

On Packet Marking Function of Active Queue Management Mechanism: Should It Be Linear, Concave, or Convex?

Hiroyuki Ohsaki and Masayuki Murata

Graduate School of Information Science and Technology

Osaka University

Toyonaka, Osaka 560-8531, Japan

E-mail: {oosaki,murata}@ist.osaka-u.ac.jp

ABSTRACT

Recently, several gateway-based congestion control mechanisms have been proposed to support the end-to-end congestion control mechanism of TCP (Transmission Control Protocol). In this paper, we focus on RED (Random Early Detection), which is a promising gateway-based congestion control mechanism. RED randomly drops an arriving packet with a probability proportional to its average queue length (i.e., the number of packets in the buffer). However, it is still unclear whether the packet marking function of RED is optimal or not. In this paper, we investigate what type of *packet marking function*, which determines the packet marking probability from the average queue length, is suitable from the viewpoint of both steady state and transient state performances. Presenting several numerical examples, we investigate the advantages and disadvantages of three packet marking functions: linear, concave, and convex. We show that, although the average queue length in the steady state becomes larger, use of a concave function improves the transient behavior of RED and also realizes robustness against network status changes such as variation in the number of active TCP connections.

Keywords: AQM (Active Queue Management) Mechanism, RED (Random Early Detection), Packet Marking Function, Steady State Performance, Transient State Performance, Robustness

1. INTRODUCTION

In recent years, AQM (Active Queue Management) mechanisms have been extensively studied by many researchers.^{1, 2} The AQM mechanism supports the congestion control mechanism of TCP (Transmission Control Protocol), which operates on an end-to-end basis. RED (Random Early Detection) is a representative AQM mechanism,³ which randomly discards an arriving packet to control the router's queue length (i.e., the number of packets in the router's buffer). Compared with the conventional Drop Tail router, the RED router is claimed to be more effective since RED keeps the average queue length small and achieves high throughput.^{3, 4} Moreover, the operation algorithm of RED is quite simple since it does not distinguish each TCP connection.

However, RED has the well-known problem that its effectiveness is greatly dependent on the setting of four control parameters, min_{th} , max_{th} , max_p , and q_w , so that the effectiveness of RED has been denied by many researchers.⁵⁻⁸ Another problem associated with RED is that its average queue length is dependent on the number of active TCP connections; i.e., the optimal setting of RED control parameters is affected by the number of active TCP connections. Many research efforts have so far been devoted to solve such problems (e.g.,⁹⁻¹⁵).

As explained above, various approaches to solving the problems associated with RED have been proposed in the literature. However, these problems originate from the fact that the RED algorithm was designed with an ad hoc approach. For example, RED randomly discards an arriving packet with a probability that is proportional to its average queue length. Upon considering the purposes of AQM mechanisms, it seems reasonable to raise the packet marking probability when the average queue length is large, and to lower it when the average queue length is small. However, little investigation has been done to discuss "whether the packet marking probability should be proportional to the average queue length or not". RED randomly discards an arriving packet with a probability which is determined from a linear function to the average queue length. More specifically, when the difference between the average queue length and min_{th} is increased by Δ , RED increases its packet marking probability linearly in relation to Δ . However, it is analytically known that TCP throughput is inversely proportional to

\sqrt{p} , which is the packet loss probability in the network.¹⁶ Also, it is known that, if the bottleneck router in the network can be modeled by a single $M/M/1$ queueing system, the average queue length is given by $\rho/(1-\rho)$, where ρ is the utilization factor.¹⁷ Hence, on the basis of the fact that the primary purpose of AQM mechanisms is to control the average queue length, it is considered that the packet marking probability of RED should not be changed linearly in relation to the average queue length. Of course, rather than using a simple queueing model such as $M/M/1$, it is necessary to model the interaction between the TCP congestion control mechanism and the network more precisely. For this reason, if the packet marking probability p_b is determined by taking account of the characteristic of the TCP window-based flow control, the steady state performances and/or transient state performances of RED can be improved. In this paper, we therefore address this problem — how the packet marking probability p_b of RED should be determined from its average queue length for achieving good steady state and transient state performances.

In this study, using an analytic approach, we therefore investigate how the function determining the packet marking probability of RED affects its performance, particularly in terms of steady state performance and transient state performance. Specifically, we utilize results of the TCP steady state analysis¹⁶ and the RED steady state analysis,¹⁸ and show how the packet marking function of RED should be determined. In numerical examples, we consider three classes of packet marking functions — linear, concave, and convex — and show how the RED performance is affected by the choice of packet marking function. We show that when the packet marking function is concave, RED achieves good transient state performance and robustness compared with the case of linear or convex function.

The organization of this paper is as follows. In Section 2, we investigate what type of packet marking function is appropriate by utilizing the analytic results presented in.^{16, 18} In Section 3, by examining several numerical examples, we show that when the packet marking function is concave, the transient behavior and robustness of RED are improved compared with the case of linear or convex function. Finally, in Section 4, we present our conclusions and discuss future works.

2. ANALYSIS

In what follows, we consider the case when the AQM mechanism of RED is operating as expected, i.e., $min_{th} \leq \bar{q} < max_{th}$. In this case, the packet marking probability p_b is determined based on the average queue length \bar{q} as

$$p_b = max_p \left(\frac{\bar{q} - min_{th}}{max_{th} - min_{th}} \right). \quad (1)$$

When the average queue length \bar{q} is large, p_b takes a value close to max_p . On the contrary, when the average queue length \bar{q} is small, p_b takes a value close to zero. As can be seen from Eq. (1), the packet marking function of RED is a linear function in relation to $(\bar{q} - min_{th})$. However, use of a linear function has not been fully validated by taking account of both steady-state and transient-state performances of RED. For example, the window-based flow control of TCP changes its window size non-linearly according to the packet loss probability in the network. For this reason, if the packet marking probability p_b is determined by taking account of the characteristic of the TCP window-based flow control, the steady state performances and/or transient state performances of RED can be improved. In what follows, we therefore discuss how the packet marking probability p_b should be determined by utilizing the analytic results of TCP¹⁶ and RED.¹⁸

In our analysis, by combining the stochastic model of the TCP window size and the deterministic model of the RED queue length, we analyze how the average queue length is affected by the RED's packet marking mechanism. Specifically, we analyze toward what value the average queue length converges with the packet marking mechanism of RED when the average queue length is given at some time. Consequently, we clarify the effect of the packet marking function on the average queue length of RED in steady state (i.e., steady state performance) the average queue length of RED in transient state (i.e., transient state performance), and robustness against variation in the number of active TCP connections.

First, the packet marking function of RED as defined by Eq. (1) is replaced by

$$p_b = max_p f \left(\frac{\bar{q} - min_{th}}{max_{th} - min_{th}} \right), \quad (2)$$




Figure 1. Evolution of TCP window size in congestion avoidance phase.

where f is a monotonically increasing function satisfying $f(0) = 0$ and $f(1) = 1$. We then introduce the notion of *queue occupancy* to indicate how many packets exist in the router's buffer. The queue occupancy is defined by

$$x \equiv \frac{\bar{q} - \min_{th}}{\max_{th} - \min_{th}}$$

for $\min_{th} \leq \bar{q} < \max_{th}$. The purpose of the remainder of this section is to investigate how the packet marking function f affects the steady state and transient state performances of RED.

We consider the expected value of the TCP window size when the packet loss probability in the network is given. In,¹⁶ the TCP throughput in steady state is derived. In the derivation process, $W(p)$, which is the expected value of TCP window size just before TCP detects a packet loss, is also derived (see Eq. (3) in¹⁶).

$$W(p) = \frac{2+b}{3b} + \sqrt{\frac{8(1-p)}{3bp} + \left(\frac{2+b}{3b}\right)^2} \quad (3)$$

In the above equation, b is the number of packets required for the destination host to return an ACK packet (usually, $b = 1$ or $b = 2$), and p is the packet loss probability in the network. In the analysis, the authors assume (1) TCP is operating in the congestion evasion phase, (2) all packet losses are detectable by duplicate ACKs (i.e., no timeout is triggered), (3) the packet loss probability in the networks is constant, and (4) the maximum window size is sufficiently large. Note that in,¹⁶ the second assumption (i.e., no timeouts) is relaxed and the more detailed result of the TCP window size is derived. However, we use a simple result given by Eq. (3) since it is more tractable than the detailed result, allowing us to know more insight on the packet marking function of RED. Also, note that Eq. (3) gives the expected value of TCP window size just before a packet loss is detected. Immediately after detecting the packet loss, the TCP window size is decreased to one half. Then, the TCP window size is slowly increased until the next packet loss is detected. In this paper, we therefore use $w(p)$ as the expected value of TCP window size (Fig. 1):

$$w(p) = \frac{1}{2} \left(\frac{W(p)}{2} + W(p) \right) = \frac{3W(p)}{4}. \quad (4)$$

Furthermore, we consider the queue length of RED in steady state when the TCP window size w is given. In,¹⁸ the queue length \bar{q} of RED is derived when the number N of TCP connections have the identical window size w .

$$\bar{q} = Nw - B\tau, \quad (5)$$

where B is the maximum transmission capacity of the RED router (i.e., the smaller value between the processing speed of the RED router and the bandwidth of the outgoing link). τ is the two-way propagation delay of the

TCP connection excluding the queuing delay at the buffer. In the analysis, almost the same assumptions as those of¹⁶ are made.

RED randomly discards an arriving packet with probability p_b . Hence, for a given packet marking probability p_b , the packet loss probability to TCP connections is given by³

$$p = \left(\frac{1}{2p_b} + \frac{1}{2} \right)^{-1}. \quad (6)$$

From Eqs. (2), (5), and (6), the average queue length \bar{q} can therefore be given by

$$\bar{q} = N w(p) - B \tau \quad (7)$$

$$= \frac{N}{4b} \left\{ 2 + b + b \sqrt{\frac{4 - 8b + b^2 + \frac{12b}{max_p f(x)}}{b^2}} \right\} - B \tau. \quad (8)$$

This equation indicates that the queue length of RED converges to \bar{q} , when the packet marking probability p_b is determined according to the function f . Note that \bar{q} is the expected value of the queue length since Eq. (4) is also the expected value. Let x^* be the queue occupancy converged by the packet dropping mechanism of RED (i.e., the average queue length in steady state if p_b is fixed). x^* is given by

$$x^* = \frac{\bar{q}(x) - min_{th}}{max_{th} - min_{th}}. \quad (9)$$

Namely, this means that the probabilistic packet marking mechanism of RED with the packet marking function f and the queue occupancy x governs the queue occupancy toward x^* . For example, by plotting Eq. (9) on the $x-x^*$ plane, the effect of the packet marking function f on the steady state and transient state performances of RED can be rigorously analyzed.

To analyze the relation between Eq. (9) and RED performance, we graphically plot the relation between queue occupancies x and x^* . An example of Eq. (9) is shown in Fig. 2. The straight line $x^* = x$ is also plotted in the figure. The following points can be inferred from this figure.

1. The queue occupancy in steady state ($t \rightarrow \infty$) is given by the intersection of the curve of Eq. (9) and the straight line $x^* = x$.
2. The steeper the gradient (dx^*/dx) of Eq. (9), the larger the impact of variation in the queue occupancy x on the average queue length of RED.
3. Conversely, when the gradient (dx^*/dx) of Eq. (9) is gentle, variation in the queue occupancy x has a negligible effect on the average queue length of RED.
4. In order for the average queue length of RED to be stable when $min_{th} < \bar{q} < max_{th}$, the gradient (dx^*/dx) of Eq. (9) must be negative.

On the basis of these observations, we find the following points with regard to the steady state and transient state performances of RED.

1. When Eq. (9) is convex (i.e., $d^2x^*/dx^2 < 0$), the average queue length of RED in steady state becomes small. Moreover, when the queue occupancy is small, transient state performance is good. Conversely, when the queue occupancy is large, transient state performance becomes degraded.
2. When Eq. (9) is linear (i.e., $d^2x^*/dx^2 = 0$), the average queue length of RED in steady state is larger than that with a convex function. Also, transient state behavior is not affected by the queue occupancy.
3. When Eq. (9) is concave (i.e., $d^2x^*/dx^2 > 0$), the average queue length of RED in steady state is large, and the transient state behavior is not good when the queue occupancy is small. On the contrary, transient state performance is good when the queue occupancy is large.

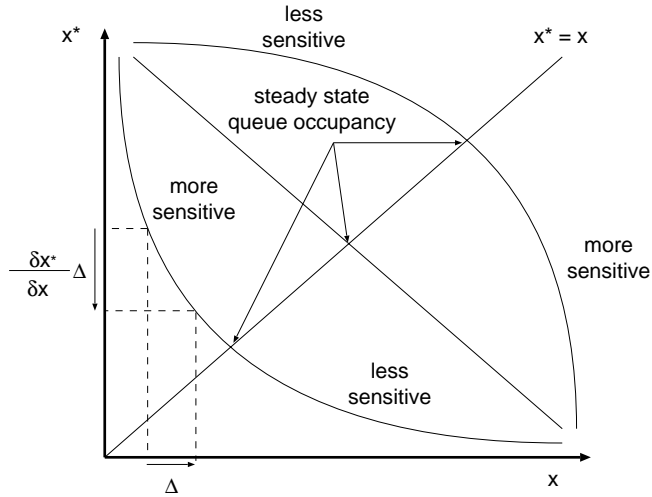


Figure 2. Queue occupancy in the x - x^* plane.

On the basis of the above observations, it is considered that the packet marking function f should be chosen such that Eq. (9) becomes a linear function in relation to x . Namely, it is desirable for the gradient $d\bar{q}/dx$ of Eq. (7) given by the following equation to be constant.

$$\frac{d\bar{q}}{dx} = \frac{-3Nf'(x)}{2max_p \sqrt{4 - 8b + b^2 + \frac{12b}{max_p f(x)} f(x)^2}} \quad (10)$$

This equation suggests that the gradient of Eq. (7) is determined by the parameter at the TCP destination host b , the number of active TCP connections N , and the maximum packet marking probability max_p . Conversely, it suggests that the gradient of Eq. (7) is almost independent of the propagation delay of TCP connections τ and the capacity of the RED router B . This result agrees with conventional research results that the average queue length \bar{q} of RED is dependent on the number of TCP connections.^{18, 13, 19} Therefore, when considering how the packet marking function f should be determined, the number of active TCP connections N should be primarily taken account of.

Moreover, the packet marking function f that makes Eq. (9) a linear function in relation to x can be derived by equating $d\bar{q}/dx$ from Eq. (10) with a constant value α and by solving this ordinary differential equation. The solution is given by

$$f(x) = 12 \left\{ max_p \left(8 - \frac{4}{b} + b \left(\frac{16\alpha^2 x^2}{N^2} - 1 \right) \right) - \frac{48\alpha b \sqrt{max_p} x C(1)}{N} + 36bC(1)^2 \right\}^{-1},$$

where $C(1)$ is a constant. As can be seen from the above equation, for determining f so that Eq. (7) becomes linear, f must be changed according to the number of active TCP connections N . Our analytic results clearly indicates that information on the number of active TCP connections is necessary for determining the packet marking probability p_b to optimize the steady state and transient state performances of RED. However, the RED router has no capability of knowing the number of active TCP connections since it does not distinguish each TCP connection. Of course, it is not impossible to estimate the number of active TCP connections, for example, by distinguishing each TCP connection as is done in SRED. However, adding such processing to the AQM router makes its implementation very complex. Since implementation simplicity is one of important features of RED, in reality, it is desirable to choose a packet marking function f such that f is as independent of the number of TCP connections as possible, and that Eq. (9) has as much linearity in relation to x as possible.

Hence, in the next section, we consider three types of packet marking functions — linear, concave, and convex — and examine which packet marking function is most suitable for optimizing RED performance. Note that the

actual steady state and transient state performances of RED such as throughput is determined by not only the type of packet marking function but also the setting of its control parameters such as min_{th} , max_{th} , max_p , and q_w . However, in what follows, we limit our attention to the packet marking function, and carefully investigate how the packet marking function affects the steady state and transient state performances of RED.

3. NUMERICAL EXAMPLES

In the following numerical examples, three types of function classes, \mathcal{F}_ϕ , \mathcal{G}_ϕ , and \mathcal{H}_ϕ , are considered as the packet marking function f .

- Linear

$$\mathcal{F}_\phi(x) = x^\phi \quad (11)$$

Note that this function is linear when $\phi = 1$, is concave when $\phi > 1$, and is convex when $\phi < 1$.

- Concave

$$\mathcal{G}_\phi(x) = \left(1 - \sqrt{1 - x^2}\right)^\phi \quad (12)$$

$\phi(> 0)$ is a parameter determining concavity. Namely, when ϕ is large, $\mathcal{G}_\phi(x)$ takes a small value. In order for \mathcal{G}_ϕ to be concave,

$$\frac{d^2\mathcal{G}_\phi(x)}{dx^2} = \frac{1}{(1 - x^2)^{\frac{3}{2}}} \left\{ \phi \left(1 - \sqrt{1 - x^2}\right)^{\phi-2} \times \left(1 + \sqrt{1 - x^2} ((\phi - 1) x^2 - 1)\right) \right\} \geq 0$$

must be satisfied. By solving the above inequalities for ϕ , we have

$$\phi \geq \lim_{x \rightarrow 0} \frac{-1 + \sqrt{1 - x^2} + x^2 \sqrt{1 - x^2}}{x^2 \sqrt{1 - x^2}} = \frac{1}{2}.$$

- Convex

$$\mathcal{H}_\phi(x) = \left(\sqrt{1 - (1 - x)^2}\right)^\phi \quad (13)$$

$\phi(> 0)$ is a parameter determining convexness. Namely, when ϕ is large, $\mathcal{H}_\phi(x)$ takes a small value. In order for \mathcal{H}_ϕ to be convex,

$$\frac{d^2\mathcal{H}_\phi(x)}{dx^2} = \phi \left(-((x - 2) x)\right)^{\frac{\phi-4}{2}} \times \left(\phi (x - 1)^2 - (x - 2) x - 2\right) \leq 0$$

must be satisfied. Similar to the previous case, by solving the above inequalities for ϕ , we have

$$\phi \leq \lim_{x \rightarrow 0} \frac{2 - 2x + x^2}{1 - 2x + x^2} = 2.$$

Note that $\mathcal{F}_\phi(x)$ with $\phi = 1.0$ is identical to Eq. (1) of the original RED.

We investigate which packet marking function is the most desirable among \mathcal{F}_ϕ (linear), \mathcal{G}_ϕ (concave), and \mathcal{H}_ϕ (convex) by showing several numerical examples. First, the queue occupancy in the x - x^* plane is shown when control parameters and system parameters are configured according to Tab. 1. Figures 3, 4, and 5 respectively show results when \mathcal{F}_ϕ (linear), \mathcal{G}_ϕ (concave), and \mathcal{H}_ϕ (convex) are used as the packet marking function. In these figures, ϕ is changed to 0.5, 1.0, and 1.5. The straight line $x = x^*$ is also plotted in all figures.

Figure 3 shows that the average queue length of RED (i.e., the intersection with the straight line $x = x^*$) increases as the value of ϕ becomes large when \mathcal{F}_ϕ (linear) is used as the packet marking function f . Moreover,

Table 1. Parameters used in numerical examples.

min_{th}	50	[packet]
max_{th}	100	[packet]
max_p	0.1	
B	1.25	[packet/ms]
τ	10	[ms]
N	10	
b	1	




Figure 3. RED queue occupancy in the $x-x^*$ plane with \mathcal{F}_ϕ (linear) for $\phi = 0.5, 1.0,$ and 1.5 .

it is also found that when the queue occupancy x is small, the gradient (dx^*/dx) of Eq. (9) is steep. On the other hand, when the queue occupancy is large, the value of Eq. (9) is almost zero. Namely, in the state where the queue occupancy is small, since RED discards arriving packets, the queue length is changed rapidly. In the state where the queue occupancy is large, the queue length rapidly decreases to min_{th} (i.e., the queue occupancy is empty) irrespective of the packet marking probability of RED.

When \mathcal{G}_ϕ (concave) is used as the packet marking function, the gradient (dx^*/dx) of Eq. (9) is small (see Fig. 4). This figure indicates that as the value of ϕ becomes larger (i.e., with increasingly stronger concavity), the average queue length of RED becomes larger. Moreover, this figure shows that Eq. (9) is closer to the straight line as compared with Fig. 3. Namely, the effect that packet losses have on the queue length of RED is almost independent of the queue occupancy.

The result when using \mathcal{H}_ϕ (convex) as the packet marking function f is shown in Fig. 5. This figure shows the relation of the queue occupancy in the $x-x^*$ plane. It is evident that the gradient (dx^*/dx) of Eq. (9) is quite steep when the queue occupancy is small. Moreover, it is also evident that, as the value of ϕ becomes large (i.e., with increasingly strong convexity), the gradient of Eq. (9) becomes even steeper. Namely, when \mathcal{H}_ϕ (convex) is used as the function f , the queue length of RED is changed rapidly when the queue occupancy is small.

Next, when \mathcal{F}_ϕ (linear), \mathcal{G}_ϕ (concave), and \mathcal{H}_ϕ (convex) are used as the packet marking function f , we show the effect of variation in the number of active TCP connections on the steady state performance and transient state performance of RED. The relation between the number of TCP connections N and the queue occupancy when using \mathcal{F}_ϕ (linear) is shown in Fig. 6. In this figure, the number of TCP connections N is varied from 1 to 20, and the parameter values shown in Tab. 1 are used. As discussed in Section 2, the intersection of the curved surface and the $x-x^*$ plane means the average queue length of RED in steady state. This indicates that the average queue length of RED in steady state becomes larger as the number of TCP connections N becomes large. In particular, it is evident that when the number of TCP connections N is small, the variation in the




Figure 4. RED queue occupancy in the $x-x^*$ plane with \mathcal{G}_ϕ (concave) for $\phi = 0.5, 1.0,$ and 1.5 .




Figure 5. RED queue occupancy in the $x-x^*$ plane with \mathcal{H}_ϕ (convex) for $\phi = 0.5, 1,$ and 1.5 .

number of TCP connections significantly affects the average queue length. Furthermore, it is also evident that the gradient (dx^*/dx) of Eq. (9) becomes small as the number of TCP connections N becomes large. This means that the transient state performance of RED is affected by changing the number of TCP connections.

Figure 7 shows the result when using \mathcal{G}_ϕ (concave) as the packet marking function f . Except for using \mathcal{G}_ϕ (concave) as the packet marking function f , the same parameter values as those for Fig. 6 are used. This figure shows that the gradient (dx^*/dx) of Eq. (9) is negligibly dependent on the number of active TCP connections when \mathcal{G}_ϕ (concave) is used. Furthermore, similar to Fig. 6, the average queue length of RED in steady state becomes larger as the number of active TCP connections N becomes large. However, the average queue length of RED is observed to increase almost linearly with the number of TCP connections N . Generally, the number of TCP connections changes according to time. For this reason, Fig. 7 would be more desirable than Fig. 6 in the sense that the average queue length does not change excessively with the number of TCP connections.

Finally, the result when using \mathcal{H}_ϕ (convex) as the packet marking function f is shown in Fig. 8. This figure shows that the average queue length in steady state becomes large as the number of TCP connections N becomes large. Moreover, it indicates that the gradient (dx^*/dx) of Eq. (9) is significantly influenced by the number of TCP connections. For this reason, considering the steady state behavior and the transient state behavior of RED, we consider that \mathcal{H}_ϕ (convex) is unsuitable for application as the function f .




Figure 6. Relation between the number of active TCP connections N and RED queue occupancy with \mathcal{F}_ϕ (linear) for $\phi = 1.0$.




Figure 7. Relation between the number of active TCP connections N and RED queue occupancy with \mathcal{G}_ϕ (concave) for $\phi = 1.0$.

From the above observations, we conclude that \mathcal{G}_ϕ (concave) is the most suitable for application as the packet marking function f . Although the average queue length in steady state can be small when either \mathcal{F}_ϕ (linear) or \mathcal{H}_ϕ (convex) is used, there is a drawback that the transient state performance is quite sensitive to the queue occupancy. On the other hand, when \mathcal{G}_ϕ (concave) is used, variation in the queue occupancy and the number of TCP connections negligibly affects the transient state behavior of RED. In summary, when \mathcal{G}_ϕ (concave) is used as the packet marking function, the transient state performance of RED and robustness to variations in the number of active TCP connections are improved.




Figure 8. Relation between the number of TCP connections N and RED queue occupancy with \mathcal{H}_ϕ (convex) for $\phi = 1.0$.

4. CONCLUSION

In this paper, we have discussed how the packet loss probability of RED should be determined from its average queue length by utilizing the results of steady state analyses of TCP¹⁸ and RED.¹⁶ By examining several numerical examples, we have investigated the performance of the RED router in three cases — when the function that determines the packet loss probability is either linear, concave, or convex. Consequently, we have found that the transient behavior and the robustness to variation in the number of TCP connections can be improved by using a concave function for determining the packet loss probability of RED.

Our analytic results have clearly shown that the control of RED, which discards arriving packets with a probability proportional to its average queue length, has several problems associated with steady state behavior and robustness. In recent years, various AQM mechanisms which solve these drawbacks of RED have been proposed. However, most of these AQM mechanisms use a function which is linear to the average queue length for determining the packet marking probability. Namely, the problems found in this paper have not been taken into account. As future works, by applying our analytic result obtained in this paper we therefore intend to design an AQM mechanism by taking account of not only steady state behavior but also transient state behavior and robustness.

Because of space limitation, comparison between analytic and simulation results cannot be included in this paper although several simulation results that we have obtained show validity of our approximate analysis. However, more extensive study using simulation would be appropriate.

REFERENCES

1. B. Braden et al., “Recommendations on queue management and congestion avoidance in the Internet,” *Request for Comments (RFC) 2309*, Apr. 1998.
2. S. Floyd and K. Fall, “Promoting the use of end-to-end congestion control in the Internet,” *IEEE Transactions on Networking*, May 1999.
3. S. Floyd and V. Jacobson, “Random early detection gateways for congestion avoidance,” *IEEE/ACM Transactions on Networking*, vol. 1, pp. 397–413, Aug. 1993.
4. Y. Zhang and L. Qiu, “Understanding the end-to-end performance impact of RED in a heterogeneous environment,” *Cornell CS Technical Report 2000-1802*, July 2000.

5. M. Christiansen, K. Jeffay, D. Ott, and F. D. Smith, "Tuning RED for web traffic," in *Proceedings of ACM SIGCOMM 2000*, pp. 139–150, Aug. 2000.
6. M. May, J. Bolot, C. Diot, and B. Lyles, "Reasons not to deploy RED," in *Proceedings of IWQoS '99*, pp. 260–262, Mar. 1999.
7. V. Misra, W.-B. Gong, and D. Towsley, "Fluid-based analysis of a network of AQM routers supporting TCP flows with an application to RED," in *Proceedings of ACM SIGCOMM 2000*, pp. 151–160, Aug. 2000.
8. C. Brandauer, G. Iannaccone, C. Diot, and T. Ziegler, "Comparison of tail drop and active queue management performance for bulk-data and Web-like Internet traffic," in *Proceedings of the Sixth IEEE Symposium on Computers and Communications, IEEE ISCC 2001*, pp. 122–129, July 2001.
9. S. Floyd, "Discussions on setting RED parameters," Nov. 1997. available at <http://www.aciri.org/floyd/REDparameters.txt>.
10. S. Floyd, "Recommendations on using the gentle variant of RED," May 2000. available at <http://www.aciri.org/floyd/red/gentle.html>.
11. D. Lin and R. Morris, "Dynamics of random early detection," in *Proceedings of ACM SIGCOMM '97*, pp. 127–137, Oct. 1997.
12. J. Aweya, M. Ouellette, D. Y. Montuno, and A. Chapman, "An adaptive buffer management mechanism for improving TCP behavior under heavy load," in *Proceedings of IEEE International Conference on Communications (ICC 2001)*, pp. 3217–3223, 2001.
13. T. J. Ott, T. V. Lakshman, and L. Wong, "SRED: Stabilized RED," in *Proceedings of IEEE INFOCOM '99*, pp. 1346–1355, Mar. 1999.
14. W.-C. Feng, D. D. Kandlur, D. Saha, and K. G. Shin, "Techniques for eliminating packet loss in congested TCP/IP networks," Tech. Rep. CSE-TR-349-97, Oct. 1997.
15. H. Wang and K. G. Shin, "Refined design of random early detection gateways," in *Proceedings of IEEE GLOBECOM '99*, pp. 769–775, 1999.
16. J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP throughput: a simple model and its empirical validation," in *Proceedings of ACM SIGCOMM '98*, pp. 303–314, Sept. 1998.
17. D. Bertsekas and R. Gallager, *Data Networks*. Englewood Cliffs, New Jersey: Prentice-Hall, 1987.
18. H. Ohsaki and M. Murata, "Steady state analysis of the RED gateway: stability, transient behavior, and parameter setting," *IEICE Transactions on Communications*, vol. E85-B, pp. 107–115, Jan. 2002.
19. M. May, T. Bonald, and J.-C. Bolot, "Analytic evaluation of RED performance," in *Proceedings of IEEE INFOCOM 2000*, pp. 1415–1424, Mar. 2000.