*Probability in the Engineering and Informational Sciences*, **16**, 2002, 367–388. Printed in the U.S.A.

# PERFORMANCE ANALYSIS OF THE RANDOM EARLY DETECTION ALGORITHM

V. SHARMA
*Department of Electrical Communication Engineering*
*Indian Institute of Science*
*Bangalore, 560012, India*
*E-mail: vinod@ece.iisc.ernet.in*

J. VIRTAMO AND P. LASSILA
*Networking Laboratory*
*Helsinki University of Technology*
*FIN-02015 HUT, Finland*
*E-mail: {Jorma.Virtamo; Pasi.Lassila}@hut.fi*

In this article we consider a finite queue with its arrivals controlled by the random early detection algorithm. This is one of the most prominent congestion avoidance schemes in the Internet routers. The aggregate arrival stream from the population of transmission control protocol sources is locally considered stationary renewal or Markov modulated Poisson process with general packet length distribution. We study the exact dynamics of this queue and provide the stability and the rates of convergence to the stationary distribution and obtain the packet loss probability and the waiting time distribution. Then we extend these results to a two traffic class case with each arrival stream renewal. However, computing the performance indices for this system becomes computationally prohibitive. Thus, in the latter half of the article, we approximate the dynamics of the average queue length process asymptotically via an ordinary differential equation. We estimate the error term via a diffusion approximation. We use these results to obtain approximate transient and stationary performance of the system. Finally, we provide some computational examples to show the accuracy of these approximations.

## 1. INTRODUCTION

Random early detection (RED) was proposed by Floyd and Jacobson [11] as an effective mechanism to control the congestion in the network routers/gateways. It

© 2002 Cambridge University Press 0269-9648/02 $12.50 

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

also helps to prevent the *global synchronization* in the transmission control protocol (TCP) connections sharing a congested router and to reduce the bias against bursty connections. It was an improvement over the previous proposals such as Random Drop and Early Random Drop (see Hashem [12], Mankin and Ramakrishnan [18], and Zhang [25]). Currently, it has been implemented in various routers and has been recommended for queue management by IETF (see Braden et al. [5]). Recently, Clark and Fang [7] have proposed the incorporation of differentiated services (Diff Serv) in the Internet by applying RED with different parameter settings to the "In" and "Out" packets of the flows arriving at a router. See Internet RFC-2475 [3] for details on the Diff Serv Architecture.

To be able to apply RED in the differentiated services environment, it is important to study its queuing behavior in detail. For example, to provide any guarantees on throughput and delays, one needs to study these as a function of the RED parameters. Previously, this has been studied via simulations in Floyd and Jacobson [11] and Lin and Morris [16]. However, it is useful to complement these efforts with an analytical study.

In a more recent reference, a simple analytical model of this mechanism has been developed in May, Bolot, Jean-Marie, and Diot [19] and Bonald and May [4]. However, RED is more complicated and several of its important features (e.g., average queue length threshold and the counter during congestion buildup time) have not been captured in this model. A more detailed analysis has been presented in Peeters and Blondia [21]. They consider a slotted system with the packet lengths having a truncated geometric distribution. Then a finite Markov chain approximation for the overall system is made. Other recent articles on the performance of RED, although not explicitly on analytical modeling of it, include Ott, Lakshman, and Wong [20] and Feng, Kandlur, Saha, and Shin [10].

A full description of the problem should consider a closed system consisting of two parts: (1) a queue receiving an aggregate flow of packets and reacting dynamically to the changes in the total flow rate under the control of the RED algorithm and (2) the population of TCP sources, each of which reacts to the packet losses at the queue as determined by the TCP flow control mechanism (see Fig. 1). In this article we focus on part 1 of the problem. It is reasonable to assume that, locally, the aggregate traffic stream can be described by a Poisson process [or maybe by some more refined model, such as the Markov modulated Poisson process (MMPP)]. In general, the arrival rate of the aggregate packet stream is time varying (or the pa-

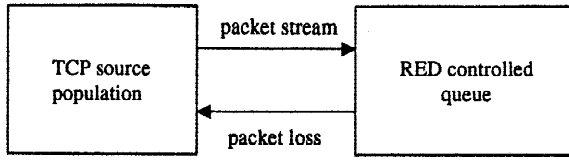



**FIGURE 1.** Interaction between the TCP population and the RED-controlled queue.

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

rameters of the MMPP are time varying). However, the typical timescale of these variations is of the order of a round-trip time, which for long-distance TCP connections is long in comparison with the timescale of the dynamics of the queue in a core network router. In this case, the queue will be in a quasiequilibrium.

So, as a first approximation for part 1, it is meaningful to consider the behavior of an RED-controlled queue receiving a Poisson or an MMPP stream of packets. This problem is studied in detail in this article and results relating to the stability of the queue and the equilibrium of the queue are provided. Given these results, from the point of view of part 2 of the system, the queue can be considered as a black box which has a given relation between the parameters of the (stationary) input process and packet losses. The analysis of the dynamics of system 2, however, is outside the scope of the present article.

When the timescale of the variations of system 2 is not very much longer than that of system 1, the stationarity assumption of the input process is no longer valid. This leads us to consider the transient behavior of the RED-controlled queue. We present a new approach based on an asymptotical approximation, which leads to an ordinary differential equation (ODE) description of the average queue length dynamics. We obtain approximate transient and stationary behavior of the system via this approximation. The error in the ODE approximation is quantified by diffusion limits.

In this article, in particular, we consider the transient where an initially empty system starts to receive the input stream at time 0. The same technique, however, can be used virtually without any modification to describe the case where input process varies smoothly in time. Again, given these results, the RED-controlled queue can be considered as a black box relating, now dynamically, the input process and packet losses. Such a model can then be incorporated in a larger model describing the interaction of parts 1 and 2.

Finally, we consider the special case of exponential packet size distribution and Poisson arrival process, which allows us to derive some exact results. The accuracy of the ODE approximation is also tested in this exponential case against simulated results.

The article is organized as follows. In the next section we explain the RED algorithm. In Section 3 we prove the stability (ergodicity) of its queuing model as well as the rates of convergence to the stationary distribution for the renewal and MMPP arrival streams with general packet length distributions. Next, we extend these results to a two traffic class case. In the process, we extend our analysis to regenerative arrival streams. In Section 4 we provide the ODE approximation and the diffusion approximations. In Section 5 we explicitly write formulas for the exponential interarrival and service times. In Section 6 we verify our approximations via computations and simulations.

## 2. THE RED ALGORITHM

This section explains the model we study in this article. Consider a queue with a finite buffer which can hold up to $K$ packets at a time. This could be the queue at an

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

output link of a router. The input to the queue is a renewal sequence or an MMPP and the packet lengths are independent and identically distributed (i.i.d.) with a general distribution. The arriving packets to the queue are subjected to admission control by the RED algorithm, explained below.

Let $q_n$ be the queue length [i.e., the number of packets in the system (including the one in service) at the time of the $n$th packet arrival]. Let $\hat{q}_n$ be the average queue length (to be defined below) at the $n$th arrival. The $n$th packet will be discarded with probability $p_n$, which depends on $\hat{q}_n$ and the past history. By defining $p_n$ appropriately, the RED algorithm controls congestion in the router before its queues overflow. Fix the constants $p_{\max}, T_{\min}$, and $T_{\max}$ appropriately such that $0 \le T_{\min} \le T_{\max} \le K$. $p_n$ is defined as

if $\hat{q}_n < T_{\min}$ then
  $p_n = 0$, $c_n = -1$
else if $\hat{q}_n > T_{\max}$ then
  $p_n = 1$, $c_n = 0$
else if $T_{\min} \le \hat{q}_n \le T_{\max}$ then
  $c_n = c_{n-1} + 1$,
  $C_n = (T_{\max} - T_{\min})/[\, p_{\max}(\hat{q}_n - T_{\min})]$ and
  $p_n = 1/(C_n - c_n)$
end if

Then, with probability $p_n$, discard the packet and $c_n = 0$,
and, with probability $1 - p_n$, queue the packet.

The average queue length $\hat{q}_n$ is defined as

$$\begin{aligned} \hat{q}_n &= (1-\beta)\hat{q}_{n-1} + \beta q_n \quad \text{if } q_n > 0 \\ \hat{q}_n &= (1-\beta)^m \hat{q}_{n-1} \qquad\quad \text{otherwise,} \end{aligned} \tag{1}$$

where $m =$ (idle time during the interarrival time of the $n$th packet)$/\bar{s}$, $\bar{s}$ is the typical transmission time for a small packet, and $0 < \beta < 1$ is an appropriately fixed constant. Typically, $\beta$ is chosen to be relatively small (e.g., Floyd and Jacobson [11] propose using $\beta \ge 0.001$ and use $\beta = 0.002$ in their simulations).

In the above algorithm, it is possible that $p_n$ can become negative or greater than 1. In that case we set $p_n = 0$ or 1 as the case may be. Also, it is possible that the algorithm decides that an arriving packet should be admitted, but $q_n = K$. In that case the packet will be discarded and we set $c_n = 0$.

Clark and Fang [7] suggest using this algorithm to provide differentiated service. The packets of different flows are classified into two classes: class 1 and class 2. The admission of class $i$ packets to the queue is controlled by the RED algorithm with parameters $T_{\min,i}, T_{\max,i}$, and $p_{\max,i}$. Class 1 is given higher priority by assigning

$$T_{\min,2} \le T_{\min,1}, \qquad T_{\max,2} \le T_{\max,1}, \qquad p_{\max,1} \le p_{\max,2}. \tag{2}$$

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

Class 1 is composed of the packets of the "premium" service class, which are in conformance with their traffic profiles. Class 2 carries the best effort traffic and the packets of the premium class, which are not in conformance with their traffic profiles. Also, for traffic class 1, we consider the average queue length $\hat{q}_n$ of only class 1 packets, whereas for class 2, the average queue length of all the packets is considered in the algorithm.

## 3. THE ANALYSIS

In this section we study the process $\{(q_n, \hat{q}_n, c_n)\}$ corresponding to the queue defined in the previous section. We show the stability of the queue and provide the rates of convergence to the stationary distribution. It is computationally very intensive to obtain these stationary distributions even for Poisson arrivals and exponential service times. Therefore, in Section 4 we will obtain asymptotic approximations for our system. As mentioned earlier, our model is much more realistic than the one studied in May et al. [19]. In Section 3.1 we study the queue with the renewal arrival stream. In Section 3.2 we consider an MMPP arrival sequence. Section 3.3 extends the results to the two traffic class case while generalizing the arrival stream to a regenerative sequence (which covers the models of long-range-dependent traffic considered in the literature).

### 3.1. Renewal Arrivals

The process $\{(q_n, \hat{q}_n, c_n)\}$ has state space

$$\{0,1,\ldots,K\} \times [0, T_{\max}] \times \{0,1,\ldots\}.$$

Actually, $\hat{q}_n$ can exceed $T_{\max}$ a little, but for the purpose of this section it will not make any difference. We will assume that the interarrival times (positive with nonzero probability) and the packet lengths form i.i.d. sequences with general distributions and these sequences are independent of each other. Then $\{(q_n, \hat{q}_n, c_n)\}$ becomes a regenerative process with the visit times to the state $(0,0,-1)$ as regeneration epochs. If we show that the regeneration length $\tau$ of the process has a finite mean and an aperiodic distribution, then this process will have a unique stationary distribution $\pi$ (Asmussen [1]). Also, starting from any initial distribution, it will converge to $\pi$ in total variation. To be able to show this we will make extra assumptions. First, once $\hat{q}_n$ is nonzero, it may not become zero again even though it may come arbitrarily close to it. Therefore, we modify the queue dynamics. We choose an arbitrarily small positive constant $\varepsilon < T_{\min}$ such that whenever $\hat{q}_n \leq \varepsilon$ we make $\hat{q}_n = 0$. If $\varepsilon$ is chosen small enough, it will not affect the dynamics and the performance of the system. In addition, to do the computations later, we will need to discretize the range $[0, T_{\max}]$ of $\hat{q}_n$ anyway, which will require us to make such a change. Next, we need to take care of the residual service time of the packet being transmitted at the time of an arrival. To control this quantity we will make various alternative assumptions—the service time can be phase type or new better than used (NBU) or bounded (in the

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

Internet and all other existing networks, these are bounded). Also, in the practical queue, the number of bits that can be stored is limited. This will enforce the boundedness assumptions on the packets. Under these assumptions we obtain the ergodicity of the process. In the following, $a$ and $s$ will denote generic interarrival and service times, respectively.

Start the process $\{(q_n, \hat{q}_n, c_n)\}$ in the state $(0,0,-1)$ at time $t = 0$. We make the assumption that the service time distribution is of phase type (we will comment on the NBU or the bounded packet length assumption later). We also assume that the interarrival times are spread out. Then, there is an $\varepsilon_1 > 0$ and a $\beta_1 > 0$ such that $\mathrm{P}\{\gamma < \beta_1\} > \varepsilon_1$ and $\mathrm{P}\{s < \beta_1\} > \varepsilon_1$, where $\gamma$ is a residual service time at the time of an arrival. Then, irrespective of the queue length at the time of (say) the $n$th arrival, with probability greater than $\varepsilon_1^K$, the queue will be empty at time $K\beta_1$ (if the $(n+1)$st arrival has not come yet). Also, let $N_1$ be an integer such that $(1-\beta)^{N_1} T_{\max} < \varepsilon$. By the spread-out property of $a$ (actually we need only that $\mathrm{P}\{a > K\beta_1 + N_1\bar{s}\} > \delta$ for some $\delta > 0$), the probability that the $(n+1)$st arrival comes after time $K\beta_1 + N_1\bar{s}$ is positive. Therefore, with a positive probability (lower bounded away from 0) whenever the next arrival (i.e., the $(n+1)$st arrival) to the system occurs, $q_{n+1} = 0$, $\hat{q}_{n+1} < \varepsilon$ (then we make $\hat{q}_{n+1} = 0$) and, hence, $c_{n+1} = -1$. Therefore, by the geometric coin-tossing argument, $E\tau^\alpha < \infty$ for any $\alpha > 0$. Also, if $(q_n = 0, \hat{q}_n = 0, c_n = -1)$, we have shown with a positive probability that this event occurs at the $(n+1)$st arrival also. Hence, $\tau$ is aperiodic. This provides the existence and uniqueness of the stationary distribution $\pi$ and also convergence to it in total variation.

If, instead of being phase type, $s$ is NBU, then, again, we can find a $\beta_1 > 0$ and an $\varepsilon_1 > 0$ satisfying the assumptions in the last paragraph. Also, for the bounded packet lengths, the condition is trivially satisfied.

The fact that $E\tau^\alpha < \infty$ holds for $\alpha > 1$ provides us with the following rate of convergence to stationary distribution (Kalashnikov [13])

$$\sup_A |\mathrm{P}\{(q_n, \hat{q}_n, c_n) \in A\} - \pi(A)| < cn^{1-\alpha}, \tag{3}$$

where $c$ is some constant and the process can start at $n = 0$ in any arbitrary distribution. Furthermore, from Sharma [22], using the boundedness of $(q_n, \hat{q}_n)$, when $E\tau^2 < \infty$, we obtain the functional central limit theorem for the process $(q_n, \hat{q}_n)$. We will use this fact to obtain diffusion approximations in the next section.

Next, let us consider the continuous-time process $\{(q(t), \hat{q}(t), c(t)), t \geq 0\}$, where $(q(t), \hat{q}(t), c(t))$ represents the state of the system at time $t$. This process is also regenerative with regeneration epochs at the arrival epochs to the state $(0,0,-1)$. Then the regeneration length $\hat{\tau}$ for this process is $\sum_{n=0}^{\tau-1} a_n$, where $\{a_n\}$ are the interarrival times, if $n = 0$ is a regeneration epoch for the process $\{(q_n, \hat{q}_n, c_n)\}$. If $Ea^\alpha < \infty$, for $\alpha \geq 1$, then we obtain $E\hat{\tau}^\alpha < \infty$. From this then follows the existence and uniqueness of the stationary distribution $\hat{\pi}$ and the rate of convergence (3) for the continuous-time process also.

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

If the arrival process is Poisson, then by PASTA, the time stationary distribution of $q(t)$ is the same as that of $q_n$. Thus, knowing $\pi$, by Little's law, we know the mean sojourn time of a packet admitted to the queue.

Next, we consider the waiting time process $\{w_n\}$ where $w_n$ is the waiting time of the $n$th packet. The process $\{(w_n, q_n, \hat{q}_n, c_n)\}$ is also regenerative with regeneration state $(0,0,0,-1)$, which occurs whenever $(q_n, \hat{q}_n, c_n) = (0,0,-1)$. Thus, we have also shown the existence and uniqueness of its stationary distribution $\pi$ and the rates of convergence to $\pi$ for this process. However, for our finite queue, the waiting times may still be unbounded (because we have not assumed the packet lengths to be bounded). However, because for $0 \le m \le \tau - 1$,

$$w_m \le \sum_{n=0}^{\tau-1} a_n = \hat{\tau},$$

we obtain

$$\mathrm{E}_\pi[w_m^\alpha] = \frac{\mathrm{E}\left[\sum_{m=0}^{\tau-1} w_m^\alpha\right]}{\mathrm{E}[\tau]} \le \frac{\mathrm{E}[\tau \hat{\tau}^\alpha]}{\mathrm{E}[\tau]} \le \frac{\mathrm{E}[\tau^p]^{1/p}, \mathrm{E}[\hat{\tau}^{q\alpha}]^{1/q}}{\mathrm{E}[\tau]}$$

for any $p > 1$ and $1/p + 1/q = 1$. Therefore, $\mathrm{E}_\pi[w_m^\alpha] < \infty$ if $\mathrm{E}[\tau^{\alpha+1}] < \infty$ and $\mathrm{E}[\hat{\tau}^{\alpha+1}] < \infty$, which holds if $\mathrm{E}[a^{\alpha+1}] < \infty$.

If we can compute $\pi$, we will be able to compute the stationary probability of packet loss also. However, because $\hat{q}$ is real-valued, even when $(q_n, \hat{q}_n, c_n)$ is Markov, it is computationally very expensive to obtain $\pi$ and $\hat{\pi}$. Therefore, in Section 4 we will use a decomposition approach to substantially ease this job.

### 3.2. MMPP Arrivals

The assumption of i.i.d. interarrival times is restrictive because it is known that, in networks, the interarrival sequence is correlated. Thus, now we extend our results to the case when the interarrival sequence is an MMPP modulated by a finite-state irreducible Markov chain $\{J(t)\}$. We assume that when $J(t) = i$, the instantaneous arrival rate is $\lambda_i$ (for convenience, we will keep the packet length distributions independent of $J(t)$ although this condition can also be generalized). Fix a state $i_0$ of $\{J(t)\}$. We will consider $(0,0,-1,i_0)$ to be a regeneration state for the process $\{(q_n, \hat{q}_n, c_n, J_n)\}$, where $J_n$ is the state of $\{J(t)\}$ at the arrival instant of the $n$th packet to the queue. Start this process in state $(0,0,-1,i_0)$ at time $n = 0$. Again we denote the intervisit time to $(0,0,-1,i_0)$ by $\tau$. Now consider the proof for the i.i.d. case given in Section 3.1. Even for the MMPP case, we can find an $\varepsilon_2 > 0$ such that an interarrival time for this process under any state of $J(t)$ is greater than $K\beta_1 + N_1\bar{s}$, with probability $\varepsilon_2 > 0$. Thus, the next visit time to state $(0,0,-1)$ by $(q_n, \hat{q}_n, c_n)$ will have a finite $\alpha$th moment for any $\alpha > 0$. Once state $(0,0,-1)$ is reached, with a probability lower-bounded by an $\varepsilon_3 > 0$, the following event will happen: In an interval of length $\delta_2$, $J(t)$ will reach state $i_0$, and in a further interval of length $\delta_2$, $J(t)$ will still

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

be in state $i_0$ while there is an arrival to the queue. That will ensure that state $(0,0,-1,i_0)$ is reached by the process with a positive probability (lower bounded away from zero) once state $(0,0,-1)$ is reached. Now, previous arguments ensure that the process $\{(q_n,\hat{q}_n,c_n,J_n)\}$ has regeneration lengths with a finite $\alpha$th moment for each $\alpha$. Aperiodicity of $\tau$ can also be shown, as earlier. Thus, we obtain existence and uniqueness of the stationary distribution and convergence rate (3) for this process. Finiteness of the moments of the waiting times can also be shown, as earlier. Now, the time stationary distribution $\hat{\pi}$ of $(q(t),\hat{q}(t),c(t),J(t))$ (the existence and uniqueness of the stationary distribution, etc. for this process can be justified as for the i.i.d. case) is related to that of $\pi$ by

$$\mathrm{P}_{\hat{\pi}}[(q(t),\hat{q}(t),c(t)) \in A, J(t) = i] = \mathrm{P}_{\pi}[(q_n,\hat{q}_n,c_n) \in A, J_n = i]\,\frac{\sum_j \pi_j \lambda_j}{\lambda_i},$$

where $\pi_j$ is the stationary probability of $\{J(t)\}$.

### 3.3. Two Traffic Classes

Next, we extend the above results to the two traffic class case described in Section 2 for Clark and Fang's scheme of differentiated services (see [7]). Now, we need to consider the process $\{(q_n,\hat{q}_{n,1},\hat{q}_n,c_{n,1},c_{n,2})\}$, where $c_{n,i}$ is the counter for class $i$, $q_n$ is the total queue length at the $n$th packet's arrival epoch (of the total traffic), $\hat{q}_n$ is the total average queue length, and $\hat{q}_{n,1}$ is the average queue length of class 1 traffic. Now, one technicality arises when we consider the two arrival streams $\{(a_{n,i},s_{n,i})\}$, $i = 1,2$, to be independent renewal sequences. Their superposition is not a renewal sequence unless both sequences are Poisson. Thus we generalize the framework which will not only allow us to handle this case but also to treat the MMPP case simultaneously (and, in fact, this allows the long-range dependence in the arrival stream). This technique was first used in Sharma [23] using the theory in Kalashnikov [13]. We assume the arrival stream of each class to be a regenerative sequence. Then, under weak conditions (see Kalashnikov [13] and Sharma [23]), we can have common regeneration epochs for both arrival streams (we denote the total arrival stream by $\{(a_n,s_n)\}$). This common regeneration length $\bar{\tau}$ has a spread-out distribution and finite $\alpha$th moment if the regeneration lengths of both individual streams have the corresponding properties. We also assume that the probability that no packet arrives to the system during a common regeneration length $\bar{\tau}$ (excluding the end points) is $\varepsilon_4 > 0$. Now, making the assumption of phase type, NBU or bounded service times of each class $i$ (both classes do not need to satisfy the same assumption), we can repeat the arguments of the i.i.d. traffic case for the single-class system to show that with a probability lower-bounded by (say) $\varepsilon_5 > 0$, irrespective of the state of the system at a common regeneration epoch, at the next such epoch the system will have the regeneration state $(0,0,0,-1,-1)$ for the process $\{(q_n,\hat{q}_{n,1},\hat{q}_n,c_{n,1},c_{n,2})\}$. Thus, all of our previous results will hold for this system.

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

## 4. ASYMPTOTIC APPROXIMATIONS

Until now we have been studying the system exactly. However, to compute the performance indices exactly requires unrealistic computational resources, even for exponential interarrival and service time distributions. Thus, in this section we develop approximation techniques that can provide asymptotically exact approximations of some of the quantities of interest. In Sections 5 and 6 we will provide some numerical results to show the accuracy of these approximations. Throughout this section the interarrival times are i.i.d. generally distributed or MMPP, whereas the service time distributions will be phase type for some of the results. For simplicity, we will restrict our study to the single-class case only.

The key parameter controlling the dynamics of the system is $\hat{q}_n$. The first time $\hat{q}_n$ exceeds $T_{\min}$ indicates the time of onset of the congestion state. When $\hat{q}_n$ exceeds $T_{\max}$, we declare that the system is too congested. Therefore, it is of interest to study the process $\{\hat{q}_n\}$ and the epochs when it exceeds $T_{\min}$ and $T_{\max}$. Furthermore, when $\beta$ is small (as suggested for practical systems), the process $\{\hat{q}_n\}$ moves very slowly as compared to $(q_n, c_n)$. Thus, the process $\{(q_n, \hat{q}_n, c_n)\}$ converges to stationarity slowly (therefore the transient period becomes important from a practical point of view). However, the $(q_n, c_n)$ component varies rapidly, and, for a given $\hat{q}_n = \hat{q}$, $(q_n, c_n)$ obeys approximately the stationary distribution of the $\{(q_n, c_n)\}$ process where $\hat{q}_n$ controlling the RED mechanism, as defined in Section 2, is fixed (i.e., $\hat{q}_n = \hat{q}$ for all $n$). (This is similar to the decomposition approach of Courtois [8].) This artificial process is denoted by $\{q_n(\hat{q}), c_n(\hat{q})\}$. It is easy to show that $\{(\hat{q}_n, c_n)\}$ has a unique stationary distribution for any $\hat{q}$. Also, starting from any initial conditions, the process converges in total variation to it. It is computationally much easier to obtain the stationary distribution of $\{q_n(\hat{q}), c_n(\hat{q})\}$ because of the reduction in the dimensionality.

Now, the actual average queue length $\hat{q}_n$, as given by (1), is not constant but moves slowly according to

$$\hat{q}_{n+1}^{\beta} = \hat{q}_n^{\beta} + \beta(q_{n+1}^{\beta} - \hat{q}_n^{\beta}) \tag{4}$$

whenever $q_n > 0$; and to emphasize the dependence of $(q_n, \hat{q}_n, c_n)$ on $\beta$ we have appended $\beta$ as a superscript to these quantities. In this section we will do the same to the other quantities as well.

Because we are really interested in the system dynamics when the system is heavily loaded (at least not lightly loaded because then the system behaves as a usual queue), $q_n = 0$ will rarely happen and, hence, (4) will provide a good enough approximation of $\{\hat{q}_n\}$ in the system. (Actually, the error thus introduced depends on $\beta$ and $\bar{s}$ also and can be nonnegligible. However, we will show that our main results—the ODE approximation—can be proved even with the inclusion of this correction. However, for simplicity in the main text in the rest of the article, we will ignore this correction.) Studying (4) exactly is not possible for our system. But in practice $\beta$ will be small and hence we study (4) as $\beta \downarrow 0$.

An advantage of identifying (4) as the key equation is that it has been well studied in the stochastic approximation literature (see Benveniste, Metivier, and

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

Priouret [2], Macchi [17], and Kushner and Yin [15]). However, the fact that $\beta$ is constant ("constant step size" case) and $\{q_n^\beta\}$ depends on $\beta$ precludes the use of most of the results available in the literature. However, this case has also been studied and we will use those results extensively in this section.

As a first-order approximation of (4), we will show that the process $\{\hat{q}_n^\beta\}$ when suitably scaled will converge in probability to the solution of a deterministic ODE as $\beta \downarrow 0$. This approximation is an analog of the law of large numbers (LLN). We will use the solution of the ODE to obtain the approximate behavior of the sample paths of $\{\hat{q}_n^\beta\}$ for small $\beta$. This, in particular, will provide an approximation of the first time to visit $T_{\min}$ and $T_{\max}$ and the distribution of the queue length (stationary as well as transient) and the probability of loss. The next step will be to obtain a second-order refinement of this approximation (analogous to the central limit theorem) by providing a diffusion approximation to the error between the sample paths of $\{\hat{q}_n^\beta\}$ and the solution of the ODE. We shall provide a justification of this approximation and verify via the numerical examples in Section 6. The theoretical justification of the above results will be provided in the Appendix. Section 4.1 provides the ODE approximation and Section 4.2 the diffusion approximation.

### 4.1. ODE Approximation

Let us consider (4) with $\hat{q}_0^\beta = a$. Define process $\{\hat{z}^\beta(t)\}$ as

$$\begin{aligned} \hat{z}^\beta(0) &= a, \\ \hat{z}^\beta(t) &= \hat{q}_{\lfloor t/\beta \rfloor} \end{aligned} \tag{5}$$

for $t > 0$, where $\lfloor t/\beta \rfloor$ is the integer part of $t/\beta$. Define the ODE

$$\begin{aligned} \frac{dz(t)}{dt} &= \mathrm{E}_\pi[q(z(t))] - z(t), \\ z(0) &= a, \end{aligned} \tag{6}$$

where $\mathrm{E}_\pi[q(z)]$ is the mean (quasi-) stationary queue length if $\{\hat{q}_n\}$ is identically equal to $z$. If $a < T_{\min}$, then, as long as $z(t) \le T_{\min}$, (6) reduces to

$$\begin{aligned} \frac{dz(t)}{dt} &= \mathrm{E}[q] - z(t), \\ z(0) &= a, \end{aligned} \tag{7}$$

where $\mathrm{E}[q]$ is the mean stationary queue length of a finite buffer queue (without the RED control).

Using the results of Kushner and Vazquez-Abad [14], we show in the Appendix that for any $0 < T < \infty$,

$$\sup_{0 \le t \le T} |\hat{z}^\beta(t) - z(t)| \xrightarrow{p} 0 \tag{8}$$

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

($\xrightarrow{p}$ denotes convergence in probability) as $\beta \to 0$, where $z(t)$ is the solution of (6). By the boundedness of $\hat{z}^\beta(t)$, this also implies the convergence of all the moments of $\hat{z}^\beta(t)$ to that of $z(t)$. In particular, we obtain that $\mathrm{Var}[\hat{z}^\beta(t) - z(t)] \to 0$ as $\beta \downarrow 0$ for any $t \geq 0$. At the end of Appendix A.1 we will also indicate how this result can be proved even when we include the idle time correction (as compared to (4)) in the dynamics of $\{\hat{q}_n^\beta\}$. From this result we can obtain the following useful approximations for our system. Let $\tau_{\min}^\beta$ and $\tau_{\max}^\beta$ be the first time $\{\hat{q}_n^\beta\}$ visits (crosses) $T_{\min}$ and $T_{\max}$ and let $\hat{T}_{\min}$ and $\hat{T}_{\max}$ be the corresponding quantities for the ODE solution $\{z(t)\}$. Then (8) provides us with $\tau_{\min}^\beta \xrightarrow{p} \hat{T}_{\min}$ and $\tau_{\max}^\beta \xrightarrow{p} \hat{T}_{\max}$ as $\beta \downarrow 0$. Hence, for small $\beta$, $\hat{T}_{\min}^\beta$ and $\hat{T}_{\max}^\beta$ can be used as approximations for our system. Also, by taking $z(n\beta)$ as an approximation of $\hat{q}_n^\beta$, we can approximate the distribution of $(q_n^\beta, c_n^\beta)$ by the stationary distribution of $(q_n(z(n\beta)), c_n(z(n\beta)))$. This also provides an approximation of the packet loss probability at the $n$th arrival instant. As mentioned earlier, because for small $\beta$ the system stays in a transient period for a long time (probably always because new flows arrive and terminate), these transient performance indices are in fact more important than the stationary distributions and moments.

This motivates us to obtain the solution of the ODE (6). In the Appendix we show that (6) has a unique, bounded, continuous solution, extendable to infinity for all initial conditions $0 \leq a \leq T_{\max}$. Also, we show that (6) has a unique global attractor $q^*$, which can easily be obtained by either iteratively finding the zero of the right side of (6), that is,

$$\mathrm{E}_\pi[q(z)] - z = 0, \tag{9}$$

or by numerically finding the solution of the ODE until it reaches the attractor (approximately). To obtain the solution of (6) for any initial $a$, we can use the standard numerical techniques. Observe that it is much easier (computationally) to obtain $\mathrm{E}_\pi[q(z(t))]$ for any $z(t)$ than to obtain the stationary mean queue length of the system. Furthermore, if $a < T_{\min}$, then to obtain the solution of the ODE until $z(t) = T_{\min}$, we can solve the simple ODE (7), which has an explicit solution

$$z(t) = ae^{-t} + \mathrm{E}[q](1 - e^{-t}).$$

If $a < T_{\min}$ and $\mathrm{E}[q] < T_{\min}$, then this solution is valid for (6) for all time and we need not consider the numerical techniques. Thus, this approximation can be used to compute various transient and stationary distributions (as explained below) of our system. However, because this solution never reaches $T_{\min}$ while $\{\hat{q}_n^\beta\}$ will almost surely cross $T_{\min}$ even for small $\beta$, we do not obtain an approximation of $\tau_{\min}^\beta$ from this solution (of course, $\tau_{\min}^\beta$ will be very large for small $\beta$). However, exactly for this case, using Freidlin–Wentzel theory in Dupuis and Kushner [9], one can obtain a sharper approximation of $\mathrm{E}[\tau_{\min}^\beta]$. (We have verified the assumptions for our system but do not provide the details in this article.)

Next, we explain how to approximate the stationary distribution and moments of the system parameters via the solution of (6). Let $\hat{q}^\beta$ denote $\{\hat{q}_n^\beta\}$ under stationarity. In the Appendix we show that $\hat{q}^\beta$ converges in total variation to $q^*$ as $\beta \downarrow 0$. Also,

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

because $\hat{q}^\beta$ is bounded, this implies $\mathrm{E}[(\hat{q}^\beta)^\alpha] \to (q^*)^\alpha$ for any $\alpha > 0$. Take the expectation of (4) under stationarity. Simplifying, we obtain $\mathrm{E}[\hat{q}^\beta] = \mathrm{E}_\pi[q^\beta]$. Thus, for small $\beta$ the stationary mean queue length can be approximated by $q^*$, which is much easier to compute. Similarly, we can obtain the stationary distribution of $(q_n^\beta, \hat{q}_n^\beta, c_n^\beta)$ as $\hat{q}^\beta \approx q^*$ and $(q_n^\beta, c_n^\beta)$ with the stationary distribution of $(q_n^\beta(q^*), c_n^\beta(q^*))$. From this we can also obtain the stationary probability of packet loss.

### 4.2. Quantifying the ODE Error

In this section we estimate the error in the ODE approximation via a diffusion approximation.

Although the following diffusion approximation should hold for the general case, our justification in the Appendix using results from Bucklew, Kurtz, and Sethares [6] is limited only to the ODE (7). The interarrival times can be general i.i.d. or MMPP, whereas the service times can be general i.i.d.

Define

$$X^\beta(t) = \frac{\hat{z}^\beta(t) - z(t)}{\sqrt{\beta}}.$$

Then $\{X^\beta(t), 0 \le t \le T\}$ converges weakly in $D[0,T]$ (the space of right-continuous real-valued functions with left-hand limits and with Skorohod topology) to a diffusion process $\{X(t), 0 \le t \le T\}$ satisfying

$$X(t) = \sigma B(t) - \int_0^t X(s)\, ds, \tag{10}$$

where $\{B(t), 0 \le t \le T\}$ is a zero mean standard Brownian motion and

$$\sigma^2 = \mathrm{Var}[q_0] + 2 \sum_{k=1}^{\infty} \mathrm{Cov}[q_0, q_k],$$

where $\{q_n\}$ is taken under stationarity.

For the exponential special case of Section 5 we can compute $\sigma^2$ more explicitly via

$$\sigma^2 = \boldsymbol{\nu} \cdot (\mathbf{u}^T + \mathbf{P}\mathbf{u}^T), \tag{11}$$

where $\boldsymbol{\nu} = [0 \cdot \pi(0), 1 \cdot \pi(1), \ldots, K \cdot \pi(K)]$ with $\pi(k)$ denoting the $k$th component of the stationary distribution of $\{q_n\}$ (in this case, the stationary distribution of an $M/M/1/K$ queue) and $\mathbf{P}$ is the transition matrix of $\{q_n\}$. The vector $\mathbf{u}$ is given by

$$\mathbf{u} = \sum_{k=0}^{\infty} (\mathbf{P}^k \mathbf{n}^T - \mathrm{E}[q]\mathbf{e}^T),$$

where $\mathbf{n} = [0, 1, \ldots, K]$ and $\mathbf{e} = [1, \ldots, 1]$.

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

Equation (10) is the well-known Langevin equation. It has a strong, unique (up to indistinguishability) solution $\{X(t)\}$. Also, $\{X(t)\}$ is a Gaussian diffusion with

$$\mathrm{E}[X(t)] = 0, \qquad \mathrm{Var}[X(t)] = \frac{\sigma^2}{2}(1 - e^{-2t}) \tag{12}$$

and has a stationary distribution that is Gaussian with mean zero and variance $\sigma^2 \int_0^\infty e^{-2s}\, ds = \sigma^2/2$. From (12) we observe that $\hat{z}^\beta(t) - z(t)$ is (for small $\beta$) approximately Gaussian with mean zero and variance proportional to $\beta$.

Although we do not have rigorous justification of the diffusion approximation of the error for the general ODE (6), the process $\{X^\beta(t)\}$ in this case should converge weakly to the diffusion satisfying the stochastic differential equation (SDE)

$$X(t) = -\int_0^t X(s)\, ds + \int_0^t \sigma(z(s))\, dB(s) + \int_0^t \frac{\partial}{\partial z} \mathrm{E}_\pi[q(z(s))] X(s)\, ds, \tag{13}$$

where

$$\sigma(z)^2 = \sum_{k=-\infty}^{\infty} \mathrm{Cov}[(q_k(z) - z), (q_0(z) - z)],$$

$\{q_k(z)\}$ is under stationarity. For the special case of $z < T_{\min}$, this equation reduces to (10). One can further carry on with the approximation procedure by observing that as $t \to \infty$, $z(t) \to q^*$ and, hence, for large $t$ we can approximate the solution of (13) by the solution of the SDE

$$X(t) = -\int_0^t X(s)\, ds + \int_0^t \sigma(q^*)\, dB_s + \int_0^t \frac{\partial}{\partial z} \mathrm{E}_\pi[q(z)]\Big|_{z=q^*} X(s)\, ds.$$

The justification of this can be obtained from Theorem 15 in Benveniste et al. [2, p. 335]. The two crucial assumptions of this theorem (A-i, A-ii) are satisfied by our system (see Appendix A.1) although we do not have a full justification. As a result of this theorem we obtain that $X^{\beta_n}(t_n)$ converges in distribution to a Gaussian random variable (r.v.) with mean 0 and variance $\sigma(q^*)^2/2$ whenever $\beta_n \downarrow 0$ and $t_n \to \infty$. This, in particular, implies that $X^\beta \sim N(0, \sigma(q^*)^2/2)$ for $\beta$ small, where $X^\beta$ denotes $X^\beta(t)$ under stationarity.

## 5. THE EXPONENTIAL SPECIAL CASE

In this section we present more specific forms of the previous results for the case where the arrival process is Poisson and the packet size distribution is exponential. In particular, we consider the conditional process $\{(q_n(\hat{q}), c_n(\hat{q}))\}$ (i.e., the process $\{(q_n, c_n)\}$ conditioned on the controlling $\hat{q}_n$ being identically equal to $\hat{q}$). The results obtained here will be used in the next section to compare with the simulations.

As noted earlier, when $\hat{q}_n < T_{\min}$, $c \equiv -1$ and the process $\{q_n\}$ is simply the queue length of an ordinary $M/M/1/K$ queue. Let $q$ be a random variable with the stationary distribution of $q_n$; that is,

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

$$p_k = \mathrm{P}\{q = k\} = \frac{\rho^k}{1 + \rho + \cdots + \rho^K},$$

and the generating function

$$g(z) = \sum_k p_k z^k = \frac{1-\rho}{1-\rho z} \frac{1-(\rho z)^{K+1}}{1-\rho^{K+1}},$$

where $\rho$ is the offered load of the system. From this we can derive the moments $\mathrm{E}[q]$ and $\mathrm{Var}[q]$. In particular, for the expectation we get

$$\mathrm{E}[q] = \frac{\rho}{1-\rho} - (K+1)\frac{\rho^{K+1}}{1-\rho^{K+1}}.$$

The case $T_{\min} \le \hat{q} \le T_{\max}$ requires a more detailed analysis. As described earlier, the packets are rejected with probability $p_{c_n}$, where

$$p_c = \frac{1}{C(\hat{q}) - c},$$

and

$$C(\hat{q}) = \frac{T_{\max} - T_{\min}}{\hat{q} - T_{\min}} \frac{1}{p_{\max}}.$$

(Here we use notation where the subscript of $p_c$ refers to the value of the counter rather than to the index $n$ of the arrival.) For $c_n > C(\hat{q}) - 1$ we define $p_c = 1$. The state space of $\{(q_n(\hat{q}), c_n(\hat{q}))\}$ is $[0, K] \times [0, \lceil C \rceil - 1]$, where $\lceil C \rceil$ is the smallest integer equal to or greater than $C$. Because the arrivals are Poisson, the stationary distribution of $\{(q_n(\hat{q}), c_n(\hat{q}))\}$ is the same as that of the corresponding continuous-time system with the state transition diagram shown in Figure 2. The transition rates for a generic state are shown on the right.

Note that for an infinite queue, $K = \infty$, the process $\{c_n\}$ alone constitutes a Markov process. Its transition diagram is given on the right-hand side of Figure 2. It is easy to see that the stationary distribution of $c_n$ (for an integer value of $C(\hat{q})$) is

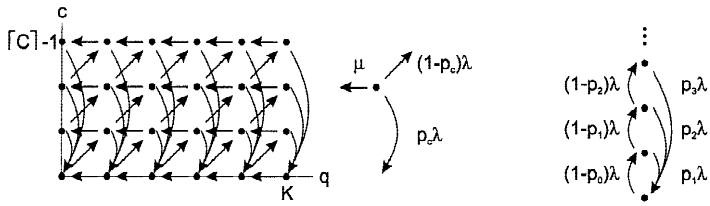



**FIGURE 2.** State transition diagrams for the processes $\{(q_n(\hat{q}), c_n(\hat{q}))\}$ and $\{(c_n(\hat{q}))\}$.

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

$$\pi_l = \mathrm{P}\{c_n = l\} = \frac{C - l}{\frac{1}{2}C(C+1)}, \qquad l = 0,1,\ldots,C-1.$$

Consequently, $p_l \pi_l$ is constant $2/C(C+1)$; that is, all values of $c_n$ are equally probable at the point where the rejection actually occurs. This was also noted in Stallings [24].

We are, however, more interested in the marginal distribution of $q_n(\hat{q})$, which can only be solved by first considering the full Markov process $\{(q_n(\hat{q}), c_n(\hat{q}))\}$. The stationary distribution of this two-dimensional process or the corresponding continuous-time process can be solved numerically, given the transition diagram of Figure 2. Let $(q(\hat{q}), c(\hat{q}))$ be a pair of random variables obeying the stationary distribution. From the two-dimensional stationary distribution, one can determine the marginal distribution of $q(\hat{q})$ and, in particular, its expectation $\mathrm{E}[q(\hat{q})]$. Finally, the attractor $q^*$ is obtained as the solution of (9), $q^* = \mathrm{E}[q(q^*)]$.

## 6. NUMERICAL RESULTS

Here we illustrate the use of some of the results obtained in the earlier sections. The system parameters are chosen to be the following: $K = 50$, $T_{\min} = 20$, $T_{\max} = 40$, and $p_{\max} = 1$. The packets are assumed to arrive according to a Poisson process with arrival rate $\lambda$ and the packet lengths are assumed to be exponentially distributed with parameter $\mu$. Unless otherwise stated, the load of the system is assumed to be $\rho = \lambda/\mu = 1.5$. For a system of this size the direct discretization approach (by discretizing $\hat{q}$ to obtain a countable state Markov chain) would not allow us to investigate the performance for the values of $\beta$ that have been suggested in the literature without running into problems with the size of the state space of the discretized process. Thus, in the following, we compare the accuracy of our approximations against simulated results.

First, we illustrate the effect of the RED control on the performance of the queue. Figure 3 plots the attractor $q^*$ as a function of the load $\rho$ obtained as a

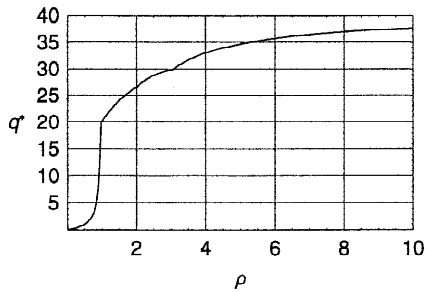



**FIGURE 3.** Value of the attractor $q^*$ as a function of the load $\rho$.

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

solution to (9). We showed in Section 4 that $\mathrm{E}_\pi[q^\beta] = \mathrm{E}_\pi[\hat{q}^\beta] \to q^*$ as $\beta \downarrow 0$. Thus, for small $\beta$, $q^*$ approximates the mean stationary queue length. Hence, from Figure 3 we can clearly see how the average queue length increases until it reaches the threshold $T_{\min}$, at which point the rapid increase is being cut down by the RED control algorithm.

Next, we show the convergence of the sample paths of $\hat{z}^\beta(t)$ when $\beta \downarrow 0$ (see Fig. 4) to the solution of the ODE (6) (see Fig. 5, left, lower curve). The solution of (6) is obtained numerically by using the formulas provided in Section 5. Floyd and Jacobson [11] suggest $\beta$ close to 0.001 (cf. the center figure in Fig. 4). From Figures 4 and 5 we can also observe that $\tau^\beta_{\min}$ can be well approximated by $\hat{T}_{\min}$. In our simulations, $T_{\max}$ was never reached. Thus, it appears that $\tau^\beta_{\min}$ is the critical parameter to estimate. The left-hand side of Figure 5 (upper curve) also shows $\mathrm{E}_\pi[q(\hat{q}_n)]$ for small $\beta$. In the simulated results, the convergence of $\mathrm{E}[q^\beta_n]$ to $\mathrm{E}[q(q^*)]$ is slower because of averaging over a large number of samples.

Then we examine the discarding behavior of the RED algorithm by looking at the time evolution of the (quasistationary) time-dependent loss probabilities due to buffer overflow and rejection due to the RED algorithm as obtained from the (quasi)-stationary distribution of the process $\{(q_n(\hat{q}), c_n(\hat{q}))\}$. The time dependency arises from the fact that $\hat{q}$ evolves in time according to the solution of the ODE (6). This is shown in Figure 6. The figure on the left is the ODE solution and the figure on the right gives the simulated results for a system with $\beta = 0.001$. The simulated results are again obtained in the same manner as earlier by averaging over 5000 sample paths. The simulation graph also includes a plot of the time evolution of $\{\hat{q}_n\}$ taken as an average over the sample paths. The results show how in the beginning, when $\hat{q} \le T_{\min}$, the queue behaves as the uncontrolled queue and the loss probabilities due to overflow are high. However, as the control starts affecting (due to increasing $\hat{q}_n$), it reduces the overflow probability, thus also reducing the probability of the global synchronization effect. This also verifies our decomposition heuristic for the probability of packet discarding.

As can be seen from the simulated results (see Figs. 5 and 6), the actual system with finite $\beta$ has an oscillatory behavior. This can be explained as follows. We have considered the $\{\hat{q}_n\}$ process in the asymptotic limit $\beta \to 0$. The changes in the process are then so slow that the $\{q_n\}$ process can be assumed to be in the equilibrium. However, the $\{q_n\}$ process follows its own dynamics, and when a change in $\{\hat{q}_n\}$ takes place, it takes a finite time before the equilibrium is attained. The related time constant is independent of $\beta$. Thus, when $\beta$ increases, the slowness of the $\{q_n\}$ process becomes, in relative terms, more important. In this regime we see overshooting and oscillations typical of control systems with time lag in the control loop.

We have also computed the approximate variance of $\hat{z}^\beta(t)$ for $t \le \tau^\beta_{\min}$ from (11). For our example, $\sigma^2 = 150$. Then, the variance of $\hat{z}^\beta(t)$ is $150(1 - e^{-2t})/(2)(10^3)$. If $q^*$ was less than $T_{\min}$ (it is not in our case), then $\sqrt{150/(2)(10^3)} \approx 0.27$ would be the standard deviation of the error under stationarity. However, this deviation cannot be seen in the example figures because the deviation is much less in magnitude than the amplitude of the oscillations. Of course, even for our present

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

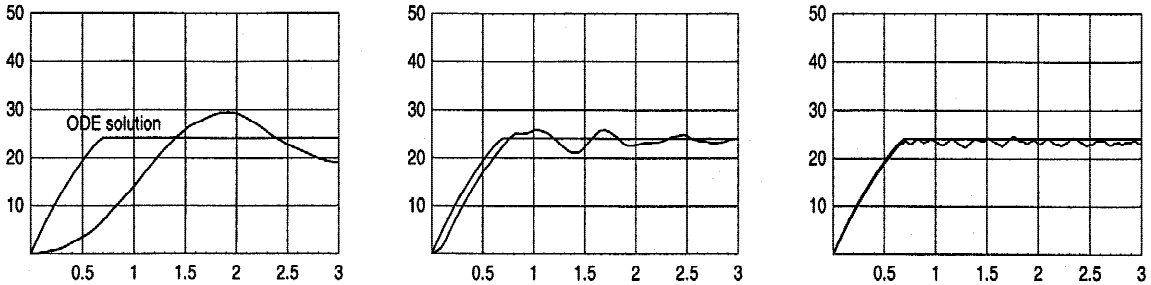



**FIGURE 4.** Convergence of sample paths of $\hat{q}_n^{\beta}$ when $\beta = 0.01$ (left), $\beta = 0.001$ (center), and $\beta = 0.0001$ (right). For reference purposes, the solution of the ODE (6) is shown in the same figures.

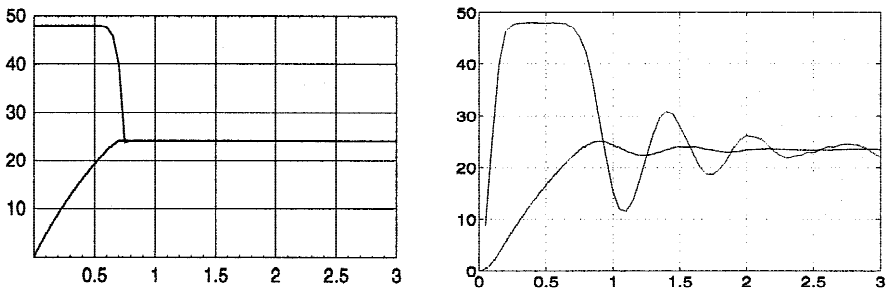



**FIGURE 5.** Time evolution of $\hat{q}_n$ and $q_n$ as obtained from the solution of the ODE (6) (left) and the same results from the simulated system (right) with $\beta = 0.001$.

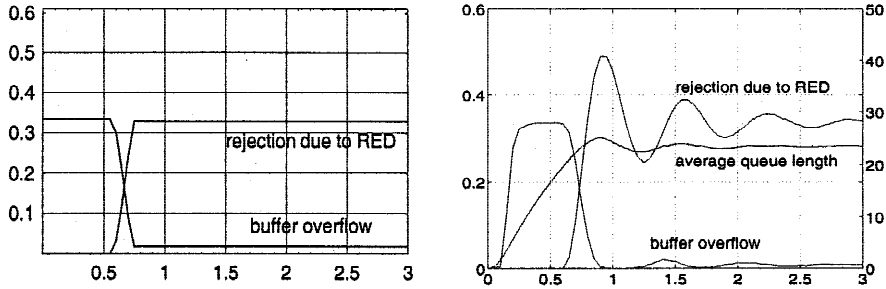



**FIGURE 6.** Analytical results (left) and simulation results (right) with $\beta = 0.001$ on the time evolution of the packet loss probabilities due to overflow or rejection.

example, the transient and stationary variances of $\hat{z}^{\beta}(t)$ could be computed along the same line using (13), but at a substantially higher computational cost for the transient variance.

*Acknowledgments*

The work of the first author was done while visiting Helsinki University of Technology. He thanks Prof. Virtamo and HUT for their kind hospitality. The work of the third author was financially supported by the Nokia Foundation.

*References*

1. Asmussen, S. (1987). *Applied probability and queues*. New York: John Wiley & Sons.
2. Benveniste, A., Metivier, M., & Priouret, P. (1990). *Adaptive algorithms and stochastic approximation*. New York: Springer-Verlag.
3. Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., & Weiss, W. (1998). *An architecture for differentiated services*. Internet RFC-2475.
4. Bonald, T. & May, M. (1999). Drop behavior of RED for bursty and smooth traffic. *Proceedings of the 7th International Workshop on Quality of Service (IWQoS'99)*, London, pp. 257–259.
5. Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S., Estrin, D., et al. (1998). *Recommendations on queue management and congestion avoidance in the Internet*. Internet RFC-2309.
6. Bucklew, J.A., Kurtz, T.G., & Sethares, W.A. (1993). Weak convergence and local stability properties of fixed step size recursive algorithms. *IEEE Transactions on Information Theory* 30: 966–978.
7. Clark, D.D. & Fang, W. (1998). Explicit allocation of best effort packet delivery service. *IEEE/ACM Transactions on Networking* 6: 362–373.
8. Courtois, P.J. (1977). *Decomposability: Queuing and computer system applications*. New York: Academic Press.
9. Dupuis, P. & Kushner, H.J. (1987). Asymptotic behavior of constrained stochastic approximation via the theory of large deviations. *Probability Theory and Related Fields* 75: 223–244.
10. Feng, W., Kandlur, D., Saha, D., & Shin, K. (1999). A self-configuring RED gateway. *Proceedings of INFOCOM'99*, pp. 1320–1328.
11. Floyd, S. & Jacobson, V. (1993). Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking* 3: 397–413.
12. Hashem, E. (1989). Analysis of random drop for gateway congestion control. Report TR-465, MIT Laboratory of Computer Science, Boston, MA.

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

13. Kalashnikov, V. (1994). *Topics on regenerative queues*. Boca Raton, FL: CRC Press.
14. Kushner, H.J. & Vazquez-Abad, F.J. (1996). Stochastic approximation methods for systems over an infinite horizon. *SIAM Journal on Control and Optimization* 34: 712–756.
15. Kushner, H.J. & Yin, G.G. (1997). *Stochastic approximation algorithms and applications*. New York: Springer-Verlag.
16. Lin, D. & Morris, R. (1997). Dynamics of random early detection. *Proceedings of SIGCOMM '97*.
17. Macchi, O. (1995). *Adaptive processing: The least mean squares approach with applications in transmission*. New York: John Wiley & Sons.
18. Mankin, A. & Ramakrishnan, K.K. (eds.). (1991). *Gateway congestion control survey*. Internet RFC-1254.
19. May, M., Bolot, J., Jean-Marie, A., & Diot, C. (1999). Simple performance models of differentiated services schemes for the Internet. *Proceedings of INFOCOM'99*, pp. 1385–1394.
20. Ott, T., Lakshman, T., & Wong, L. (1999). SRED: Stabilized RED. *Proceedings of INFOCOM'99*, pp. 1346–1355.
21. Peeters, S. & Blondia, C. (1999). A discrete time analysis of random early detection with responsive best effort traffic. *Proceedings of 7th IFIP ATM & IP Workshop*, Antwerp.
22. Sharma, V. (1995). Reliable estimation via simulation. *Queuing Systems* 19: 169–192.
23. Sharma, V. (1998). Some limit theorems for regenerative queues. *Queuing Systems* 30: 341–363.
24. Stallings, W. (1997). *High-speed networks—TCP/IP and ATM design principles*. Englewood Cliffs, NJ: Prentice-Hall.
25. Zhang, L. (1989). A new architecture for packet switching network protocol. Report TR-455, MIT Laboratory of Computer Science, Boston, MA.

## APPENDIX

### A.1. ODE Approximation

In this Appendix we supply the proofs for the asymptotic approximations in Section 4.1. First we show that $\hat{z}^\beta$ converges weakly to the unique bounded solution of the ODE (6) in space $D[0,T]$ with the Skorohod topology. For this we use the result in Kushner and Vazquez-Abad [14]. In the following we show that our system satisfies the assumptions in this reference. At the end of this section we will discuss the solution of the ODE (6).

In the notation of Kushner and Vazquez-Abad [14], we take $\beta = \epsilon$, $\hat{q}_n^\beta = \theta_n^\epsilon$, $Y_n^\epsilon = q_{n+1}^\beta - \hat{q}_n^\beta$, and $(q_n^\beta, c_n^\beta) = \xi_n^\epsilon$. Now we will further assume that the service times of the packets are exponentially distributed (see comments later in this Appendix for generalization). Then, the boundedness of $q_n^\beta$ and $\hat{q}_n^\beta$ provides the uniform integrability of $\{Y_n^\epsilon\}$. For tightness of $\{\xi_n^\epsilon, \theta_n^\epsilon\}$ we need the tightness of $\{(q_n^\beta, \hat{q}_n^\beta, c_n^\beta), \beta > 0, n \geq 0\}$. Now, we further assume that $c_n^\beta$ is upper-bounded by some (arbitrarily large) constant $c$. In a practical system this will always be required. Given $\hat{q}_n^\beta$ fixed (say $\equiv \hat{q}$) $(q_n^\beta, c_n^\beta)$ is a Markov chain ($\stackrel{\triangle}{=} \{(q_n(\hat{q}), c_n(\hat{q}))\}$) with a transition matrix independent of $\beta$. Thus, condition (3.6) in Kushner and Vazquez-Abad [14] is satisfied. For condition (3.7), because $q_n^\beta$ and $c_n^\beta$ are integer-valued, we only have to prove the continuity of the transition probabilities

$$\mathrm{P}\{q_{n+1}^\beta = q', c_{n+1}^\beta = c' \,|\, q_n^\beta = q, c_n^\beta = c, \hat{q}_n^\beta \equiv \hat{q}\} \tag{A.1}$$

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

as a function of $\hat{q}$. From the algorithm in Section 2 we see that the probability of rejection of an arrival is a continuous function of $\hat{q}$, except possibly at $\hat{q} = T_{\max}$. If $p_{\max} = 1$, then we have continuity at $T_{\max}$ also. The continuity has been recommended from a practical point of view also in "Gentle" RED. Another possibility is to keep $p_{\max} < 1$ at $T_{\max}$, but then linearly increase the probability of rejection from $p_{\max}$ to 1 for $\hat{q} \in [T_{\max}, T_{\max} + \delta]$ for some small $\delta > 0$. This will not change the dynamics of the system much. Denoting by $\mu(\cdot|\hat{q})$ the stationary distribution of (A.1) (which can be easily shown to exist uniquely), the boundedness of the state space of this Markov chain provides us (3.8) in Kushner and Vazquez-Abad [14]. Thus, all of the assumptions for Theorem 3.1 in Kushner and Vazquez-Abad [14] are satisfied. This provides us with the tightness of the processes $\{\hat{z}^{\beta}(t), t \in [0,T]\}, \beta < 1$. We show later in this Appendix that (6) has a unique solution for each initial condition. This implies that $\{\hat{z}^{\beta}(t), t \in [0,T]\}$ converges weakly to this unique solution. This, of course, also implies that the first visit times to $T_{\min}$ and $T_{\max}$ by $\hat{z}^{\beta}(t)$ converge in distribution to the corresponding visit times of the solution of the ODE (6).

The above convergence holds if the service time distributions are phase type, using the same arguments (now $\xi_n^{\epsilon}$ will also include the phase of the packet being transmitted). One can try to extend the results to the case of general service time distributions by taking into account the residual service time of the packet being transmitted, but we have not studied this extension. This generalization is more of theoretical interest only, because any distribution can be arbitrarily closely approximated by phase-type distributions. Also, in the context of the ODE (7), this generalization is obtained directly from Bucklew et al. [6].

From Kushner and Vazquez-Abad [14] we also obtain under the same assumptions that $\hat{q}^{\beta}(\beta t_{\beta} + t)$ converges in distribution to $q^*$ if $\beta t_{\beta} \to \infty$ as $\beta \downarrow 0$. By taking $t_{\beta}$ large enough, we can ensure that $\hat{q}^{\beta}(\beta t_{\beta})$ is within $\delta$ variation distance from its stationary distribution $\pi^{\beta}$ for $\delta$ arbitrarily small. This ensures that

$$\lim_{\beta \downarrow 0} \sup_{x} |\pi^{\beta}(x) - \delta_{q^*}(x)| < \delta,$$

where $\delta_{q^*}$ is the probability measure with an atom of weight 1 at $q^*$. Taking $\delta \downarrow 0$, we obtain that $\sup_x |\pi^{\beta}(x) - q^*| \to 0$ as $\beta \downarrow 0$. Because $\hat{q}^{\beta}$ is bounded, this ensures that all stationary moments of $\hat{q}^{\beta}$ also converge to that of $q^*$.

If we include the idle time correction to (4), then we can use the result of Kushner and Vazquez-Abad [14] as follows. For simplicity we only discuss the case of exponential service times. Define $m_n$ as the $m$ used in the definition of RED in Section 2 at the time of the $n$th arrival. Include $m_n$ as one of the components in $\{\xi_n^{\epsilon}\}$. Given $q_n^{\beta}$, $m_{n+1}^{\beta}$ is independent of everything else. Also, $m_{n+1}^{\beta} \leq q/\bar{s}$, where $\bar{s} > 0$. Thus, if $\mathrm{E}[a^{1+\alpha}] < \infty$ for some $\alpha > 0$, $\{m_n^{\beta}\}$ will be uniformly integrable. Continuity of the transition matrix of $\{(q_n^{\beta}(\hat{q}), \hat{q}_n^{\beta}(\hat{q}), c_n^{\beta}(\hat{q}))\}$ as a function of different parameters can be easily seen. The existence and uniqueness of the stationary distribution of this process follows in the same way as that of the

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

process without $m_n$ because the regeneration epochs for both processes are the same. Also, tightness of the stationary distribution holds by the boundedness of $q_n^\beta, c_n^\beta, m_n^\beta \leq a/\bar{s}$ and $\mathrm{E}[a^{\alpha+1}] < \infty$. Thus, we have verified all the assumptions required.

Now we show that the ODE (6) has a unique solution for each initial condition $a$ and the solution extends up to infinity and is bounded. We also show that this ODE has a unique global attractor to which each such solution converges. The justification for convergence to the solution of (6) provided above is limited to the case where the service times are phase type and the counter $c_n$ is upper bounded by a constant $c$. Thus we limit our discussion to this case. Under these conditions the process $\{(q_n(\hat{q}), c_n(\hat{q}))\}$ for $\hat{q} \leq T_{\max}$ is a finite-state irreducible Markov chain independent of $\beta$. Also, as observed earlier, its transition matrix is a continuous function of $\hat{q}$ (for continuity at $\hat{q} = T_{\max}$, we need $p_{\max} = 1$, which we assume here). Then, by the finite-state space, the stationary distribution of this Markov chain and, hence, also $\mathrm{E}_\pi[q_n(\hat{q})]$ are continuous functions of $\hat{q}$ for all $\hat{q} \leq T_{\max}$. We need Lipschitz continuity with respect to $\hat{q}$ for obtaining a unique local solution of the ODE. For this we modify the rejection probability of RED around $T_{\min}$ and $T_{\max}$ a little, to make it differentiable. Then $\mathrm{E}_\pi[q_n(\hat{q})]$ also becomes differentiable. This implies the Lipschitz continuity. This provides a local solution of the ODE (6) for each initial condition $0 \leq a \leq T_{\max}$. Next, observe that $\mathrm{E}_\pi[q_n(T_{\max})] = 0$ when $p_{\max} = 1$ (or at $T_{\max} + \delta$ in the case of the other linear approximation mentioned) and that $\mathrm{E}_\pi[q_n(\hat{q})]$ is a nonincreasing continuous function of $\hat{q}$ with $0 < \mathrm{E}_\pi[q_n(0)] < \infty$. Thus (see Fig. A.1) the equation $\mathrm{E}_\pi[q_n(\hat{q})] = \hat{q}$ has a unique solution $q^*$. This is the unique attractor of the ODE (6). Also, observe that when $z(t) = \hat{q} < q^*$, then $\mathrm{E}_\pi[q_n(\hat{q})] - \hat{q} > 0$ and the solution $z(t)$ of (6) strictly increases. Similarly, when $z(t) = \hat{q} > q^*$, then $\mathrm{E}_\pi[q_n(\hat{q})] - \hat{q} < 0$ and, hence, $z(t)$ strictly decreases. Thus, irrespective of $z(0) = a \leq T_{\max}$, $z(t)$ converges monotonically to $q^*$. This, along with the existence of (locally) a solution of (6) shown earlier, provides the existence and uniqueness of

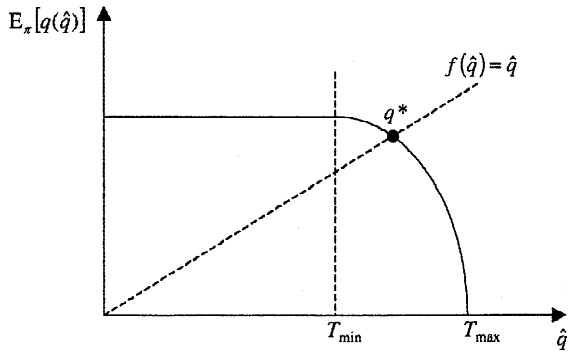



**FIGURE A.1.** Attractor of the ODE.

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press

a solution of (6) for each $a \leq T_{\max}$ and convergence to $q^*$ (thus, also the solution exists for all time and is bounded). Furthermore, the solution is a continuous function.

### A.2. Diffusion Approximation

Next we justify the diffusion approximation to the error in the ODE approximation only in the limited context of (7). We use Theorem 2 of Bucklew et al. [6]. In their terminology, $\beta = \mu$, $\hat{q}_n^\beta = W_k$, $q_{n+1} = Y_k$, $H(W_k, Y_k) = Y_k - W_k$, and there are no $\psi_k$ and $u_k$. Then, by the boundedness of $q_n$ and $\hat{q}_n^\beta$, it is easy to verify assumptions C1, C2, and C3 in [6]. Also,

$$L_\mu(t) = \sum_{k=0}^{\lfloor t/\beta \rfloor - 1} (q_{k+1} - \mathrm{E}_\pi[q])\sqrt{\beta};$$

hence C4 in [6] follows from the functional central limit theorem in Section 3.

https://doi.org/10.1017/S0269964802163078 Published online by Cambridge University Press