A Singular Perturbation Approach to Analysing a RED queue

Eitan Altman, K. Avrachenkov and B. J. Prabhu
INRIA-Sophia Antipolis
06902 Sophia-Antipolis, France

Abstract

Several Active Queue Management (AQM) techniques for routers in the Internet have been proposed and studied during the past few years. One of the widely studied proposals, Random Early Detection (RED), involves dropping an incoming packet with some probability based on the estimated average queue length at the router. The analytical approaches to obtaining average drop probabilities in a RED enabled queue have been either based on using the instantaneous queue size for calculating the drop probability or have considered averaging with a fluid approximation. In this paper, we use a singular perturbation based approach to analyse a RED enabled queue with drop probabilities based on the estimated average queue size as has been proposed in the standard RED. The singular perturbation approach is motivated by the fact that the instantaneous and the estimated average queue lengths evolve at two different time scales. We present an analytical method to calculate the average queue size and the average drop probability for the non responsive flows. We also provide analytical expressions for the Poisson arrivals and exponential service times case. Our model is derived under several approximations, and is validated through simulations.

Key words - Singular perturbation, RED.

I. Introduction

Active Queue Management (AQM) involves controlling the congestion at the routers in conjunction with the end-to-end congestion control which is based on controlling congestion at the source [1]-[2]. The use of a droptail queue for closed loop sources (e.g, TCP) has been observed to result in inefficient use of the outgoing link capacity at a bottleneck router [3]. Under heavy traffic conditions, packet losses in a droptail queue have been observed to lead to synchronized evolution of windows at different TCP sources which use the same bottleneck link. This synchronized increase and decrease in window size at different sources leads to inefficient use of the outgoing link capacity at the bottleneck link [3].

One of the proposals, Random Early Detection (RED), to overcome this synchronized window evolution involves dropping packets at random before the congestion actually sets in [1]. Two of the design aims of RED are to accept occasional bursts and to maintain a reasonable average queue length when the system is heavily loaded with the objective of efficient use of the system capacity. Towards this end, the routers maintain a variable corresponding to the exponentially weighted moving average of the instantaneous queue length. An incoming packet is dropped with some probability which is a function of this estimated average queue length. So far, the various approaches to analysing a RED enabled queue have either involved the use instantaneous queue size, instead of the estimated average, to compute the drop probabilities [4] or have considered averaging with a fluid approximation [5],[6]. In [7], the authors compute the joint distribution of the instantaneous and the averaged queue length of an $M/M/1/K$ queue as a solution of a set of differential equations. An analytical expression for the case of buffer size 1 and 2 is also provided. However, for other cases they present numerical methods to solve the equations.

The aim of our study is to use singular perturbation techniques [8]-[9] to provide an analytical expression for computing the drop probabilities and average queue length in a RED enabled queue with averaging as has been proposed in RED. Furthermore, we make use of the transition probability matrices in order to preserve the stochastic nature of the system as opposed to a fluid approximation which is deterministic. The use of singular perturbation technique is motivated by the fact that two different time scales are involved in the evolution of the instantaneous and the average queue length. Indeed, for very small values of the

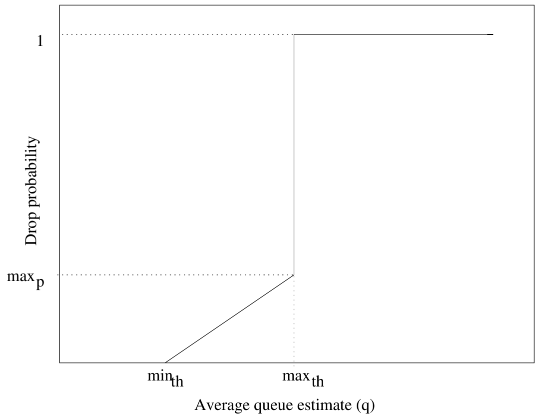

Fig. 1. RED drop probability function

averaging parameter, the average queue length can be assumed to vary much more slowly compared to the instantaneous queue length. We analyse an open loop system in which packets arrive at a RED enabled queue in batch-poisson process. The study of close loop systems involving TCP clients will be considered in the ensuing work.

The rest of the paper is organized as follows. In Section II we present the system model and the analytical method to compute the drop probabilities and the average queue length. In Section III we compare the results obtained using the analysis and with simulations. Finally, in Section IV we present the conclusions.

II. System Model and Analysis

We consider a RED enabled queue at which packets arrive according to a Poisson process. The maximum buffer size at this queue is assumed to be B. The packet service times are assumed to be exponential. Let Q_n and $\hat{q}_n$ be the instantaneous and estimated average queue lengths, respectively, before the n^{th} arrival. The estimated average queue length, $\hat{q}_n$, is an exponentially weighted moving average of the instantaneous queue length, Q_n, and is given by :

$$\hat{q}_{n+1} = (1-a)\hat{q}_n + aQ_{n+1}, \tag{1}$$

where the averaging parameter, a, is a small number. The drop probability for the $(n+1)^{th}$ arrival is a given function of $\hat{q}_{n+1}$. Figure 1 shows the drop function proposed in the basic RED algorithm [1].

With this formulation, the couple (q, Q) is a Markov chain. Its analysis is difficult, partly because the support of q is a countable subset of the interval $[0, B]$ where B is the maximum queue size.

We propose an approximation that reduces the support of q to a finite set given by $G = \{0, ..., mB\}$, where m is some integer. In the following discussion, $\hat{q}_n$ will refer to the discretized estimated average queue length. We assume below that a is sufficiently small such that $amB < 1$. The modified update equation is then written as

$$\hat{q}_{n+1} = (1-a)\hat{q}_n + amQ_{n+1}, \quad q_n \in \{0, 1, ..., mB\}. \tag{2}$$

Then we replace equation(2) by the following approximation:

$$\hat{q}_{n+1} = \hat{q}_n + C_n, \tag{3}$$

where

$$C_n = \begin{cases} 1 & \text{w.p. } a(mQ_{n+1} - \hat{q}_n)^+ \ , \\ -1 & \text{w.p. } a(\hat{q}_n - mQ_{n+1})^+ \ , \\ 0 & \text{otherwise} \ , \end{cases}$$

with $x^+ = max(x, 0)$.

We note that with this definition, we obtain

$$\begin{aligned} E[\hat{q}_{n+1} - \hat{q}_n | \hat{q}_n, Q_{n+1}] &= 1 \cdot a(mQ_{n+1} - \hat{q}_n)^+ + (-1) \cdot a(\hat{q}_n - mQ_{n+1})^+ \\ &= a(mQ_{n+1} - \hat{q}_n), \end{aligned} \tag{4}$$

so that (2) becomes the mean field of (3). With this approximation we only allow for transitions to neighbouring states.

Our interest is in obtaining the steady state distribution of the couple $(\hat{q}, Q)$ which will allow us to study the various performance measures (e.g., average queue length, average drop probability, etc.) of the system.

Let $\tilde{P}$ denote the transition probability matrix for the two dimensional Markov chain (q, Q) with $\tilde{\pi}(a)$ as its stationary distribution. Let A_i be defined as

$$A_i := \{a_{jk}\}_i = P(Q_{n+1} = k | Q_n = j, \hat{q}_n = i), \qquad \forall i \in 0, 1, ..., mB. \tag{5}$$

A_i is a $(B+1) \times (B+1)$ matrix which denotes the transition probability matrix of instantaneous queue length, Q, for a given value of the average queue length, $q = i$. Let C_i^+, C_i^-, C_i^0 denote matrices given by

$$\begin{aligned}
C_i^+ := \{c_{jj}\} &= p(\hat{q}_{n+1} = i+1 | Q_{n+1} = j, \hat{q}_n = i) \\
&= a(m \cdot j - i)^+, \\
C_i^- := \{c_{jj}\} &= p(\hat{q}_{n+1} = i-1 | Q_{n+1} = j, \hat{q}_n = i) \\
&= a(i - m \cdot j)^+, \\
C_i^0 := \{c_{jj}\} &= p(\hat{q}_{n+1} = i | Q_{n+1} = j, \hat{q}_n = i) \\
&= 1 - a(m \cdot j - i)^+ - a(i - m \cdot j)^+, \\
& \quad \forall i \in \{0, 1, ..., mB\} \qquad \forall j \in \{0, 1, ..., B\}.
\end{aligned}$$

C_i^+i, C_i^-, and C_i^0 are $(B+1) \times (B+1)$ matrices which denote the transition probability of $q_{n+1} = q_n + 1$, $q_{n+1} = q_n - 1$, and $q_{n+1} = q_n$, respectively, when $Q_{n+1} = j$. We note that C_i's are diagonal matrices which satisfy the equality

$$C_i^+ + C_i^- + C_i^0 = I. \tag{6}$$

Using the identity

$$\begin{aligned}
p(\hat{q}_{n+1}, Q_{n+1} | \hat{q}_n, Q_n) &= p(\hat{q}_{n+1} | \hat{q}_n, Q_n, Q_{n+1}) \cdot p(Q_{n+1} | \hat{q}_n, Q_n) \\
&= p(\hat{q}_{n+1} | \hat{q}_n, Q_{n+1}) \cdot p(Q_{n+1} | \hat{q}_n, Q_n),
\end{aligned}$$

we can write $\tilde{P}$ as

$$\tilde{P} = \begin{bmatrix}
A_0(C_0^- + C_0^0) & A_0 C_0^+ & 0 & \dots & \dots \\
A_1 C_1^- & A_1 C_1^0 & A_1 C_1^+ & 0 & \dots \\
0 & A_2 C_2^- & A_2 C_2^0 & A_2 C_2^+ & \dots \\
\vdots & \vdots & \ddots & \vdots & \vdots \\
\vdots & \vdots & \vdots & \ddots & \vdots \\
0 & \dots & \dots & A_{mB} C_{mB}^- & A_{mB}(C_{mB}^0 + C_{mB}^+)
\end{bmatrix}. \tag{7}$$

The diagonal rows entries of $\tilde{P}$ are matrices which give the transition probability matrix of Q for a given value of $\hat{q}_n$, and $\hat{q}_{n+1} = \hat{q}_n$ whereas the off-diagonal entries give the transition probability matrix of Q for a given value of $\hat{q}_n$ and $\hat{q}_{n+1} = \hat{q}_n + 1$, or $\hat{q}_{n+1} = \hat{q}_n - 1$. Using equation(6), we can rewrite $\tilde{P}$ as

$$\tilde{P} = P + PC, \tag{8}$$

where P is given by

$$P = \begin{bmatrix}
A_0 & 0 & \dots & 0 \\
0 & A_1 & \dots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & \dots & 0 & A_{mB}
\end{bmatrix},$$

and C is given by

$$C = \begin{bmatrix}
-C_0^+ & C_0^+ & 0 & \dots \\
C_1^- & -(C_1^- + C_1^+) & C_1^+ & \dots \\
\vdots & \vdots & \ddots & \vdots \\
0 & \dots & C_{mB}^- & -C_{mB}^-
\end{bmatrix}.$$

We note that $\tilde{P}$ is a function of the averaging parameter a. In [10], the recommended value of a is of the order of 0.002. This suggests that the averaging parameter is, in general, small. This assumption on a allows us to obtain the steady state probabilities of $\tilde{P}$ using a singular perturbation approach which we outline below and which is computationally less expensive than algorithms to compute the invariant vectors of matrices.

To use the singular perturbation approach, we first note that P is the matrix containing the transition probability matrices of Q for a given value of q when the averaging parameter $a = 0$. Thus, P has mB ergodic classes each corresponding to a particular value of $\hat{q}$. The stationary distribution of each of these classes can be computed separately. This system of decomposable Markov chain is the unperturbed system which can be analysed separately for different values of q. In general, the stationary distribution of the unperturbed system (i.e., $a = 0$) is not the same as the stationary distribution of the original system when $a \to 0$ [11]. This motivates us to use singular perturbation technique to determine the stationary distribution of our system when $a \to 0$.

Later, we also show that the distributions of the instantaneous queue length and the average queue length can be computed for the general case (i.e. $a \nrightarrow 0$). We note that the use of the approximation given in (3) has resulted in a level dependent Quasi-Birth-Death (QBD) type of transition matrix which is as given in (7). We can now use the algorithm, given in [12], for computing the steady state probability vector for level dependent QBD.

A. *Limiting case*

In this subsection we obtain the steady state probability vector of $\tilde{P}$ when $a \to 0$. We note that all the elements of the matrices C_i^+ and C_i^- have a as a common multiple and thus we can replace C_i^+ by aD_i^+ and C_i^- by aD_i^- where D_i^+ and D_i^- have elements independent of a. Thus we can rewrite equation (8) as

$$\tilde{P} = P + aPD. \tag{9}$$

This is the standard formulation for the singularly perturbed Markov chains [8]. The singular perturbation approach allows us to find the stationary distribution of $\tilde{P}$ as $\lim_{a\to 0} \tilde{\pi}(a)$.

We present the method for finding the limiting stationary distribution of $\tilde{P}$, the details of which can be found in [8], [9]. Let π_i denote the stationary distribution of A_i, and l et V be a $(mB+1)\times(mB+1)(B+1)$ matrix such that in the i^{th} row the entries, corresponding to the columns $\hat{q}_n = i$, are given by the rows of π_i, and is zero elsewhere. π_i is the stationary distribution of the unperturbed Markov chain of Q when packets are dropped with a probability depending on $\hat{q} = i$. Let W be a $(mB+1)(B+1) \times (mB+1)$ matrix such that in the i^{th} column the entries corresponding to the rows $\hat{q}_n = i$ are 1, and is zero elsewhere.

$$V = \begin{bmatrix} \pi_0 & \underline{0} & \dots & \underline{0} \\ \underline{0} & \pi_1 & \dots & \underline{0} \\ \vdots & \vdots & \ddots & \vdots \\ \underline{0} & \dots & \dots & \pi_{mB} \end{bmatrix},$$
$$W = \begin{bmatrix} \underline{1}' & \underline{0} & \dots & \underline{0} \\ \underline{0} & \underline{1}' & \dots & \underline{0} \\ \vdots & \vdots & \ddots & \vdots \\ \underline{0} & \dots & \dots & \underline{1}' \end{bmatrix},$$

where $\underline{0}$ and $\underline{1}$ are $(B+1)\times 1$ vectors of 0 and 1, respectively.

Let S denote the generator matrix of the aggregated Markov chain. When a goes to 0, the stochastic sequence $\hat{q}_n$ converges weakly to a Markov chain induced by the aggregated transition matrix. Then S is given by

$$S = VPDW.$$

Since π_i is the stationary distribution of A_i, VP reduces to V, i.e.,

$$S = VDW. \tag{10}$$

Denote

$$
\begin{aligned}
f_i^+ &= \pi_i D_i^+ \underline{1}' && (11)\\
&= \sum_{j=\lceil \frac{i}{m} \rceil}^{B} \pi_i(j)(mj-i). && (12)\\
f_i^- &= \pi_i D_i^- \underline{1}' && (13)\\
&= \sum_{j=0}^{\lfloor \frac{i}{m} \rfloor} \pi_i(j)(i-mj). && (14)\\
& && (15)
\end{aligned}
$$

where $\pi_i(j)$ is the j^{th} element of π_i.

Using equations (11) and (13) we can rewrite S as

$$
S = \begin{bmatrix} -f_0^+ & f_0^+ & 0 & \dots \\ f_1^- & -(f_1^- + f_1^+) & f_1^+ & \dots \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & f_{mB}^- & -f_{mB}^- \end{bmatrix}.
$$

The stationary distribution of this generator matrix can be obtained by solving $\gamma S = 0$. The stationary distribution, γ, is given by

$$
\gamma := \{\gamma(i)\} = \gamma(0) \frac{\prod_{j=0}^{i-1} f_j^+}{\prod_{j=1}^{i} f_j^-}, \qquad i \in \{0, 1, ..., mB\}, \tag{16}
$$

where

$$
\gamma(0) = \left(\sum_{i=0}^{mB} \frac{\prod_{j=0}^{i-1} f_j^+}{\prod_{j=1}^{i} f_j^-} \right)^{-1}.
$$

Proposition 1: Let $\tilde{\pi}(a)$ be the stationary distribution of $\tilde{P}$ for a given value of a, and let π_i denote the stationary distribution of A_i. Then, the limiting stationary distribution of $\tilde{P}$, $\tilde{\pi}(0)$ is given by

$$
\tilde{\pi}(0) = [\gamma_1 \pi_1 \ \gamma_2 \pi_2 \ ... \ \gamma_{mB} \pi_{mB}], \tag{17}
$$

where γ_i is given in (16).
We note that the states $q > m \cdot max_{th}$ are transient and the steady state probability of $q > m \cdot max_{th}$ is 0. This can be shown as follows. From equation (12) we have

$$
f_i^+ = \sum_{j=\lceil \frac{i}{m} \rceil}^{B} \pi_i(j)(mj-i).
$$

Since the packet drop probability is 1 for $q \geq m \cdot max_{th}$, the steady state probability, $\pi_i(j)$, becomes

$$
\pi_i(j) = \begin{cases} 1 & j = 0 \\ 0 & j > 0 \end{cases} \ \forall i \geq m \cdot max_{th}. \tag{18}
$$

This is true since all the packets arriving at the queue would be dropped and the queue would be empty. Since, $\pi_i(j)$ is as given in equation (18), and $i \geq m \cdot max_{th} > 0$, we get that $f_i^+ = 0, \quad \forall i \geq m \cdot max_{th}$. However, from equation (14) we get $f_i^- > 0 \quad \forall i$. From equation (16) it follows that $\gamma(i) = 0 \quad \forall i > m \cdot max_{th}$.

If the arrival process to the queue is Poisson with rate λ and the service times are exponentially distributed with mean μ, then the stationary distribution of A_i, π_i, is given by

$$
\pi_i := \{\pi_i(j)\} = \pi_i(0) \rho_i^j, \tag{19}
$$

where

$$
\pi_i(0) = \left(\sum_{j=0}^{B} \rho_i^j \right)^{-1},
$$

and $\rho_i = \lambda \cdot (1 - h(\hat{q}_n = i)) \cdot \mu^{-1}$. Here $h(\cdot)$ is the given drop function.

From the above analysis, using equation (16), we are able to obtain the stationary distribution of the average queue length variable, q. We use the PASTA property which states that the arrivals see this average queue length distribution. Hence, we can calculate the average packet drop probability as

$$P_{avg} = \sum_{i=0}^{mB} \gamma(i)h(i), \tag{20}$$

For batch arrivals, we make the assumption that every packet in a batch is dropped with same drop probability. With this assumption, we can write the probability of dropping at least one packet in batch as

$$P_{avg} = \sum_{i=0}^{m(B-N)} (1 - (1 - \gamma(i))^N) \cdot h(i) + \sum_{i=m(B-N)+1}^{mB} h(i), \tag{21}$$

B. Approximate analysis

We present an approximate analysis which allows us to obtain the mean value of q without having to find the distribution.

Proposition 2: The equilibrium value of q, q^*, can be obtained by solving

$$q^* = mE(Q(q^*)), \tag{22}$$

where $E(Q(q^*))$ is the expected queue length when we drop with constant probability corresponding to q^*.

Using equations (12) and (14) we get

$$f_i^+ - f_i^- = m \sum_{j=0}^{B} j\pi_i(j) - i = mE(Q_i) - i, \tag{23}$$

where $E(Q_i)$ is the expected queue length in the unperturbed Markov chain with $\hat{q}_n = i$. The term $f_i^+ - f_i^-$ can be thought of as an indicator for transitions of the average queue length, q. We note that in states where $mE(Q_i) > i$, $f_i^+ - f_i^-$ is positive, hence indicating a tendency to increase q. Similarly, when $mE(Q_i) > i$, $f_i^+ - f_i^-$ is negative, hence indicating a tendency to decrease q. The equilibrium value of q will then occur at the point where $f_i^+ - f_i^- = 0$. Hence, q^* can be obtained by solving (22). For the existence and uniqueness of the solution, we assume that the drop probability function (this is the case in the gentle variant of RED [13]) is a continuous non-decreasing function, $h(\hat{q})$, on the estimated average queue length. The effective offered load, ρ^*, for a given equilibrium value of q^* is then given by

$$\rho^* = \rho(1 - h(q^*)). \tag{24}$$

Since ρ^* is a continuous non-increasing function on q^*, we can express q^* as

$$q^* = h^{-1}\big(1 - \frac{\rho^*}{\rho}\big) := g(\rho^*). \tag{25}$$

We note that g is continuous non-increasing function on ρ^*. Since the average queue length, $E(Q(\rho^*))$, in the stationary regime is a continuous non-decreasing function of the offered load, then we can say that $E(Q(q^*))$, is a continuous non-increasing function of q^*. As q^* and $EQ(\cdot)$ the map to the same interval i.e., $[0, B]$, we can say that the solution to equation (22) exists and is unique.

C. Exact analysis

In subsection II-A we presented an anaylsis for the limiting case $a \rightarrow 0$. In this subsection we online an algorithm with which we can compute the steady state probabilities when $a \nrightarrow 0$. However, a still has to satisfy the condition $amB < 1$.

The transition probability matrix of the two dimensional Markov (q, Q) given in equation (7) can be written as

$$\tilde{P} = \begin{bmatrix} L_0 & F_0 & 0 & \dots & \dots \\ B_1 & L_1 & F_1 & 0 & \dots \\ 0 & B_2 & L_2 & F_2 & \dots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & \dots & B_{mB} & L_{mB} \end{bmatrix}, \tag{26}$$

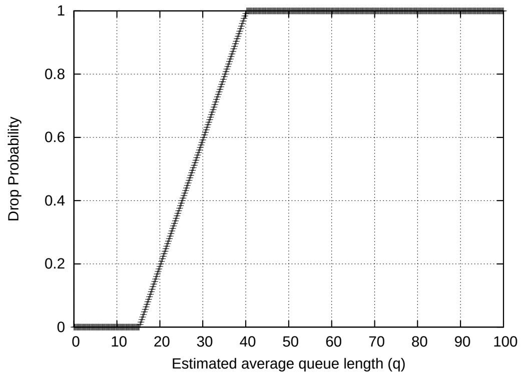

Fig. 2. RED drop probability function

where B_i, L_i, F_i denote the backward, local, and forward transition matrices, respectively. This is the transition matrix for a level dependent QBD process. Let the $\tilde{\pi}$ be stationary distribution of $\tilde{P}$. Let $\tilde{\pi} = [\tilde{\pi}_0 \tilde{\pi}_1 ... \tilde{\pi}_{mB}]$, where $\tilde{\pi}_i$'s are vectors of size $1 \times B$. $\tilde{\pi}_i$ is the steady state probability vector of Q for $q = i$. $\tilde{\pi}$ can be found using the following algorithm [12].

1) Compute the S_i matrices using the following recursion

$$\begin{aligned} S_0 &= L_0 \\ S_i &= L_i + B_i (I - S_{i-1})^{-1} F_{i-1}, \qquad 1 \leq i \leq mB. \end{aligned}$$

2) $\tilde{\pi}_{mB}$ is computed by solving

$$\begin{aligned} \tilde{\pi}_{mB} &= \tilde{\pi}_{mB} S_{mB}, \\ \tilde{\pi}_{mB} \cdot \underline{1} &= 1. \end{aligned}$$

3) $\tilde{\pi}_i$, $0 \leq i \leq mB$ is computed using the recursion

$$\tilde{\pi}_i = \tilde{\pi}_{i+1} B_{i+1} (I - S_n)^{-1},$$

4) $\tilde{\pi}$ is found by normalization

$$\tilde{\pi} = \frac{\tilde{\pi}}{\tilde{\pi} \cdot \underline{1}}$$

III. Numerical results

In this section we present the results obtained through analysis as described in the previous section and the results obtained through simulations. The arrival process is assumed to be Poisson batch arrival process with rate λ and fixed batch sizes, N. The service times are assumed to exponential with mean $1/\mu$. The offered load, ρ, is defined as $\rho = \frac{\lambda N}{\mu}$. The analytical results were obtained by numerically calculating the values using MATLAB. The buffer length, B, is taken to be 100 whereas the discretization parameter, m, is taken to be 5. The drop probability function is taken to be of the form

$$P_{drop} = \begin{cases} 0, & \hat{q} < min_{th} \\ \frac{\hat{q} - min_{th}}{max_{th} - min_{th}}, & min_{th} \leq \hat{q} \leq max_{th} \\ 1, & \hat{q} > max_{th} \end{cases} \tag{27}$$

We note that all the values of min_{th}, max_{th}, and $\hat{q}$ are scaled by a factor m in the analysis in order to discretize the estimated average queue size. The unscaled values for min_{th} and max_{th} are 15 and 40, respectively, and the drop probability function versus the unscaled estimated average queue length is as shown in figure 2. We shall use the average drop probability and the probability of at least drop in a burst as performance measures. The simulations were performed using C, and approximately 10^8 packets were generated during the simulations.

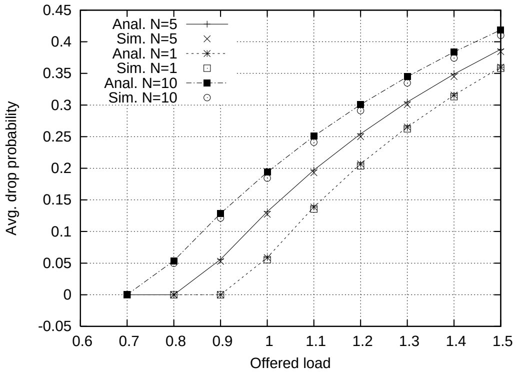

Fig. 3. Average packet drop probability versus offered load

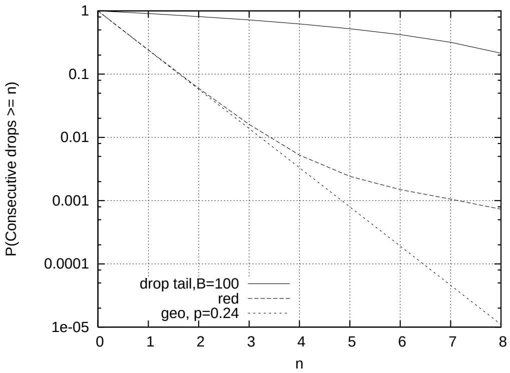

Fig. 4. Probability of greater than or equal to n consecutive drops. $\rho = 1.1$. Burst size = 10.

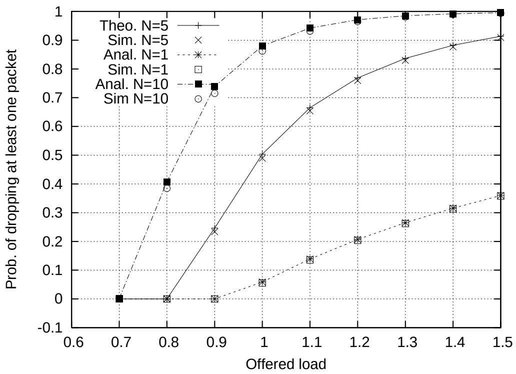

Fig. 5. Probability of drop of at least one packet in a burst versus offered load

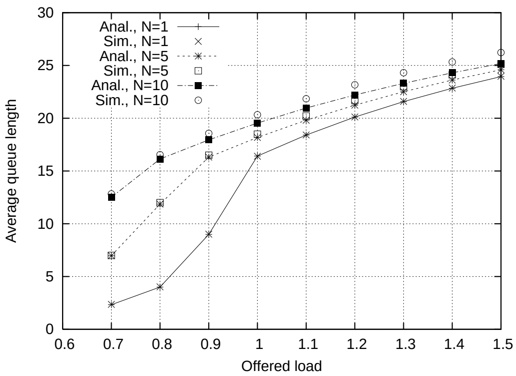

Fig. 6. Average queue length versus offered load

A. Limiting case

First, we present the results for the limiting case. In the simulations, the averaging parameter, a, was taken to be 0.0001. Figures 3 shows the analytical as well as the simulation results for average packet drop probability as a function of the offered load for different values of burst sizes. There is a fairly good correspondence between the analytical and simulation results.

In [4] it has been argued that the number of drops becomes infinite with positive probability as $a \to 0$. However, for the drop function as shown in Figure 1, the number of consecutive drops appears to follow a geometric behaviour for small values of the number of consecutive drops. We plot the probability of greater than or equal to n consecutive drops as a function of n for a drop tail and a RED queue in Figure 4. The decay rate of the tail is much faster in the case of a RED queue when compared to that of a drop tail queue. For small values of n, the decay rate is close to p^n, where p is the average packet drop probability. Most of the mass can be seen to concentrated within small values of n and, so the geometric approximation appears to be a reasonable one. A continuous drop function ensures that the drops are independent and, thus, reduces the synchronization of drops for different flows. This argument supports using gentle RED. The tail drop probability during this simulation scenario was 2×10^{-4} which also suggests that synchronized losses are reduced.

As mentioned above, the probability of dropping consecutive packets becomes independent when the drop function is continuous. This suggests that the probability of drops in a burst is binomially distributed. The probability of at least one drop in a burst can thus be approximated by $1-(1-p)^N$, where p is the average drop probability and N is the burst size. This behaviour is seen in Figure 5 in which the probability of dropping at least one packet in a burst is plotted as a function of the offered load.

Next, we validate the approximate analysis for obtaining the average queue length. In order to obtain the average queue length we need to solve the fixed point equation as given in equation (22). For example, the average queue length as observed by a incoming packet in a $M/M/1$ queue with finite buffer B is given

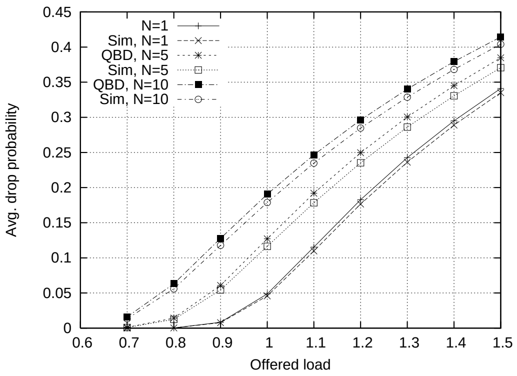

Fig. 7. Average drop probability versus offered load.

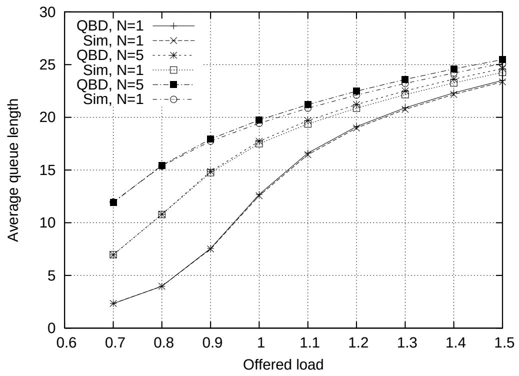

Fig. 8. Average queue length versus offered load.

by

$$E(Q_{q^*}) = \frac{\rho_{q^*}}{(1-\rho_{q^*})} - (B+1)\frac{\rho_{q^*}^{(B+1)}}{(1-\rho_{q^*}^{(B+1)})}, \tag{28}$$

where

$$\rho_{q^*} = \rho(1 - \frac{q^* - min_{th}}{max_{th} - min_{th}}).$$

We can now obtain the equilibrium value of average queue length by solving

$$q^* = \frac{\rho_{q^*}}{(1-\rho_{q^*})} - (B+1)\frac{\rho_{q^*}^{(B+1)}}{(1-\rho_{q^*}^{(B+1)})}. \tag{29}$$

Similarly, for the batch arrivals we use the expression for the infinite buffer case as an approximation. The equilibrium value for batch arrivals is obtained by solving

$$q^* = \frac{\rho_{q^*}}{(1-\rho_{q^*})}\frac{N+1}{2}. \tag{30}$$

For the values of min_{th}, max_{th}, and B assumed in this section, we plot in Fig. 6 the average queue as obtained from solving equations (29) and (30) numerically using MATLAB and the results obtained through simulations.

B. General Case

Next, we consider the case when $a \nrightarrow 0$. In the rest of this subsection we assume that the discretization parameter, m, is 1. The drop function is the same as in the previous section. We need that amB be less than 1. Thus, a has to be chosen such that $a < 0.01$ in order to use the QBD algorithm. We choose $a = 0.009$ which is close to 0.01. In Figure 7 and Figure 8 we plot the average drop probability and the average queue length, respectively, as a function of the offered load. We note that the behaviour is similar to that observed in the limiting case.

In order to see the effect of choosing an averaging parameter which is not necessarily small, we plot in Figure 9 the distribution function of the average queue length, q, using the QBD analysis, simulations, and the singular perturbation (SP) analysis. The offered load was 1.2 and the burst size was taken to be 10. As can be observed from the figure, the effect of increasing a is to increase the variance of q. We cannot, however, reduce the variance of q by taking the limit. There is a limiting distribution of q (as was seen from the limiting case) and hence a limiting variance which is different from 0.

IV. CONCLUSIONS

We used a singular perturbation based approach to find the limiting stationary distribution of the average queue length and the packet drop probability in a RED enabled queue. For Poisson arrival and exponential service times the limiting stationary distribution was given as closed form expressions. The equilibrium point of the average queue length was found to be the fixed point solution of a function depending on the average

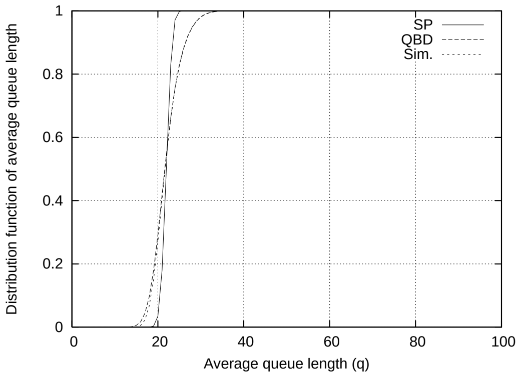

Fig. 9. Distribution function of the average queue length. $N = 10$ and offered load is 1.2. $a = 0.009$.

queue length of a system with no averaging. The analytical results were observed to match fairly accurately with results obtained through simulations. We also showed that by using a continuous drop function we can reduce the number of consecutive losses thereby avoiding synchronized losses. We showed that for the case when the averaging parameter is not small, the stationary distribution can be computed using an algorithm for level dependent QBDs. However, for small values of the averaging parameter, the computational complexity of using this algorithm was much greater than that by using the singular perturbation approach. In this work, we considered scenario in which packet generation process is independent of packet drops. Future work will involve the study of a RED queue with closed loop TCP traffic sources.

References

[1] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, August 1993.

[2] R. Pan, B. Prabhakar, and K. Psounis. Choke: A stateless active queue management scheme for approximating fair bandwidth allocation. In *Proc. of IEEE INFOCOM 2000*, 2000.

[3] L. Zhang and D. Clark. Oscillating behavior of network traffic: a case study simulation. *Internetworking:research and experience*, 1(2):101–112, 1990.

[4] T. Bonald, M. May, and J. Bolot. Analytic evaluation of RED performance. In *Proc. of IEEE INFOCOM 2000*, 2000.

[5] V. Mishra, W. Gong, and D. Towsley. Fluid based analysis of a network of AQM routers supporting TCP flows with an application to RED. In *Proc. of ACM SIGCOMM 2000*, 2000.

[6] C. Hollot, V. Misra, D. Towsley, and W. Gong. A control theoretic analysis of RED. In *Proc. of IEEE INFOCOM 2001*, 2001.

[7] E. Kuumola, J.A.C. Resing, and J. Virtamo. Joint distribution of instantaneous and averaged queue length in an M/M/1/K system. In P. Tran-Gia and J. Roberts, editors, *Proceedings of the 15th ITC Specialist Seminar "Internet Traffic Engineering and Traffic Management"*, pages 58–67, July 2002.

[8] P.J. Schweitzer. Perturbation series expansion of nearly completely-decomposable markov chains. In O.J. Boxma, J.W. Cohen, and H.C. Tijms, editors, *Teletraffic Analysis and Computer Performance Evaluation*. Elsevier Science Publishers B.V. (North Holland), 1986.

[9] K.E. Avrachenkov, J.A. Filar, and M. Haviv. Singular perturbations of Markov chains and decision processes. A survey. In E.A. Feinberg and A. Shwartz, editors, *Handbook of Markov Decision Processes: Methods and Applications*, International Series in Operations Research and Management Science, pages 113–153. Kluwer Academic Publishers, January 2002.

[10] S. Floyd. Red: Discussions of setting parameters, November 1997. Available at http://www.icir.org/floyd/red.html.

[11] P.J. Schweitzer. Perturbation theory and finite Markov chains. *J. Appl. Probability*, 5:410–413, 1968.

[12] D.P. Gaver, P.A. Jacobs, and G. Latouche. Finite birth-and-death models in randomly changing environments. *Adv. Appl. Probability*, 16:715–731, 1984.

[13] S. Floyd. Recommendation on using the gentle variant of RED, March 2000. Available at http://www.icir.org/floyd/red.html.