

TECHNICAL RESEARCH REPORT

TCP Traffic Modeling via Limit Theorems

by Peerapol Tinnakornsriruphap, Armand M. Makowski

CSHCN TR 2002-13
(ISR TR 2002-23)



The Center for Satellite and Hybrid Communication Networks is a NASA-sponsored Commercial Space Center also supported by the Department of Defense (DOD), industry, the State of Maryland, the University of Maryland and the Institute for Systems Research. This document is a technical report in the CSHCN series originating at the University of Maryland.

Web site <http://www.isr.umd.edu/CSHCN/>

TCP Traffic Modeling via Limit Theorems

Peerapol Tinnakornsriruphap *, Armand M. Makowski
Department of ECE and Institute for Systems Research,
University of Maryland, College Park, MD 20742

June 4, 2002

Abstract

Traditional TCP traffic modeling has focused on “micro-scale” modeling of TCP, i.e., detailed modeling of a single TCP flow. While micro-scale models of TCP are suitable for understanding the precise behavior of individual flows, they are not well suited to the situation when a large number of TCP flows interact with each other as is the case in realistic networks. In this survey, we present several works which focus on “macro-scale” modeling of TCP, where the aggregate behavior of TCP traffic can be simplified by applications of limit theorems.

Keywords: TCP, Traffic modeling, limit theorems, resource allocation

1 Introduction

One of the key mechanisms for the operation of the best-effort service Internet is the congestion control mechanism in TCP [1]. While there are several variations of the basic TCP congestion-control mechanism, they all have in common the *additive increase/multiplicative decrease* (AIMD) algorithm. This AIMD algorithm enables TCP congestion-control to be robust under diverse conditions. Unfortunately, the self-clocking feedback mechanism of the AIMD algorithm does introduce some additional complexity into the behavior of network traffic. There has been a number of studies in order to gain insights into this complex behavior. At the present, the relationship between the throughput of a single TCP and its round-trip and loss probability is fairly well understood [3] [4] [6] [7].

There are, however, certain aspects concerning TCP that are not well understood and which cannot be analyzed with the aforementioned models. This includes issues such as buffer behavior at a bottleneck router and the aggregate throughput when many TCP flows compete for the bandwidth of a link. While one could extend the micro-scale models in order to answer these questions, the resulting models are not scalable. More specifically, since each flow is modeled in great details, when the number of flows becomes large, the size of the state space for such models explodes and the analysis becomes intractable. Even numerical calculations or simulations of such models are very complicated, computationally prohibitive, and would not provide additional advantages over full-scale simulation with existing simulation packages (e.g., NS [8]). While one could make certain assumptions to simplify the analysis, it is not clear what are the irrelevant details that can be omitted while still providing reasonably accurate analysis.

To make matters worse, recent developments in Active Queue Management (AQM) techniques have introduced additional complexity in transport protocols. The development of AQM originated from the well-known fact that with simple Tail-Drop gateways, TCP congestion-control leads to undesirable behavior, i.e., global synchronization. When several TCP flows compete for bandwidth in a Tail-Drop gateway, it has been observed experimentally that packets from many flows are usually discarded simultaneously [9], resulting in a poor utilization of the network. AQM algorithms such as Random Early Detection (RED) [10] and Explicit Congestion Notification (ECN) [11] have been proposed to help alleviate this problem by randomly dropping/marking packets depending on queue size, thereby avoiding heavy congestion and preventing global synchronization. As can easily be imagined, the introduction of AQM further exacerbate the difficulty of understanding issues associated with buffer behavior and aggregate TCP traffic. There have been attempts to model the interactions between TCP and AQM, but so far the analytically tractable models are either too crude or too simplistic as we now discuss:

One class of models which has received much attention recently is the one where the TCP throughput emerges as the solution to a utility maximization problem [12] [13] [14]. With this class of model, TCP is shown to be “fair” in the sense that the throughput (or rate) of each TCP flow shares the bandwidth following a utility maximization problem. Moreover, the AQM mechanism can be interpreted as the feedback mechanism to communicate the information from the network to

*Corresponding author, phone: (301) 405-6844, fax: (301) 314-8586 email:{peerapol, armand}@isr.umd.edu

the users (TCP flows). While this model can provide some insights on the interaction between TCP flows and the network, there are certain features of the actual systems which are missing when TCP is viewed in this utility maximization context. Most glaring is the absence of the notion of “packet,” key to the operation of the congestion-control mechanism of TCP. Moreover, the numerical calculation of solutions is still very complex and computationally prohibitive when the number of TCP flows is large.

Other analytical efforts to model interactions between TCP congestion-control and the bottleneck router/buffer usually involve various ad-hoc assumptions to make the analysis feasible. For example, Hollot et al. linearized the TCP mechanism and studied the system using a control-theoretic approach [15]. The problem is that TCP is highly non-linear and the regime where such an approximation is accurate, is usually small. In [16], an analytical framework for multiple TCP flows sharing a RED gateway is developed under several potentially unrealistic assumptions. In [17], a simple analysis was carried out with TCP connections operating as Poisson processes under “slow” and “fast” rates.

The problem of modeling network traffic and of understanding its impact on various performance metrics is not unique to TCP networks. Early traffic modeling efforts in telephony revealed the suitability of Poisson processes to model the time patterns in the stream of call requests to a telephone exchange. This culminated with the pioneering work of A.K. Erlang who proposed the $M|M|c|c$ model for dimensioning call systems [18]. Similarly, it is well known that the multiplexed packet traffic generated by a large number of bursty data sources is well described by a Poisson process. It turns out in retrospect that this so-called “Poisson modeling” is easily justified with the help of limit theorems due to Palm [19] and Khintchine [20]. According to these results, under very weak technical conditions, the superposition of point processes is a point process which converges to a Poisson (point) process when the number of superposed point processes grows large.

These success stories point to the possibility of systematically applying limit theorems in order to derive traffic models. The advantages of doing so are three-fold. First, model simplification typically occurs when applying limit theorems, thereby filtering out irrelevant details without relying on ad-hoc assumptions. Second, since the knowledge and understanding of limit theorems are central to the Theory of Probability, it is reasonable to expect the existence of suitable limit theorems (with very weak assumptions) which can be applied to the traffic model of interest. Finally, the resource allocation problem is interesting only in networks operating at high utilization such as when the number of users is large. In such scenario, the limit behavior will be more accurate as the number of users increases.

Having taken this notion to heart, we survey various approaches to TCP modeling via limit theorems. The paper is organized as follows: Section 2 discusses some of the issues associated with micro-scale models of multiple TCP flows. In Section 3, we outline a general framework for macro-scale modeling of TCP traffic by limit theorems. Specific examples on the limiting models without packet-level operations are presented in Section 4. An outline of models with packet-level operations will be presented in Section 5 with two detailed examples in Section 6 (a rate-based model) and 7 (a window-based model). The Law of Large Numbers and the Central Limit Theorem of these models will be presented in Section 8 and 9, respectively. Section 10 discusses the application of the limiting results to the network dimensioning problem. Simulation results demonstrating the limit theorems are shown in Section 11. Section 12 discusses other limit theorems on TCP. Section 13 concludes the paper and discusses further challenges.

A word on the notation in use: Equivalence in law or in distribution between random variables (rvs) is denoted by $=_{st}$. The indicator function of an event A is simply $\mathbf{1}[A]$. We use \xrightarrow{P}_n and \implies_n to denote convergence in probability and convergence in distribution (with n going to infinity), respectively. Unless specified otherwise, $X_i^{(N)}(t)$ represents the quantity X of flow i (out of N flows) at time t .

2 Micro-scale TCP Traffic modeling

In this section, we review some of the work on micro-scale TCP traffic modeling as well as the efforts made to extend these micro-scale models to the situation with many interacting flows. The resulting models suffer from either one of the following difficulties: (i) These models are typically too complex and therefore not tractable analytically. Sometimes this can also translate in numerical calculations being computationally prohibitive; (ii) Reduction in complexity can be achieved by making simplifying assumptions which are often ad-hoc, difficult to justify, and sometimes downright unrealistic. We argue that these assumptions should emerge in a natural way from the traffic modeling process, rather than being enforced for the main reason that it enables the analysis to go through. It will soon become apparent that a modeling methodology based on limit theorems provides a direct remedy to this issue – Indeed, model simplifications occur in the limit, without the need for any ad-hoc assumptions.

2.1 TCP modeling by Markov Process

One of the earliest and most popular efforts to model TCP made use of Markov chain modeling. The size of the TCP congestion window acts as the state of the chain and the loss probability (either independent or dependent on the state)

determines the transition probabilities. Padhye et al. have derived an approximation for the steady-state distribution of a Markov chain modeling a single TCP connection with a fixed loss probability [3]. Altman et al. [6] have extended the model to cover more general loss processes, e.g., continuous-time Markov chain with different loss probability in each state.

2.2 Extension of micro-scale TCP model to many TCP flows

In this section, we point to various efforts for extending the micro-scale TCP model to many TCP flows. In the models to be described, there is no attempt to consider a limiting model.

Hasegawa et al. [21] consider a Markov chain for N TCP connections using either Tail-Drop or RED gateway. The state space in this system is the vector of window size of all of N connections and the queue size, which is a function of the sum of the size of congestion windows. For each connection, the transition probability depends on the current window size and queue size, hence the transition probability depends on the the window size of *all* connections. The authors then derive a fixed-point solution of the average window size in steady-state. However, this fixed-point solution is very complicated and requires solving a large system of non-linear equations. It is not clear how this could be accomplished effectively and whether it would offer any analytical understanding to the problem.

Another example is the work by Garetto et al. [22] where a closed-queueing network model is proposed to investigate the interaction between many TCP flows. A closed network of $M|M|\infty$ queues is introduced where each queue represents a state of TCP algorithm. The number of users in each queue represents the number of TCP connections in that state. The service rate of the user in each queue depends on the state of TCP. After completing service in a queue, each user's transition to the next queue depends on the loss probability. This closed-queueing network is used in a two-tier model which can be described as follows: Given a loss probability, a numerical calculation of the steady-state distribution of the closed-queueing network yields an approximated traffic load. This load will be input into a pseudo-network model (e.g., $M|M|1|B$) to determine the new loss probability for the closed-queueing network. If the result from the successive iterations converges, it is claimed that the convergence will be to the same working point as given by the model where each TCP flow is modeled in details.

Several questions arise concerning such a model. First, the model implicitly assumes that the dynamics of the TCP congestion control mechanism converges to steady state faster than the network dynamics; this is the opposite of the real Internet where TCP reacts to events in the network. Next, while the numerical calculations are not as complex as would be the case in a detailed model of TCP flows, they are far from simple. The example considered in [22] consists of 11 different queue types and the number of queues (M_q) is 357. Calculations for the steady-state distribution of the queue are typically of the order $O(M_q^2)$ for each iteration. Finally, the "interaction" between TCP flows is developed in an abstract model where its deviation from the actual interaction process cannot be quantified. It also impossible to draw any analytical conclusion concerning these interactions.

Baccelli et al. [23] propose fixed-point methods for the simulation of the sharing of a local loop by a large number of interacting homogeneous TCP connections. The analysis uses a detailed description of one TCP connection and a simplified description of the interaction with other connections. It is again difficult to quantify the effect of such simplifications and the accuracy of the results can only be verified by comparing them to results of extensive simulations.

2.3 Fluid approximation of TCP congestion-control

Mathis et al. [4] use a continuous-time fluid flow approximation to the discrete time process of window behavior. Assume the congestion signals TCP to back off according to a Poisson process where the k^{th} signal occurs at time epoch τ_k ($k = 1, 2, \dots$). Then, the approximate evolution of the congestion window evolution of a single TCP flow is governed by

$$\frac{dW}{dt} = \frac{1}{W(t)}, \quad (1)$$

except around the points τ_k ($k = 1, 2, \dots$) where

$$W(\tau_k^+) = W(\tau_k^-)/2. \quad (2)$$

These equations can be used to derive the fact that that the average window size is of the order of $1/\sqrt{p}$. This model is suitable only for a single flow because the congestion notification is assumed to be Poisson and each notification is independent of each other (an assumption similar to [7] which uses stochastic differential equation to model TCP). When more than one flow utilizes the same bottleneck link, this assumption is not helpful for capturing the interaction between flows. Bonald [24] considers a similar model for several TCP flows with a major difference that congestion occurs when the sum of congestion windows exceeds the bandwidth-delay product plus the bottleneck buffer. Under the assumption

that at every congestion epoch all TCP flows simultaneously back off, TCP can be shown to be fair and an explicit closed-form formula for the utilization can be derived. However, the assumed total synchronization between flows assumption is unrealistic.

One approach which has received much attention recently is to model TCP throughput as the solution to a utility maximization problem. The interest in this model originates from the work of Kelly [12] who showed that the utility maximization problem of the system composing of a network and users can be decomposed into two separate problems, namely the network problem and the users problem, assuming the utility function of each user is a concave function of the user’s received throughput. If the network maximizes its revenue and then users subsequently maximize their utilities, the recursive maximization sequences will converge to the solution of the system utility maximization problem. Subsequent work by Kelly et al. [13] shows that the user problem can be solved by a rate control algorithm which can be implemented as a congestion-control algorithm in TCP. Recent work by Low [14] shows that the throughput of AIMD TCP congestion-control algorithm solves the user problem with a certain utility function. This enables efficient modeling of a large network with multiple users and makes it possible to observe the interactions between the network and traffic flows. Furthermore, AQM mechanisms such as RED can be modeled as the feedback functions from the network to the users.

While there are many advantages of viewing the system as a utility maximization problem and of modeling TCP traffic as rate-controlled fluid flows, there are definite drawbacks as well. Most important of all is the absence of “packets” from the system model, since the congestion-control mechanism of TCP relies on packet-level operations. For example, it is not possible under this model to derive the queue length distribution at routers which is an important question for any successful network dimensioning. Additionally, while the solution to the maximization problem might accurately describe the steady-state solution of TCP, the distributed solution does not capture the short-term dynamics of TCP well because of the absence of the window mechanism. Finally, the numerical calculation of the solution is still very complex as it suffers from state space explosion when the number of TCP flows becomes large. Therefore, this type of models appears more suitable for understanding the “big picture” and the qualitative behavior of congestion-control algorithms, rather than for effective and accurate network dimensioning.

3 Limiting behavior of TCP with a large number of flows

In this section, we outline the type of results that can be expected from modeling TCP with a large number of homogeneous TCP flows via limit theorems. We also describe common properties of the models.

There are two common characteristics that can be expected when taking the limit as the number of users to infinity. First, a properly normalized aggregated quantity (e.g., queue size and throughput) converges and a simplified recursive dynamics emerges for the limiting quantity. Second, the effect of any single flow on other flows becomes less pronounced as the number of flows grows large due to the homogeneity of the flows. Therefore, it is reasonable to expect the throughput of a flow to be asymptotically independent of other flows.

Although there are different ways to apply a limiting process to a TCP model, they have some necessary characteristics in common: First of all, only the congestion in the bottleneck link is considered with capacity of the bottleneck link and the feedback information (if AQM is utilized) scaling with the number of users. For example, if C denotes the capacity per user of at the bottleneck router, then its overall capacity should be scaled up to NC when N TCP flows share the bottleneck link.

Next, in order to describe the interaction between the bottleneck router and TCP flows, all models use a recursion similar to Lindley’s recursion for some quantity of interest (e.g., queue length or transmission rate); let $\zeta(t)$ denote the value of this quantity at the discrete time epoch t . During the timeslot $[t, t + 1)$, it is possible for the quantity to experience an increase (e.g., the queue length increases due to new incoming packets) or a reduction (e.g., the packets in the queue are serviced/transmitted resulting in smaller queue size). We denote the amount of increase and decrease as $\alpha(t)$ and $\beta(t)$, respectively. As the quantity of interest in the network is typically non-negative, we can write the following recursion

$$\zeta(t + 1) = [\zeta(t) + \alpha(t) - \beta(t)]^+. \tag{3}$$

In this paper, we classify the limiting models into two classes. The first class is the models that ignore the packet-level operations of TCP, and consider only the evolution of the transmission rate, while models in the second class take into account detailed packet-level operations. It is important to make this distinction because the actual operations of TCP, of RED and of the network rely on the packet-level operations. The omission of this level of details simplifies the analysis but at the expense of additional distortion in the results.

4 Models without packet-level operations

Since all of the models in this section consider only the transmission rate of the TCP flows, we use a common notation $x_i^{(N)}(t)$ to denote the transmission rate of flow i ($i = 1, \dots, N$) at time t .

4.1 Shakkottai and Srikant [25]

Consider a discrete-time model of N homogeneous proportional fair congestion-control flows (such as the primal algorithm in [13]) utilizing a bottleneck router with either rate-based marking (such as Virtual Queue marking) or queue-based marking (such as RED), i.e.,

$$x_i^{(N)}(t+1) = \left(x_i^{(N)}(t) + \Delta - \beta x_i^{(N)}(t-d) f^{(N)}(x_i^{(N)}(t-d) + e_i^{(N)}(t-d)) \right)^+,$$

where Δ and β are positive constants which determine the rate at which a flow increases and decreases its transmission rate, f is the marking function, d is the round-trip delay between the flow and the bottleneck router, and $e_i^{(N)}$ is a “noise” process, representing short-lived and uncontrolled flows. A natural way to obtain a simplified limiting model resorts to rescaling the length of timeslots to be inversely proportional to the number of flows N . Then, as N gets large, the average transmission rate is expected to converge almost surely (under the appropriate assumptions on the “noise” process and the function f) to a deterministic quantity. This quantity can be described by a functional differential equation, thereby justifying deterministic fluid approximations of the algorithms described in Section 2.3. However, the aforementioned approach does not apply well to the Internet with TCP congestion-control because (i) the model ignores the complexity of the window-based implementation and consider only the transmission rate of the flow and (ii) the proportional fair congestion-control algorithm is additive-increase/additive-decrease based (when packetized), whence continuous in its fluid limit and thus “nicer” than TCP congestion-control which may experience abrupt transmission rate changes. It is still an open technical problem on whether there exists a fluid limit for AIMD TCP congestion-control.

4.2 Hong and Lebedev [26]

Consider next the model where the throughput of each connection evolves at time epochs where congestion occurs. At every epoch, each connection draws a $\{0.5, 1\}$ random number representing the fraction of throughput that the connection retains after congestion. This fraction depends on the throughput of the flow just before congestion occurred. If T_n is the n^{th} congestion epoch and $\gamma_{i,n}(x)$ is a $\{0.5, 1\}$ -valued rv depending on the rate x , then

$$x_i^{(N)}(T_n^+) = \gamma_{i,n} \left(x_i^{(N)}(T_n^-) \right) \cdot x_i^{(N)}(T_n^-)$$

with obvious notation. Let NC be the capacity of the bottleneck router, and assume that the residual capacity at time T_n^+ , namely $NC - \sum_{i=1}^N x_i^{(N)}(T_n^+)$, is divided evenly among all users. Then,¹

$$x_i^{(N)}(T_{n+1}^-) = x_i^{(N)}(T_n^+) + C - \frac{1}{N} \sum_{i=1}^N x_i^{(N)}(T_n^+)$$

The asymptotic behavior of the average rate at each epoch can be described by the following theorem.

Theorem 1 Denote $x_{i,n}^{(N)}$ the transmission rate of connection i right after the n^{th} -epoch. Let $g(y) = \mathbf{P}[\gamma_{i,n}(y) = 0.5]$. If g is continuous and non-decreasing with $\inf_{y \in \mathbf{R}^+} g(y) > 0$ and $\sup_{y \in \mathbf{R}^+} g(y) < 1$, then

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N x_{i,n}^{(N)} = s_n \quad \text{in } L_1 \text{ and a.s.}$$

where $s_n = \mathbf{E} \left[x_n^{(\infty)} \right] = \lim_{N \rightarrow \infty} \mathbf{E} \left[x_{i,n}^{(N)} \right]$. Moreover, the sequence $\{x_n^{(\infty)}, n = 0, 1, \dots\}$ is stationary and ergodic, and can be characterized by the relation

$$x_{n+1}^{(\infty)} \stackrel{st}{=} \gamma_{n+1} \left(x_n^{(\infty)} + C - \mathbf{E} \left[x_n^{(\infty)} \right] \right) \cdot \left(x_n^{(\infty)} + C - \mathbf{E} \left[x_n^{(\infty)} \right] \right),$$

where $\gamma_{n+1}(x)$ is a rv which is identically distributed to $\gamma_{i,n+1}(x)$.

This result suggests the existence of a simpler process for describing the asymptotic behavior of the average rate. Although the model unrealistically assumes that the multiplicative decrease in a TCP flow depends only on its own transmission rate and not on that of other flows, we will show in the next section that similar results also exist in more elaborate models with many TCP flows sharing a RED gateway.

¹Although the authors do not limit the transmission rate to be only non-negative, the extension of the model to incorporate such restriction should not change the nature of the result.

4.3 Adjih et. al [27]

The authors consider a continuous-time model where partial differential equations of the free buffer space in a Tail-Drop gateway and the density of the window size distribution are specified. Under the assumption that the buffer in the bottleneck router scales with the number of users, some asymptotic results on both the free buffer space and the window distribution are established by the use of mean-field approximations.

5 Models with detailed packet-level operations

Tinnakornsrisuphap and Makowski have recently introduced two discrete-time models, referred hereafter as the rate-based model [28] and the window-based model [29], respectively, when a large number of TCP flows compete for bandwidth resources through a RED gateway.

These two models are conceptually similar, and are also organized around the behavior of the RED buffer contents observed at discrete epochs (which are indexed by $t = 0, 1, \dots$). They differ through the specific mechanism used for generating incoming packets in response to the congestion that develops in the RED buffer shared by the TCP flows. Specifically, in the rate-based model, the packet transmission rate of a source increases/decreases in response to random packet dropping at the RED gateway with a probability which depends on the queue size. In the window-based model, the AIMD algorithm controls the window size in response to the bottleneck ECN/RED gateway marking incoming packets with a probability which depends on the queue size.

Fix $N = 1, 2, \dots$, and consider the situation where N TCP flows are active. For simplicity time is assumed slotted. Let $Q^{(N)}(t)$ denote the number of packets in the buffer at the beginning of timeslot $[t, t + 1)$. In that timeslot, source i generates $A_i^{(N)}(t + 1)$ packets according to a mechanism to be specified shortly; these incoming packets are put in the buffer at the beginning of timeslot $[t, t + 1)$. Let the rv

$$A^{(N)}(t + 1) := \sum_{i=1}^N A_i^{(N)}(t + 1) \quad (4)$$

represent the aggregate number of packets offered to the network by the N sources at the beginning of timeslot $[t, t + 1)$, so that $Q^{(N)}(t) + A^{(N)}(t + 1)$ packets are available for transmission during that timeslot. Since the outgoing link operates at the rate of NC packets/timeslot, $[Q^{(N)}(t) + A^{(N)}(t + 1) - NC]^+$ packets will not be transmitted during timeslot $[t, t + 1)$, and will remain in the buffer, their transmission being deferred to subsequent timeslots. The number $Q^{(N)}(t + 1)$ of packets in the buffer at the beginning of timeslot $[t + 1, t + 2)$ is therefore given by

$$Q^{(N)}(t + 1) = [Q^{(N)}(t) - NC + A^{(N)}(t + 1)]^+ \quad (5)$$

In order to fully specify the model, we need to specify the *joint* statistics of the rvs $\{A_i^{(N)}(t + 1), i = 1, \dots, N; t = 0, 1, \dots\}$. This will be done in some details in Sections 6 and 7 for each of the models mentioned earlier. Throughout, let $f^{(N)} : \mathbb{R}_+ \rightarrow [0, 1]$ denote the *marking/dropping probability* function of the RED gateway. Moreover, we find it convenient to use the collection of i.i.d. $[0, 1]$ -uniform rvs $\{U_i(t + 1), V_i(t + 1), V_{i,j}(t + 1), i, j = 1, \dots; t = 0, 1, \dots\}$ which are assumed independent of the rvs $Q^{(N)}(0)$ and other initial conditions.

6 A rate-based model for TCP+RED

According to the rate-based model, a source either transmits or is idle in a given timeslot. So, let $B_i^{(N)}(t + 1)$ be a $\{0, 1\}$ -valued rv that encodes the number of packets generated by source i . Moreover, let $R_i^{(N)}(t + 1)$ represent the possibility that the packet generated by source i at the beginning of timeslot $[t, t + 1)$ is rejected, i.e., $R_i^{(N)}(t + 1) = 1$ (resp. $R_i^{(N)}(t + 1) = 0$) if the packet is rejected by (resp. accepted into) the RED buffer. Set

$$B_i^{(N)}(t + 1) = \mathbf{1} [U_i(t + 1) \leq \alpha_i^{(N)}(t)] \quad (6)$$

where $\alpha_i^{(N)}(t)$ is an $[0, 1]$ -valued rv which denotes the (conditional) *transmission rate* of traffic source i at the beginning of timeslot $[t, t + 1)$, and let

$$R_i^{(N)}(t + 1) = \mathbf{1} [V_i(t + 1) \leq f^{(N)}(Q^{(N)}(t))] \quad (7)$$

denote the indicator function of the event that the incoming packet from source i will be rejected. Thus,

$$A_i^{(N)}(t+1) = (1 - R_i^{(N)}(t+1))B_i^{(N)}(t+1). \quad (8)$$

To select the transmission rates we argue as follows: Suppose that source i generates no packet during timeslot $[t, t+1)$ ($B_i^{(N)}(t+1) = 0$), then the transmission rate of source i in the next timeslot remains unchanged. If on the other hand, a packet is produced by source i at the beginning of timeslot $[t, t+1)$, then either the packet is successfully transmitted ($R_i^{(N)}(t+1) = 0$), or it is dropped ($R_i^{(N)}(t+1) = 1$). In the former case, the transmission rate of source i in the next timeslot is *increased* to $\alpha_i^{(N)}(t)^{1-\varepsilon}$ ($0 < \varepsilon < 1$), while this transmission rate is *decreased* by a factor γ ($0 < \gamma < 1$) to $\gamma\alpha_i^{(N)}(t)$ in the latter case.

Under the constraint that transmission rates are bounded to the unit interval, these rules attempt to *emulate* the additive increase and multiplicative decrease, respectively, of the TCP congestion-control by conservatively increasing the transmission rate if the transmission is successful and reducing the transmission rate by the factor γ when the event of a packet loss. This can be summarized into the single equation

$$\begin{aligned} & \alpha_i^{(N)}(t+1) \\ = & \alpha_i^{(N)}(t)^{1-\varepsilon}(1 - R_i^{(N)}(t+1))B_i^{(N)}(t+1) + \gamma\alpha_i^{(N)}(t)R_i^{(N)}(t+1)B_i^{(N)}(t+1) + \alpha_i^{(N)}(t)(1 - B_i^{(N)}(t+1)) \end{aligned} \quad (9)$$

7 A window-based model for TCP+RED+ECN

Assume that each connection has an infinite amount of data to transmit and that in each timeslot it transmits as much as is allowed by its congestion window in that timeslot. So, for $i = 1, \dots, N$, let $W_i^{(N)}(t)$ be an integer-valued rv that encodes the number of packets generated by source i (and hence its congestion window) at the beginning of timeslot $[t, t+1)$, so that

$$A_i^{(N)}(t+1) = W_i^{(N)}(t). \quad (10)$$

We assume the integer $W_i^{(N)}(t)$ to be in the range $\{1, \dots, W_{\max}\}$ for some finite integer W_{\max} .

The AIMD algorithm controls the window size and the bottleneck ECN/RED gateway *marks* the incoming packets with a probability which depends on the queue size. We represent this possibility by the $\{0, 1\}$ -valued rv $M_{i,j}^{(N)}(t+1)$ ($j = 1, \dots, W_i^{(N)}(t)$) with the interpretation that $M_{i,j}^{(N)}(t+1) = 0$ (resp. $M_{i,j}^{(N)}(t+1) = 1$) if the j th packet from source i is marked (resp. not marked) in the RED buffer. The process by which packets are marked is described first: For each $i = 1, \dots, N$ and $j = 1, 2, \dots$, we define the marking rvs

$$M_{i,j}^{(N)}(t+1) = \mathbf{1} \left[V_{i,j}(t+1) > f^{(N)}(Q^{(N)}(t)) \right],$$

so that the rv $M_{i,j}^{(N)}(t+1)$ is the indicator function of the event that the j th packet from source i is *not* marked in timeslot $[t, t+1)$.

Thus, in a round-trip, each packet coming into the router is marked/dropped with identical (conditional) probability which depends only on the queue length at the beginning of the timeslot. This model approximates the case where the memory of the queue averaging mechanism is long, which is the case for the recommended parameter settings of RED [30]. Next we introduce the rvs

$$M_i^{(N)}(t+1) = \prod_{j=1}^{W_i^{(N)}(t)} M_{i,j}^{(N)}(t+1), \quad (11)$$

so that $M_i^{(N)}(t+1) = 1$ (resp. $M_i^{(N)}(t+1) = 0$) corresponds to the event that no packet (resp. at least one packet) from source i has been marked in timeslot $[t, t+1)$. The evolution of the window mechanism for source i can now be described through the recursion

$$\begin{aligned} & W_i^{(N)}(t+1) \\ = & \min \left(W_i^{(N)}(t) + 1, W_{\max} \right) M_i^{(N)}(t+1) + \min \left(\left\lceil \frac{W_i^{(N)}(t)}{2} \right\rceil, W_{\max} \right) (1 - M_i^{(N)}(t+1)). \end{aligned} \quad (12)$$

This equation emulates the interaction between TCP and RED as follows: If no packet from source i is marked in timeslot $[t, t+1)$, then the congestion window size in the next timeslot is increased by 1. On the other hand, if one or more packets

are marked in timeslot $[t, t + 1)$, then the congestion window in the next timeslot is reduced by half. The size of the congestion window is limited by the maximum window size W_{\max} ².

8 A Law of Large Numbers

We are interested in determining the limiting behavior of the rate-based and window-based models as the number N of TCP sources becomes large. The discussion is carried out under the following assumptions (A1)-(A3):

(A1) There exists a continuous function $f : \mathbb{R}_+ \rightarrow [0, 1]$ such that for each $N = 1, 2, \dots$,

$$f^{(N)}(x) = f(N^{-1}x), \quad x \geq 0;$$

(A2) For each $N = 1, 2, \dots$, the queue dynamics start with the conditions

$$Q^{(N)}(0) = 0;$$

(A3) For some non-random α in $(0, 1]$ and some W in $\{1, \dots, W_{\max}\}$,

$$\alpha_i^{(N)}(0) = \alpha \quad \text{and} \quad W_i^{(N)}(0) = W$$

for each $i = 1, \dots, N$.

Assumption (A1) is a structural conditions while (A2)-(A3) are made essentially for technical convenience as it implies that for each $N = 1, 2, \dots$ and all $t = 0, 1, \dots$, the rvs $\alpha_1^{(N)}(t), \dots, \alpha_N^{(N)}(t)$ (resp. $W_1^{(N)}(t), \dots, W_N^{(N)}(t)$) are *exchangeable*. Assumptions (A2)-(A3) can be omitted but at the expense of a more cumbersome discussion.

Theorem 2 *Assume (A1)-(A3) in the rate-based model of Section 6. Then, for each $t = 0, 1, \dots$, there exist a (non-random) constant $q(t)$ and a $[0, 1]$ -valued rv $\alpha(t)$ such that the following holds:*

(i) *The convergence*

$$\frac{Q^{(N)}(t)}{N} \xrightarrow{P} {}_N q(t) \quad \text{and} \quad \alpha_1^{(N)}(t) \xrightarrow{P} {}_N \alpha(t)$$

takes place;

(ii) *For any integer $I = 1, 2, \dots$, the rvs $\{\alpha_i^{(N)}(t), i = 1, \dots, I\}$ become asymptotically independent as N becomes large, with*

$$\lim_{N \rightarrow \infty} \mathbf{P} \left[\alpha_i^{(N)}(t) \leq x_i, i = 1, \dots, I \right] = \prod_{i=1}^I \mathbf{P} [\alpha(t) \leq x_i]$$

for any x_1, \dots, x_I in $[0, 1]$.

Moreover, with initial conditions $q(0) = 0$ and $\alpha(0) = \alpha$, it holds that

$$q(t+1) = [q(t) - C + (1 - f(q(t)))\mathbf{E}[\alpha(t)]]^+$$

and

$$\alpha(t+1) \underset{st}{=} \alpha(t)^{1-\varepsilon} (1 - R(t+1))B(t+1) + \gamma\alpha(t)R(t+1)B(t+1) + \alpha(t)(1 - B(t+1))$$

where $B(t+1) = \mathbf{1}[U(t+1) \leq \alpha(t)]$ and $R(t+1) = \mathbf{1}[V(t+1) \leq f(q(t))]$ for i.i.d. $[0, 1]$ -uniform rvs $\{U(t+1), V(t+1), t = 0, 1, \dots\}$.

Theorem 3 *Assume (A1)-(A3) in the window-based model of Section 7. Then, for each $t = 0, 1, \dots$, there exist a (non-random) constant $q(t)$ and an $\{1, \dots, W_{\max}\}$ -valued rv $W(t)$ such that the following holds:*

(i) *The convergence*

$$\frac{Q^{(N)}(t)}{N} \xrightarrow{P} {}_N q(t) \quad \text{and} \quad W_1^{(N)}(t) \Rightarrow_N W(t)$$

²If $W_i^{(N)}(0)$ lies in the range $\{1, \dots, W_{\max}\}$ for each $i = 1, \dots, N$, then so does $W_i^{(N)}(t)$ for each $t = 0, 1, \dots$ and the minimum with W_{\max} in the second term of (13) can be omitted.

takes place;

(ii) For any integer $I = 1, 2, \dots$, the rvs $\{W_i^{(N)}(t), i = 1, \dots, I\}$ become asymptotically independent as N becomes large, with

$$\lim_{N \rightarrow \infty} \mathbf{P} \left[W_i^{(N)}(t) = k_i, i = 1, \dots, I \right] = \prod_{i=1}^I \mathbf{P} [W(t) = k_i]$$

for any k_1, \dots, k_I in \mathbf{N} .

Moreover, with initial conditions $q(0) = 0$ and $W(0) = W$, it holds that

$$q(t+1) = [q(t) - C + \mathbf{E} [W(t)]]^+ \quad (13)$$

and

$$W(t+1) =_{st} \min(W(t) + 1, W_{\max}) M(t+1) + \min\left(\left\lceil \frac{W(t)}{2} \right\rceil, W_{\max}\right) (1 - M(t+1))$$

where

$$M(t+1) = \mathbf{1} \left[V(t+1) \leq (1 - f(q(t)))^{W(t)} \right]$$

for i.i.d. $[0, 1]$ -uniform rvs $\{V(t+1), t = 0, 1, \dots\}$.

Theorems 2 and 3 suggest that a bottleneck queue driven by a random marking/dropping algorithm, under a large number of TCP sources, can be characterized by a two-dimensional recursion for the evolution of the limiting normalized queue length $q(t)$ and the limiting transmission rate $\alpha(t)$ (resp. window size $W(t)$). The convergence result $\frac{Q^{(N)}(t)}{N} \xrightarrow{P} Nq(t)$ is a byproduct of the convergence

$$\frac{A^{(N)}(t+1)}{N} = \frac{\sum_{i=1}^N A_i^{(N)}(t+1)}{N} \xrightarrow{P} Na(t) \quad (14)$$

for some constant $a(t)$ determined by the model. Specifically, for the rate-based model we have $a(t) = (1 - f(q(t)))\mathbf{E} [\alpha(t)]$, while for the window-based model, $a(t) = \mathbf{E} [W(t)]$.

This result, while similar to the Weak Law of Large Numbers, cannot be obtained by a straightforward application of the classical Law of Large Numbers. Indeed, the summands in (4) (under (A2)-(A3)) are *identically* distributed and *correlated* rvs whose common distribution *varies* with N . However, as the number of sources increases, the dependency between any pair of sources becomes weaker so that the aggregate behavior eventually becomes deterministic. This provides some indication as to the possibility that RED might indeed be more effective in breaking global synchronization between TCP flows as the number of flows become larger.

One advantage of modeling a RED gateway over a Tail-Drop gateway is that in RED gateway, there is a fixed orderly structure on how the packets are being marked/dropped depending on the probability function. In a Tail-Drop gateway, it is more difficult to accurately model how incoming packets are dropped, for such events usually depend on the precise timing of the packets arrival and departure. However, we conjecture as long as such a mechanism does not favor any single flow then a similar convergence should also exist.

9 CLT complements

The approximation implied by Theorems 2 and 3 can be improved with the help of a Central Limit Theorem-type result. More specifically, since $Q^{(N)}(t)/N \xrightarrow{P} Nq(t)$, it is natural to investigate the distribution of $\sqrt{N}(Q^{(N)}(t)/N - q(t))$ as N becomes large. The following theorem presents such a result for the window-based model in the setup of Theorem 3 but with Assumption (A1) strengthened to read as Assumption (A1bis), where

(A1bis) There exists a twice continuously differentiable function $f : \mathbf{R}_+ \rightarrow [0, 1]$ such that for each $N = 1, 2, \dots$,

$$f^{(N)}(x) = f(N^{-1}x), \quad x \geq 0.$$

For each $k = 1, \dots, W_{\max}$, we now define

$$L_k^{(N)}(t) = \frac{1}{N} \sum_{i=1}^N \mathbf{1} [W_i^{(N)}(t) = k] - \mathbf{P} [W(t) = k];$$

$$L_0^{(N)}(t) = \frac{Q^{(N)}(t)}{N} - q(t).$$

Theorem 4 Assume (A1bis)-(A3). Then, for each $t = 0, 1, \dots$, there exists an $\mathbb{R}^{W_{\max}+1}$ -valued rv $(L_0(t), L_1(t), \dots, L_{W_{\max}}(t))$ such that

$$\sqrt{N} \left(L_0^{(N)}(t), L_1^{(N)}(t), \dots, L_{W_{\max}}^{(N)}(t) \right) \Rightarrow_N (L_0(t), L_1(t), \dots, L_{W_{\max}}(t)). \quad (15)$$

Moreover, the distributional recurrence

$$L_0(t+1) =_{st} \begin{cases} 0 & K(t) > 0 \\ L_0(t) + \bar{L}(t) & K(t) < 0 \\ (L_0(t) + \bar{L}(t))^+ & K(t) = 0 \end{cases}$$

holds with $L_0(0) = 0$, where we have set

$$\bar{L}(t) := \sum_{k=1}^{W_{\max}} k \cdot L_k(t) \text{ and } K(t) = C - q(t) - \mathbf{E}[W(t)].$$

Theorem 4 is established in [29], and a similar result can be established for the rate-based model. The convergence (15) suggests the approximation

$$Q^{(N)}(t) \simeq Nq(t) + \sqrt{N}L_0(t). \quad (16)$$

We can interpret $K(t)$ as the residual capacity per user in the limit in timeslot $[t, t+1)$. If there exists extra capacity for the average user rate to increase ($K(t) > 0$), then there is no fluctuation in the limiting queue. On the other hand, when there is congestion ($K(t) < 0$), the non-trivial limiting distribution can be found. Some technical difficulties arise in the special case where $K(t) = 0$.

As part of the proof of Theorem 4 we find that

$$L_k^{(N)}(t+1) \Rightarrow_N c_k(t) f'(q(t)) L_0(t) + \xi(t)$$

for each $k = 1, \dots, W_{\max}$ for some Gaussian term $\xi(t)$ and non-random constant $c_k(t)$. Thus, the magnitude of the queue fluctuation is proportional to the derivative of f around the limiting (normalized) queue size $q(t)$. Therefore, it is safer to utilize a smooth marking probability function as suggested in the ‘‘gentle’’ option as opposed to the original recommendation in [10] and [30]. Some oscillatory behavior with RED was reported when the average packet drop rate exceeds max_p in the absence of RED’s ‘‘gentle’’ modification [31]; these observations are certainly compatible with the analytical CLT results presented here.

10 Application to network dimensioning

We briefly discuss how to apply the convergence results of Section 8 and 9 to the network dimensioning problem. In [29], it is shown for the window-based model that if the limiting queue $q(t)$ converges to a constant $q > 0$ in steady-state (i.e., t going to infinity), then $f(q)$ is the marking probability in steady state. If W^* denotes the rv representing the limiting window size in steady state, then it turns out [29] that

$$C = \mathbf{E}[W^*] = \mathbf{E}[W^{f(q)}], \quad (17)$$

where $\mathbf{E}[W^{f(q)}]$ is the average window size in steady state for a TCP flow with *fixed* marking/dropping probability $f(q)$. As a result, the steady state throughput of the limiting behavior can be calculated from a well-known TCP throughput model with fixed loss probability, e.g., [3] [4].

We now consider a simple application of this limiting result; An ISP currently services up to N_1 TCP flows at peak hour through an ECN/RED access gateway connecting to the core network with the link speed of $N_1 C$ packets/second. The network manager can roughly determine the buffer utilization in the ECN/RED gateway as follows:

(i) Determine the marking probability per flow ($p = f(q)$) from the relation $C = \mathbf{E}[W^p]$ by using a TCP throughput formula such as the one in [3];

(ii) Calculate the limiting queue length q in steady state by solving $p = f(q)$;

(iii) Approximate the queue length distribution in steady state via the CLT complement. If the steady state exists, the CLT complement determines the distribution of the queue size fluctuations around q . The delay and overflow distributions can also be approximated via the CLT complement;

While these limiting results apply only to TCP flows with identical round-trip, there are situations where they could be useful. For example, the buffer dimensioning problem in an intercontinental Internet link where it is typically a bottleneck, its large propagation delay dominates the round-trip and the number of flows is extremely large.

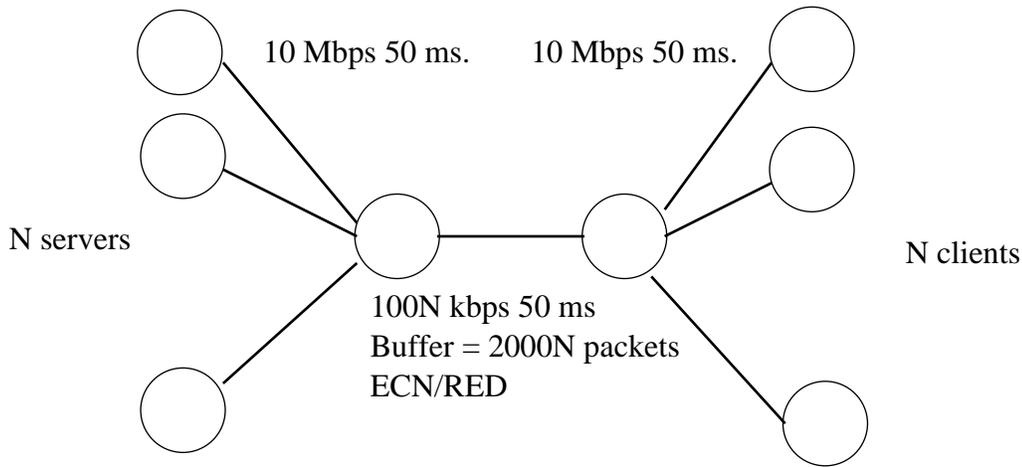


Figure 1: The simulation setup in NS-2.

11 Simulation results

In this section, we present results from (Monte-Carlo) simulations of the model presented in Section 7³ and from NS-2 simulations [8] to demonstrate that the behaviors suggested by both Theorem 3 and 4 do exist. For the NS-2 simulations, we use the system shown in Figure 1. Each server establishes a TCP Reno connection to a corresponding client, thereby competing for the capacity in the ECN/RED gateway. Each TCP has a fixed packet size of 1500 bytes and maximum window size of 200 packets. The marking probability function in the ECN/RED gateway is specified as

$$f^{(N)}(Q) = f(N^{-1}Q) \quad (18)$$

with $f : \mathbb{R}_+ \rightarrow [0, 1]$ taken to be

$$f(x) = \min(0.01(x-1)^+, 1), \quad x \geq 0.$$

The “time constant” parameter w_q for the Exponential Weighted Moving Average is set to 0.002, similar to the recommended value in [30]. Every round equals the round-trip propagation delay of 200 milliseconds. At the beginning of each round, we collect the instantaneous queue length in the ECN/RED buffer for a total duration of 200 seconds. Figure 2 shows the queue length normalized by the number of connection (N) as a function of time. We note a behavior similar to that discussed in Theorem 3 as fluctuations in the normalized queue length decrease with the number of connections increasing. Moreover, Figure 2 also suggests the existence of a steady-state for the limiting model, with a steady-state normalized queue length being constant at approximately 4.85 packets/user, corresponding to the steady-state marking probability of $0.0385 = f(4.85)$.

To simulate the model described in Section 7, we use the same parameter setup as in the NS simulation, i.e., $W_{\max} = 200$, simulation time of 1000 timeslots and the same marking function. The capacity per user (C) of the bottleneck router can be calculated from (17). We use the following approximation

$$\mathbf{E}[W^p] \simeq \min\left(W_{\max}, \sqrt{\frac{3}{2p}}\right) \quad (19)$$

(given as Eqn. 33 in [3]) to calculate the steady-state throughput when the marking probability p is 0.0385 (obtained from the NS simulation). A simple calculation yields $C = 6.24$ packets/timeslot. The simulation result is shown in Figure 3. By a quick comparison to Figure 2, note a qualitative similarity where the fluctuation decreases as the number of users increases. Further inspection reveals that the average normalized queue length is around 4.93 packets/user, very close to 4.85 packets/user produced by the NS simulation. Therefore, the model appears to capture the essential behavior of queue dynamics in ECN/RED gateways.

To gauge the rate of convergence, we assume that the queue is in steady state after the first 100 samples and that the process is Gaussian and ergodic. Therefore, the steady-state standard deviation of the queue can be approximated from the sample standard deviation of the queue at timeslot 101 and after. The comparison between the sample standard deviation from the model and from NS simulation is displayed in Figure 4. It is clear that both follow a similar trend. We also expect from Theorem 4 that the standard deviation will decrease as function of $N^{-0.5}$ for large N . Let S_N denote the sample standard deviation of the normalized queue when the number of users is N . We can see from Figure 4 that $N^{-0.5}S_1$ provides a good approximation of the standard deviation S_N for large N .

³The simulation results for the rate-based model are qualitatively similar to the window-based model and can be found in [28].

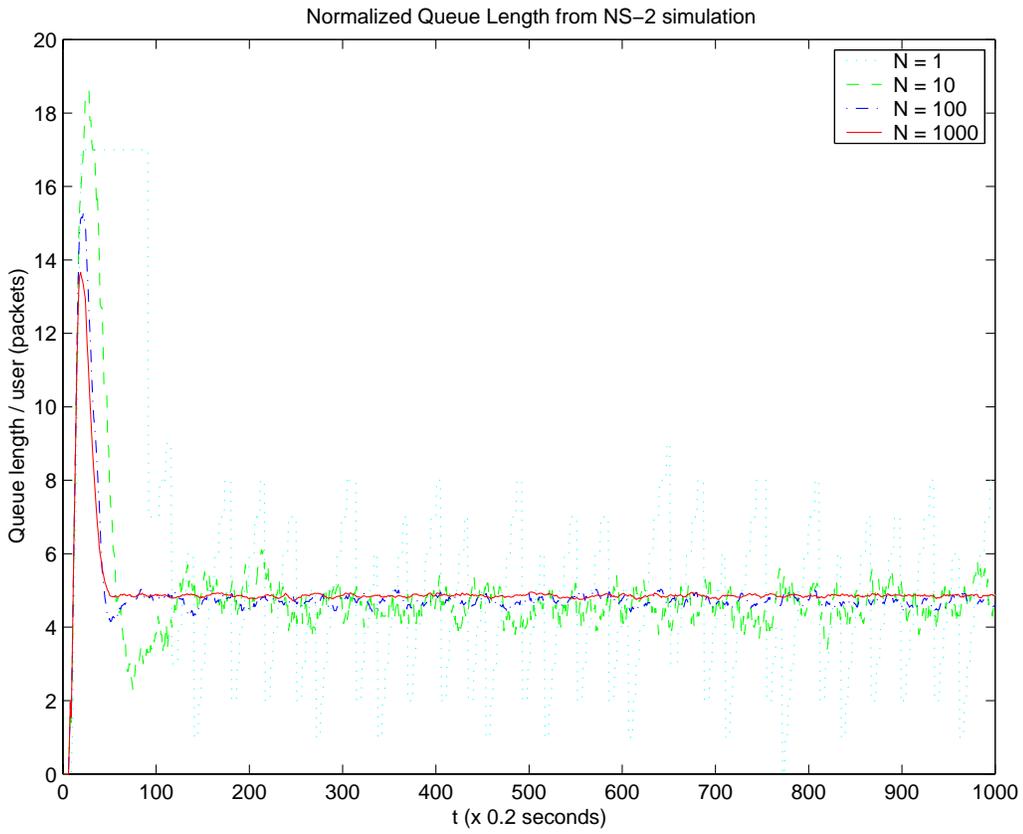


Figure 2: The normalized queue length of the ECN/RED gateway in NS-2 simulation.

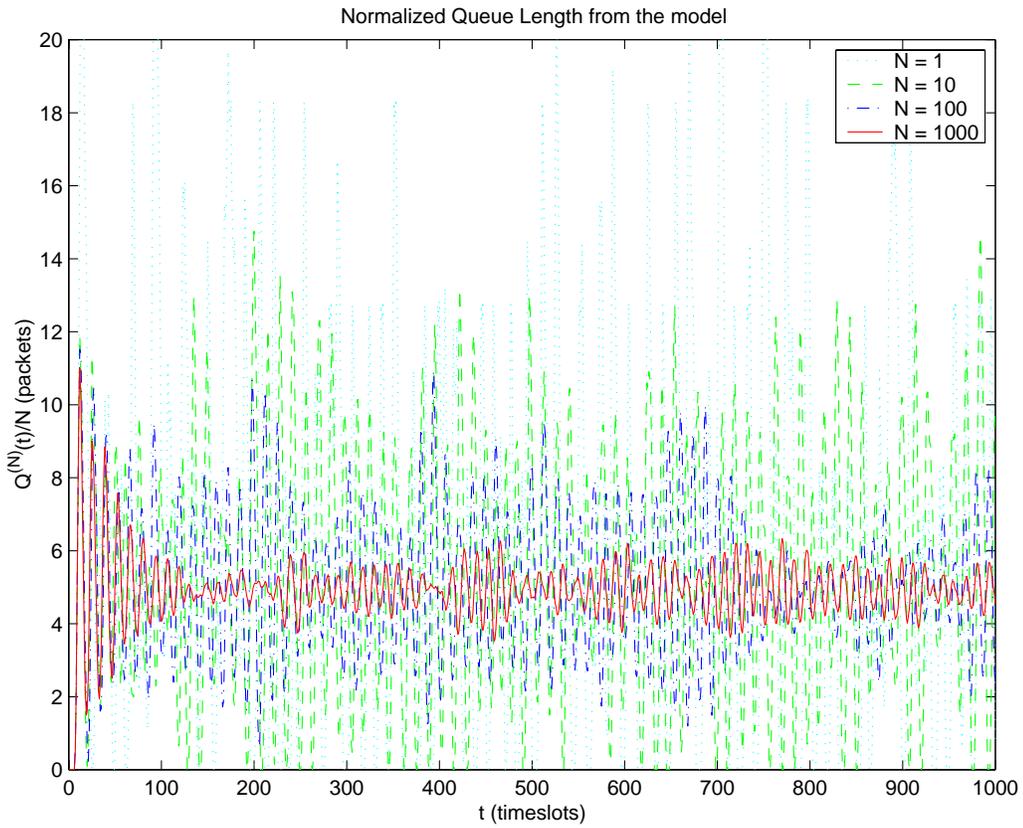


Figure 3: The normalized queue length of the model.

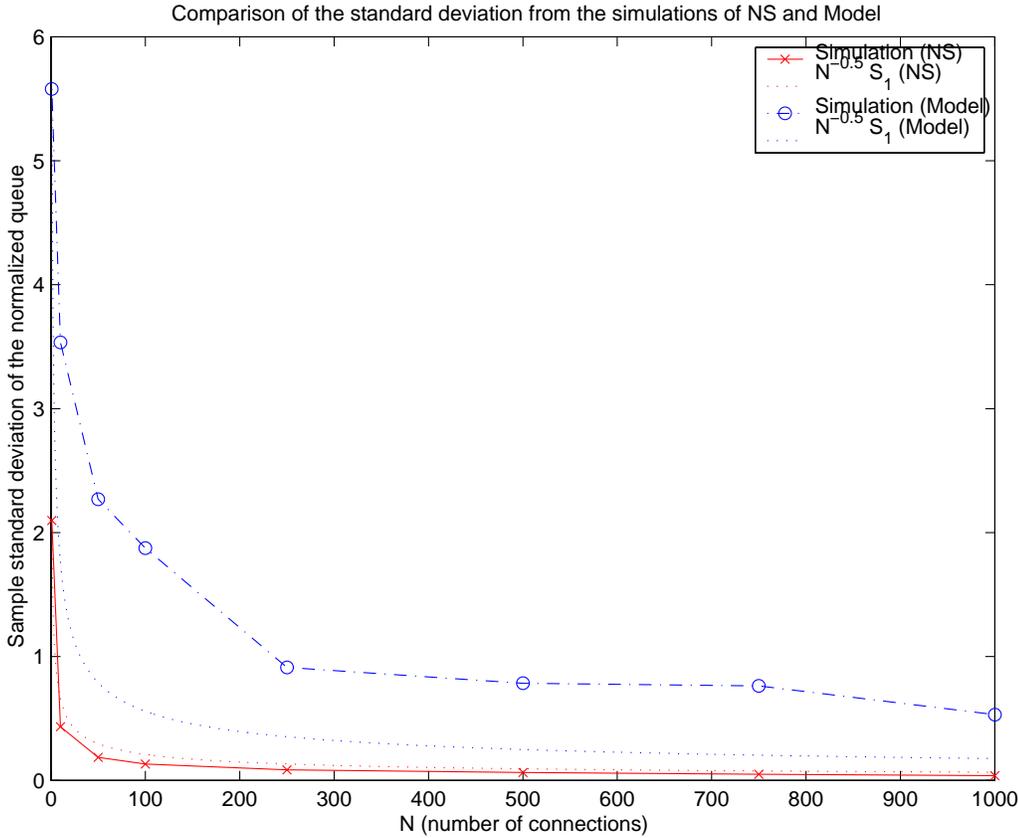


Figure 4: Sample standard deviation of the normalized queue.

12 Other limit theorems on TCP/RED

In this section, we outline two modeling approaches via limiting theorems which use asymptotics in a parameter other than the number of TCP flows.

Dumas et al. [32] analyze the average throughput for a single TCP connection with a constant loss rate. Even though formulae are well known from previous work [2] [3] [4] [5] [6] [7], they all require ad-hoc approximations. However, it is shown in [32] that as the loss probability approaches zero, the *exact* asymptotic throughput formula can be identified as follows:

Theorem 5 *Let ρ^α be the steady-state throughput of a TCP flow when α is the loss probability of the flow. If R denotes the round-trip delay, then*

$$\lim_{\alpha \rightarrow 0} \sqrt{\alpha} \rho^\alpha = \frac{\kappa}{R} \quad \text{where } \kappa \simeq 1.31.$$

In other words, we recover the asymptotics

$$\rho^\alpha \sim \frac{\kappa}{R\sqrt{\alpha}}.$$

Surprisingly, while the value 1.31 for the constant is derived solely from the structure of the problem, it is nevertheless in agreement with the values obtained from experiments and simulations [33]. On the other hand, ad-hoc approximations suggest the values $\sqrt{3/2} \approx 1.22$ [3] or $\sqrt{2}$ [7] for this constant. This evidence provides one more data point in support of the robustness of traffic modeling via limit theorems.

In a different vein, Sharma and Purkayastha [34] consider an ODE approximation of the average queue process of a RED gateway shared by multiple TCP connections when the exponential average parameter approaches zero. In the limit, a simplification in the model occurs as the complex behavior of the average queue length can be approximated by an ODE as the memory of the queue increases and well-established numerical methods can be unleashed to evaluate various metrics of interest.

13 Conclusions

In this paper, we have outlined existing problems with micro-scale TCP traffic modeling and have advocated the need for a “robust” traffic modeling approach. Robustness is typically achieved by looking at the system at the appropriate scale where some parameters are naturally very large or very small in relation to other parameters. This approach via limit theorems has been used with some resounding success in a number of setups, e.g., Poisson models of call requests at a telephone exchange in classical telephony and of multiplexed bursty data source at a multiplexer in data networks. In the limit, irrelevant details are removed and a simpler yet robust model results. Natural parameters that are good candidates for taking a limit are the number of TCP flows, the drop-probability of a TCP flow, and the exponential average parameter in RED.

While this survey demonstrated the feasibility of limit theorems to TCP modeling, challenges remain. First, the model needs to accommodate heterogeneous TCP flows. Second, short-lived flows and uncontrolled flows (e.g. UDP flows) need to be incorporated. While the CLT result provides some insights on the relationship of the dropping/marketing probability function and queue fluctuations, further studies can be done to solve the optimization problem where the target queue size and queue fluctuation can be achieved by a certain feedback function.

References

- [1] Van Jacobson, “Congestion avoidance and control,” in *Proceedings of SIGCOMM’88 Symposium*, Aug. 1988, pp. 314–332.
- [2] J. Padhye, V. Firoiu, and D. Towsley, “A stochastic model of TCP Reno congestion avoidance and control,” Tech. Rep. 99-02, Department of Computer Science, University of Massachusetts, Amherst, 1999.
- [3] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, “Modeling TCP Reno performance: a simple model and its empirical validation,” in *IEEE/ACM Transactions on Networking*, Apr. 2000.
- [4] M. Mathis, J. Semske, J. Mahdavi, and T. Ott, “The macroscopic behavior of TCP congestion avoidance algorithm,” *Computer Communication Review*, vol. 27, no. 3, July 1997.
- [5] T. V. Lakshman and U. Madhow, “The performance of TCP/IP for networks with high bandwidth-delay products and random loss,” *IEEE/ACM Transactions on Networking*, June 1997.
- [6] E. Altman, K. Avrachenkov, and C. Barakat, “TCP in presence of bursty losses,” in *Proceedings of SIGMETRICS*, 2000.
- [7] V. Misra, W. Gong, and D. Towsley, “Stochastic differential equation modeling and analysis of TCP window size behavior,” in *Proceedings of Performance*, 1999.
- [8] “UCB/VINT/LBNL Network Simulator (ns) version 2.0,” <http://www.isi.edu/nsnam>.
- [9] L. Zhang and D. Clark, “Oscillating behavior of network traffic: A case study simulation,” *Internetworking: Research and Experience*, vol. 1, no. 2, pp. 101–112, 1990.
- [10] Sally Floyd and Van Jacobson, “Random early detection gateways for congestion avoidance,” *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397–413, Aug. 1995.
- [11] Sally Floyd, “TCP and explicit congestion notification,” *Computer Communication Review*, vol. 24, no. 5, pp. 10–23, Oct. 1994.
- [12] Frank P. Kelly, “Charging and rate control for elastic traffic,” *European Transactions on Telecommunications*, vol. 8, pp. 33–37, 1997.
- [13] Frank Kelly, A.K. Maulloo, and D.K.H. Tan, “Rate control in communication networks: shadow prices, proportional fairness and stability,” *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [14] Steven H. Low, “A duality model of TCP and queue management algorithms,” in *Proceedings of ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, Sept. 2000.
- [15] C. Hollot, V. Misra, D. Towsley, and W. Gong, “A control theoretic analysis of RED,” in *Proceedings of IEEE INFOCOM*, 2001.
- [16] A. Abouzeid and S. Roy, “Analytic understanding of RED gateways with multiple competing TCP flows,” in *Proceedings of IEEE GLOBECOM*, 2000.
- [17] H. M. Alazemi, A. Mokhtar, and M. Azizoglu, “Stochastic modeling of random early detection gateways in TCP networks,” in *Proceedings of IEEE GLOBECOM*, 2000.
- [18] A.K. Erlang, “The theory of probabilities and telephone conversation,” *Nyt Tidsskrift Mat.*, 1909.

- [19] C. Palm, "Analysis of the Erlang traffic formulae for busy-signal arrangements," *Ericsson Tech.*, 1938.
- [20] A.Y. Khintchine, "Mathematisches uber die erwartung vor einemoffenthchen schalter," *Mat. Sb.*, 1932.
- [21] G. Hasegawa and M. Murata, "Analysis of dynamic behaviors of many TCP connections sharing Tail-Drop/RED routers," in *Proceedings of IEEE GLOBECOM*, Nov. 2001.
- [22] Michele Garetto, Renato Lo Cigno, Michela Meo, and Marco Ajmone Marsan, "A detailed and accurate closed queueing network model of many interacting TCP flows," in *Proceedings of IEEE INFOCOM*, 2001, pp. 1706–1715.
- [23] Francois Baccelli, Dohy Hong, and Zhen Liu, "Fixed point methods for the simulation of the sharing of a local loop by a large number of interacting TCP connections," Tech. Rep. RR-4154, INRIA, Apr. 2001.
- [24] Thomas Bonald, "Comparison of TCP Reno and TCP Vegas via fluid approximation," Tech. Rep. 3563, INRIA, Nov. 1998.
- [25] Sanjay Shakkottai and R. Srikant, "Mean FDE models for Internet congestion control under a many-flows regime," Tech. Rep., Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Aug. 2001.
- [26] Dohy Hong and Dmitri Lebedev, "Many TCP user asymptotic analysis of the AIMD model," Tech. Rep. RR-4229, INRIA, July 2001.
- [27] Cdric Adjih, Philippe Jacquet, and Nikita Vvedenskaya, "Performance evaluation of a single queue under multi-user TCP/IP connections," Tech. Rep. RR-4141, INRIA, Mar. 2001.
- [28] Peerapol Tinnakornrisuphap and Armand Makowski, "Queue dynamics of RED gateways under large number of TCP flows," in *Proceedings of IEEE GLOBECOM*, 2001.
- [29] Peerapol Tinnakornrisuphap and Armand Makowski, "Limit behavior of ECN/RED gateways under large number of TCP flows," Tech. Rep., The Institute for Systems Research, University of Maryland, 2002.
- [30] Sally Floyd, "RED: Discussions of setting parameters," <http://www.icir.org/floyd/REDparameters.txt>.
- [31] Victor Firoiu and Marty Borden, "A study of active queue management for congestion control," in *Proceedings of IEEE INFOCOM*, 2000.
- [32] Vincent Dumas, Fabrice Guillemin, and Philippe Robert, "Limit results for markovian models of TCP," in *Proceedings of IEEE GLOBECOM*, 2001.
- [33] J. Mahdavi and S. Floyd, "TCP-friendly unicast rate-based flow control," Technical Note sent to the end2end-interest mailing list, Jan. 1997, Available from http://www.psc.edu/networking/papers/tcp_friendly.html.
- [34] Vinod Sharma and Punyaslok Purkayastha, "Performance analysis of TCP connections with RED control and exogenous traffic," in *Proceedings of IEEE GLOBECOM*, 2001.