Modeling TCP latency. *Proceedings of INFOCOM 2000*.
14. Dumas, V., Guillemin, F., & Robert, P. (2001). Limit results for Markovian models of TCP. *Proceedings of GLOBECOM 2001*.
15. Dumas, V., Guillemin, F., & Robert, P. (2002). A Markovian analysis of additive-increase multiplicative-decrease (AIMD) algorithms. *Advances in Applied Probability* 34(1): 85–111.
16. Feldmann, A., Gilbert, A., Huang, P., & Willinger, W. (1999). Dynamics of IP traffic: A study of the role of variability and the impact of control. *Proceedings of SIGCOMM 1999*.
17. Floyd, S. (1991). Connections with multiple congested gateways in packet-switched networks Part 1: One-way traffic. *Computer Communication Review* 21(5): 30–47.
18. Floyd, S. & Fall, K. (1999). Promoting the use of end-to-end congestion control in the internet. *IEEE/ACM Transactions on Networking* 7(4): 458–472.
19. Floyd, S., Handley, M., Padhye, J., & Widmer, J. (2000). Equation-based congestion control for unicast applications. *Proceedings of SIGCOMM 2000*.
20. Floyd, S. & Jacobson, V. (1992). On traffic phase effects in packet-switched gateways. *Internetworking: Research and Experience* 3(3): 115–156.
21. Garetto, M., Cigno, R., Meo, M., & Marsan, M. (2001). A detailed and accurate closed queuing network model of many interacting TCP flows. *Proceedings of INFOCOM 2001*.
22. Goyal, M., Guérin, R., & Rajan, R. (2002). Predicting TCP throughput from non-invasive network sampling. *Proceedings of INFOCOM 2002*.
23. Hollot, C., Misra, V., Towsley, D., & Gong, W. (2001). A control theoretic analysis of RED. *Proceedings of INFOCOM 2001*.
24. Hollot, C., Misra, V., Towsley, D., & Gong, W. (2001). On designing improved controllers for AQM routers supporting TCP flows. *Proceedings of INFOCOM 2001*.
25. Kelly, F., Maulloo, A., & Tan, D. (1998). Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* 49(3): 237–252.
26. Kumar, A. (1998). Comparative performance analysis of versions of TCP in local network with a lossy link. *IEEE/ACM Transactions on Networking* 6(4): 485–498.

27. Lakshman, T.V. & Madhow, U. (1997). The performance of networks with high bandwidth-delay products and random loss. *IEEE/ACM Transactions on Networking* 5(3): 336–350.
28. Lakshman, T.V., Madhow, U., & Suter, B. (1997). Window-based error recovery and flow control with a slow acknowledgment channel: A study of TCP/IP performance. *Proceedings of INFOCOM 1997*.
29. Lakshman, T.V., Madhow, U., & Suter, B. (2000). TCP/IP performance with random loss and bi-directional congestion. *IEEE/ACM Transactions on Networking* 8(5): 541–555.
30. Low, S.H. (2000). A duality model of TCP and queue management algorithms. Extended version of paper presented at *Proceedings of ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*.
31. Low, S., Peterson, L., & Wang, L. (2001). Understanding TCP Vegas: A duality model. *Proceedings of SIGMETRICS 2001*.
32. Mathis, M., Semke, J., Mahdavi, J., & Ott, T. (1997). The macroscopic behavior of the TCP congestion avoidance algorithm. *Computer Communication Review* 27(3): 67–82.
33. Misra, A. & Ott, T. (1999). The window distribution of idealized TCP congestion avoidance with variable packet loss. *Proceedings of INFOCOM 1999*.
34. Misra, V., Gong, W., & Towsley, D. (1999). Stochastic differential equation modeling and analysis of TCP window size behavior. *Proceedings of Performance 1999*.
35. Misra, V., Gong, W., & Towsley, D. (2000). A fluid-based analysis of a network of AQM routers supporting TCP flows with an application to RED. *Proceedings of SIGCOMM 2000*.
36. Mitzenmacher, M. & Rajaraman, R. (2001). Towards more complete models of TCP latency and throughput. *Journal of Supercomputing* 20(2): 137–160.
37. Ott, T., Kemperman, J.H.B., & Mathis, M. (1996). The stationary behavior of ideal TCP congestion avoidance, unpublished manuscript.
38. Padhye, J., Firoiu, V., Towsley, D., & Kurose, J. (1998). Modeling TCP throughput: A simple model and its empirical validation. *Proceedings of SIGCOMM 1998*.
39. Padhye, J., Firoiu, V., Towsley, D., & Kurose, J. (2000). Modeling TCP Reno performance: A simple model and its empirical validation. *IEEE/ACM Transactions on Networking* 8(2): 133–145.
40. Roughan, M., Erramilli, A., & Veitch, D. (2001). Network performance for TCP networks. Part 1: Persistent sources. *Proceedings of Seventeenth International Teletraffic Congress*.
41. Sahu, S., Nain, P., Towsley, D., Diot, C., & Firoiu, V. (2000). On achievable service differentiation with token bucket marking for TCP. *Proceedings of SIGMETRICS 2000*.
42. Savari, S. & Telatar, E. (1999). The behavior of stochastic processes arising in window protocols. *Proceedings of the 1999 IEEE International Symposium on Information Theory*.
43. Schwefel, H. (2001). Behavior of TCP-like elastic traffic at a buffered bottleneck router. *Proceedings of INFOCOM 2001*.
44. Sikdar, B., Kalyanaraman, S., & Vastola, K.S. (2001). An integrated model for the latency and steady-state throughput of TCP connections. *Performance Evaluation* 46(2–3): 139–154.
45. Sikdar, B., Kalyanaraman, S., & Vastola, K.S. (2001). Analytic models and comparative study of the latency and steady-state throughput of TCP Tahoe, Reno and SACK. *Proceedings of GLOBECOM 2001*.
46. Willinger, W., Taqqu, N., Sherman, R., & Wilson, D. (1997). Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking* 5(1): 71–86.
47. Yang, Y. & Lam, S. (2000). General AIMD congestion control. *Proceedings of ICNP 2000*.
48. Yang, Y., Kim, M., & Lam, S. (2001). Transient behaviors of TCP-friendly congestion control protocols. *Proceedings of INFOCOM 2001*.
49. Zorzi, M., Chockalingam, A., & Rao, R. (2000). Performance analysis of TCP on channels with memory. *IEEE-JSAC* 18: 1289–1300.
50. URL: <http://www.cs.unc.edu/Research/dirt/proj/tcpmodel>