LRED: A Robust and Responsive AQM Algorithm Using Packet Loss Ratio Measurement

Chonggang Wang, *Member, IEEE*, Jiangchuan Liu, *Member, IEEE*, Bo Li, *Senior Member, IEEE*, Kazem Sohraby, *Senior Member, IEEE*, and Y. Thomas Hou, *Senior Member, IEEE*

Abstract— Active queue management (AQM) is an effective means to enhance congestion control, and to achieve tradeoff between link utilization and delay. The *de facto* standard, Random Early Detection (RED), and many of its variants employ queue length as a congestion indicator to trigger packet dropping. Despite their simplicity, these approaches often suffer from unstable behaviors in a dynamic network. Adaptive parameter settings, though might solve the problem, remain difficult in such a complex system. Recent proposals based on analytical TCP control and AQM models suggest the use of both queue length and traffic input rate as congestion indicators, which effectively enhances stability. Their response time generally increases however, leading to frequent buffer overflow and emptiness. In this paper, we propose a novel AQM algorithm that achieves fast response time and yet good robustness. The algorithm, called Loss Ratio based RED (LRED), measures the latest packet loss ratio, and uses it as a complement to queue length for adaptively adjusting the packet drop probability. We develop an analytical model for LRED, which demonstrates that LRED is responsive even if the number of TCP flows and their persisting times vary significantly. It also provides a general guideline for the parameter settings in LRED. The performance of LRED is further examined under various simulated network environments, and compared to existing AQM algorithms. Our simulation results show that, with comparable complexities, LRED achieves shorter response time and higher robustness. More importantly, it trades off the goodput with queue length better than existing algorithms, enabling flexible system configurations.

Index Terms—Active queue management, congestion control, TCP, packet loss ratio.

———————————— ◆ ————————————

1 INTRODUCTION

Buffer management for Internet routers plays an important role in congestion control [1][2]. However, the two main objectives of buffer management, namely, high link utilization and low packet queuing delay, often conflict with each other. Specifically, given that most end-nodes employ the responsive Additive Increase and Multiplicative Decrease (AIMD) TCP congestion control, a small buffer generally achieves a low queuing delay, but suffers from excessive packet losses and low link utilization, and vice versa. In addition, a simple policy like the widely used First-In-First-Out (FIFO) Tail-Drop often causes strong cor-

relations among packet losses, resulting in the well-known "TCP synchronization" problem [3].

To mitigate such problems, Active Queue Management (AQM) has been introduced in recent years [2] [3] [5]-[14]. The basic idea is to actively trigger packet dropping (or marking provided explicit congestion notification (ECN) [4] is enabled) before buffer overflow. Obviously, the drop probability should depend on the degree of congestion. The *de facto* AQM standard, Random Early Detection (RED) [5], and many of its variants employ queue length as a congestion indicator to trigger packet dropping. Despite their simplicity, these approaches often suffer from unstable behaviors in a dynamic network. Adaptive parameter settings, though might solve the problem, remain difficult in such a complex system. Recent proposals based on analytical TCP control and AQM models suggest the use of both queue length and traffic input rate as congestion indicators, which effectively enhances stability. Their response time generally increases however, leading to frequent buffer overflow and emptiness.

In this paper, we argue that packet loss ratio, which has never been explored in previous AQM studies, is another important index. The packet loss ratio is measured as the fraction of the packets dropped by the router and is updated over time. Intuitively, for a well-designed AQM algorithm, the loss ratio should be close to the desired drop probability in a steady-state, and it deviates from the desired drop probability if the buffer is (or tends to) overflow

————————————————

- *Chonggang Wang is with the Department of Electrical Engineering, University of Arkansas, Fayetteville, AR 72701. E-mail, cgwang@ieee.org; Phone: +1-479-575-8655; Fax: +1-479-575-7967.*
- *Jiangchuan Liu is with the Schoold of Computer Science, Simon Fraser University, Burnaby, British Columbia, Canada V5A 1S6. E-mail: jcliu@cs.sfu.ca; Phone: +1-604-291-4336; Fax: +1-604-291-3045.*
- *Bo Li is with the Department of Computer Science, Hong Kong University of Science and Technology, Hong Kong, P.R.China. E-mail: bli@cs.ust.hk; Phone: +852-2358-6976; Fax: +852-2358-1477. The Corresponding Author.*
- *Kazem Sohraby is with the Department of Electrical Engineering, University of Arkansas, Fayetteville, AR 72701. E-mail, sohraby@uark.edu; Phone: +1-479-575-2635; Fax: +1-479-575-7967.*
- *Y. Thomas Hou is with the Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA 24061. E-mail: thou@vt.edu; Phone: +1-540-231-2950; Fax: +1-540-231-8292.*

or empty. In other words, an increasing packet loss ratio implies that congestion occurs, and a decreasing implies that the congestion is relieved.

Given the above observations, we propose a novel AQM algorithm, LRED, which incorporates the packet loss ratio as a complement to queue length for congestion estimation. We stress two salient features of this hybrid design: First, the calculation is simple and fast for both measures, which is desirable for high-throughput routers; second, it enables a multi-granular update for the packet drop probability: upon each packet arrival, we can use the instantaneous queue length mismatch to update the drop probability, so as to keep the queue length around an expected value; on a coarser grain, we can adjust the drop probability according to the packet loss ratio, so as to expedite control response. We have developed an analytical model for LRED, which suggests that this multi-granular update improves not only the stability of the system, but also its responsiveness.

The performance of LRED is further examined under various simulated network environments, and compared to existing AQM algorithms, including Adaptive Virtual Queue (AVQ) [10], Proportional-Integra (PI) [12], and Random Exponentially Marking (REM) [13]. Our simulation results reveal that, with comparable complexities, LRED achieves shorter response time and better robustness. More importantly, it enables better tradeoff between the goodput and queue length than existing algorithms, leading to flexible system configurations.

The remainder of this paper is organized as follows. Section 2 briefly introduces the existing AQM algorithms. In Section 3, we offer an overview of the LRED algorithm. Section 4 develops an analytical model for the combined TCP/AQM system and discusses its general properties. Based on this model, we analyze the stability and responsiveness of LRED and other AQM algorithms in Section 5. The simulation results for LRED are presented in Section 6. Finally, we conclude the paper and give out future works in Section 7.

2 RELATED WORK

There have been numerous proposals on AQM in the past decade, and RED [5] is probably the most widely studied algorithm. RED uses the average queue length to calculate the packet drop probability and to regulate the queue length accordingly. Specifically, when the average queue length is greater than a pre-configured threshold (min_{th}), RED begins to drop newly arrived packets; the dropping probability increases linearly to the average queue length with a slope of max_p. Despite its simplicity, the optimal parameter configuration for RED remains a daunting task. Hence, several enhancements, like S-RED [7] and ARED [8], have been introduced to adaptively configure the parameters. Another variation is BLUE [9], which calculates the packet drop probability based only on two events: buffer overflow and emptiness. When the buffer is overflow, it increases packet drop probability by δ_1, and, when empty, decreases by δ_2. Nevertheless, adaptive settings in such a complex system are still difficult. BLUE can not well regulate the queue length to an expected value.


Fig. 1. Packet loss ratio measurement in LRED.

On the other hand, AVQ [10] uses input rate $x(t)$ to control packet drop and to achieve an expected link utility γ. Basically, the packet drop probability is proportional to the mismatch between $x(t)$ and γ. Through maintaining a virtual queue, AVQ deterministically drops packets upon each new packet arrival, realizing the same effect of probabilistic packet drop. AVQ achieves lower average queue length and higher link utility than RED and its variants [10].

Recently, some advanced algorithms use the queue length and input rate jointly to achieve better performance. One example is PI [12], which regulates the queue length to an expected value q_0 according to the queue mismatch and its integral. The latter is closely related to the input rate mismatch. If the network states are known *a priori*, optimal parameters of PI can be determined through a control-theoretical model. However, in dynamic networks, PI may have to use a conservative setting to ensure stability, yielding long response time. Another example is REM [13], which uses a linear combination of the queue mismatch and input rate mismatch to calculate the drop probability, and the input rate mismatch is equivalently simplified to the queue variance between two adjacent length samples.

LRED employs the packet loss ratio as a complement to the queue length for AQM. To the best of our knowledge, it has never been explored in the previous studies. The use of packet loss ratio enables LRED to catch network dynamics in time, thus achieving fast control response and better performance in terms of goodput, average queue length, and packet loss ratio.

3 OVERVIEW OF LRED

Similar to most existing AQM algorithms, LRED keeps a packet drop probability p, which is adaptively updated upon each packet arrival, and the incoming packets are dropped with probability p from the tail of the queue. The calculation of the drop probability is relatively simple, following two design rules: 1) when the queue length is close to an expected steady-state length q_0, the drop probability should be close to the packet loss ratio; and 2) when the queue length becomes larger (or smaller), the drop probability should be increased (or decreased) to regulate the queue length.

To achieve an adaptive yet stable estimation for the packet loss ratio, LRED estimates the ratio in every measurement period (t_m). Let $l(k)$ be the packet loss ratio of the k-th measurement, which is calculated as the number of dropped packets over the total number of packets arrived during the latest M periods (see Fig. 1); that is:
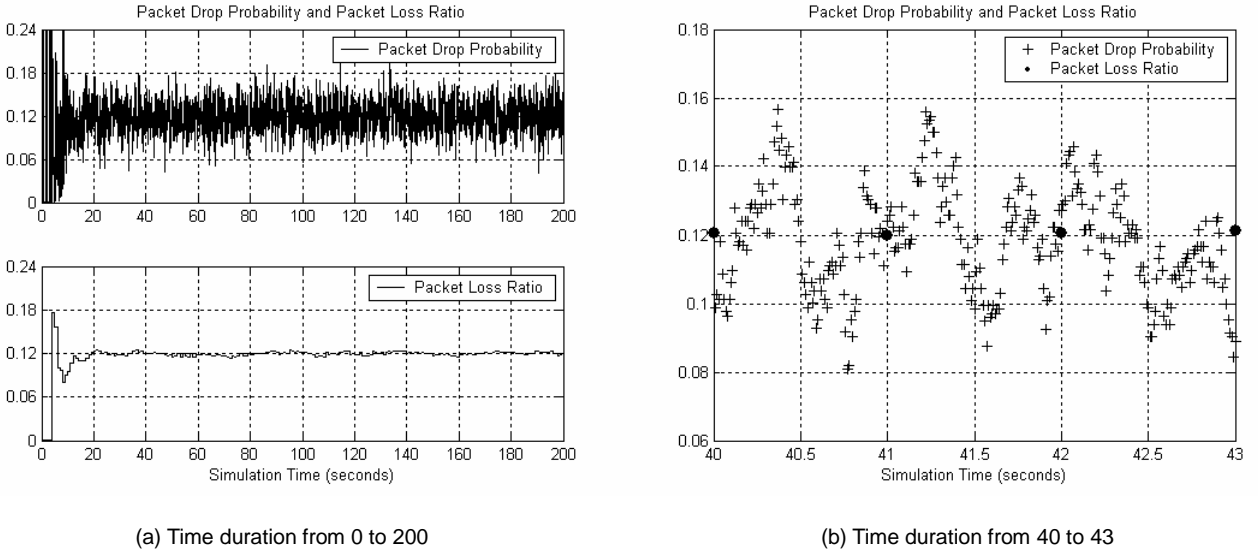
(a) Time duration from 0 to 200

(b) Time duration from 40 to 43

Fig. 2. Evolution of packet drop probability (p) and packet loss ratio ($l(k)$) in a simulation.

$$l(k) = \Big(\sum_{i=0}^{M-1} N_d(k-i)\Big) / \Big(\sum_{i=0}^{M-1} N_a(k-i)\Big), \qquad (1)$$

where $N_d(k)$ is the number of packets dropped in the k-th measurement period, and $N_a(k)$ is the number of packets arrived in the k-th measurement period. The estimated packet loss ratio $\overline{l(k)}$ is calculated based on $l(k)$ using an exponential weighted moving average (EWMA), as follows:

$$\overline{l(k)} = \overline{l(k-1)} \cdot w_m + (1 - w_m) \cdot l(k), \qquad (2)$$

where w_m is a weighting factor. In order to promptly catch loss changes, we set w_m to a small value.

Given the estimated packet loss ratio and the instantaneous queue length, a straightforward way to meet the design rules (to control the queue length around an expected steady-state value) is to use a linear function, $p = \overline{l(k)} + \alpha(q - q_0)$, where p is the packet drop probability and q is the instantaneous queue length. It is however often difficult to choose an optimal α. In particular, if $\overline{l(k)}$ is large and α is set too small, the packet drop probability for a small queue length q would still be quite high, resulting in low link utility. To avoid this, in LRED, we let p increase with the measured loss ratio; in other words, we have:

$$p = \overline{l(k)} + \beta\sqrt{\overline{l(k)}}(q - q_0), \qquad (3)$$

where $\beta > 0$ is a pre-configured constant. Intuitively, the adjustment of the packet drop probability should be related to traffic rate change, in order to avoid buffer overflow or buffer emptiness. In Eq. (3), this adjustment is set to $\beta\sqrt{\overline{l(k)}}(q - q_0)$, which uses square-root function of $\overline{l(k)}$. This square-root function can guarantee that $\beta\sqrt{\overline{l(k)}}(q - q_0)$ is proportional to the percentage of traffic rate change, if we assume traffic flows are TCP. This is because the TCP throughput is reversely proportional to the square-root of the stable packet loss ratio [15]. Furthermore, given that $(q - q_0)$ depends on both the value of stable traffic rate and the percentage of traffic rate change, the product of $(q - q_0)$ and $\sqrt{\overline{l(k)}}$ depends on the percentage of traffic rate

change only. Our analysis in Section 5 also shows that Eq. (3) ensures stable control. In the case of non-uniform packet sizes, different drop probabilities may have to be applied to packets with different sizes. Specifically, larger packets might have higher drop probabilities. In this case, we can maintain the expected queue length counted in bytes, and therefore improve the fairness among TCP connections with different packet size.

It can be seen from Eqs. (2) and (3) that LRED updates $\overline{l(k)}$ and p at different time scales. $\overline{l(k)}$ is updated every measurement period (see Fig. 1), while the packet drop probability p is re-calculated each time a new packet comes. In a steady-state, $\overline{l(k)}$ could converge to a stable value. However, the AIMD mechanism in TCP implies that a TCP sender will always try to decrease (or increase) its sending rate if it detects packet loss (or not); therefore the queue length in routers will unavoidably fluctuate and the packet drop probability p thus fluctuates around $\overline{l(k)}$. This is illustrated in Fig. 2, where we can see $\overline{l(40)} \approx \overline{l(41)} \approx \overline{l(42)} \approx \overline{l(43)} \approx 0.12$; When $t = 41.0$, $\overline{l(41)} \approx 0.12$; When $t < 41.0$, the packet drop probability fluctuates around 0.12.

We can observe from Eq. (3) and Fig. 2 that: 1) In between two adjacent instants k and $k+1$, when p is updated according to Eq. (3), $\overline{l(k)}$ remains constant and therefore p dependents only on the control error $(q - q_0)$; or more precisely, it is proportional to $(q - q_0)$ only; 2) In a steady-state, $\overline{l(k)}$ will be close to the stable value, and therefore p is still proportional to the control error $(q - q_0)$ in this case. As such, we believe that LRED is closer to a proportional controller.

Detailed operations for LRED can be found in Fig. 3 where q_0 is assumed to 100. We will further discuss its parameter settings and prove its stability with the square-root function in Section 5.

```
/* Parameters */
q_o : the expected value of queue length;              q  : the current instantaneous queue length;
p  : the calculated packet drop probability;           lossRatio: the estimated packet loss ratio;
β  : weighting factor (=0.001 in this paper);          t_m : the measurement period;
M : the number of the latest periods to measure lossRatio ;   w_m : the measurement weight;

/* Initialization */
arrPktNum=0;      dropPktNum=0;      allArrNum=0;       allDropNum=0;    index=0;         q_o =100;
M=4;              w_m =0.1;          β =0.001;          t_m =1.0;        lossRatio=0.0;

LossRatioMeasure() /* Called every t_m seconds */      Enqueue() /* Called upon each packet arrival*/
1       dropNum[index]=dropPktNum;                      1       arrPktNum++;
2       arrNum[index]=arrPktNum;                        2       p = lossRatio + β√lossRatio(q − q_o) ;
3       arrPktNum=0;                                    3       p = max(0, min(1, p)) ;
4       dropPktNum=0;                                   4       random=uniformRandom(0, 1);
5       index++;                                        5       if (buffer is full) {
6       if (index==M)         index=0;                  6               Drop the packet;
7       for (i=0; i<M; i++) { allDropNum+=dropNum[i]; } 7               dropPktNum++;
8       for (i=0; i<M; i++) { allArrNum+=arrNum[i]; }   8       }
9       lossRatioTemp=allDropNum/allArrNum;             9       else if (random> p )  {Enqueue the packet; }
10      lossRatio=lossRatio* w_m +lossRatioTemp*(1- w_m ); 10    else {   Drop the packet;
11      allDropNum=0;                                   11              dropPktNum++;
12      allArrNum=0;                                    12      }
```

Fig. 3. Pseudocode of the LRED algorithm.

4 MODEL OF COMBINED TCP/AQM SYSTEMS

In this section, we first review an analytical model for a combined TCP/AQM system [11][14]. We then make several important observations on the model, which serve as the basis for designing a stable AQM algorithm, in particular, for LRED proposed in this paper.

4.1 The Combined TCP/AQM Models

We consider an abstract network of a single bottleneck with multiple TCP connections. Its status can be represented by a 3-tuple (N, C, R), where N is the number of TCP flows, C is the bottleneck link capacity, and R is the round trip time (RTT). With the assumption that TCP flows are long-lived persistent and packet size is constant, the system equation for the TCP congestion window (w) and the queue length (q) can be approximated as [11]:

$$\dot{w} = f(w, p) = 1 / R - (w(t)w(t - R) / \eta R) * p(t - R) , \quad (4)$$
$$\dot{q} = g(w, p) = (N / R) * w(t) - C , \quad (5)$$

where p is the packet drop probability, and η is a parameter depending on TCP implementations. Let b be the number of packets acknowledged by a received acknowledgement (ACK), we have $\eta = 3 / 2b$ [15]. Eqs. (4) and (5) give a model of combined TCP/AQM system with long-lived flows and constant packet size. It is worth noting that the actual traffic mix in real Internet might be more complex; for example, flows can start and end all the time and the packet sizes may vary for different applications. Yet, existing studies have shown that this model offers a good approximation [11][12][14].

If we let $f(w, p) = 0$ and $g(w, p) = 0$, the TCP congestion window w_0 and packet drop probability p_0 in a steady-state can be calculated as:

$$w_0 = \frac{RC}{N}, \quad p_0 = \frac{\eta}{w_0^2} = \frac{\eta N^2}{R^2 C^2} < 1 . \quad (6)$$

For ease of exposition, we assume each ACK acknowledges only one packet; therefore $b = 1$ and $\eta = 1.5$. The steady-state throughput of a single TCP flow given by the above model becomes $C / N = \sqrt{3/2} / (R\sqrt{p_0})$, which is consistent with the well-known TCP throughput equations [15].

From Eqs. (4) and (5), it is clear that the combined TCP/AQM system is non-linear. Let $\delta w = w - w_0$ and $\delta p = p - p_0$ be the mismatches (i.e., the deviations from the steady-state values) for the TCP congestion window and the packet drop probability, respectively. Eqs. (4) and (5) can be locally linearized around a stable point (w_0, p_0) by assuming $w(t) \approx w(t - R)$, as follows:

$$\delta\dot{w} = \frac{\partial f}{\partial w}\delta w + \frac{\partial f}{\partial p}\delta p = -[K_{11}\delta w + K_{12}\delta p(t - R)], \quad (7)$$

$$\delta\dot{q} = \frac{\partial g}{\partial w}\delta w + \frac{\partial g}{\partial p}\delta p = K_{21}\delta w , \quad (8)$$

where $K_{11} = 2N /(R^2 C)$, $K_{12} = RC^2 /(\eta N^2)$, and $K_{21} = N / R$ [11]. All of them are positive constants related to the network parameters.

Applying Laplace transform to Eqs. (7) and (8), we have the following system equations:

$$s * W(s) = -[K_{11} * W(s) + K_{12} * P(s) * e^{-Rs}], \quad (9)$$

$$s * Q(s) = K_{21} * W(s) . \quad (10)$$

In Eqs. (9) and (10), there are three unknown variables: $W(s)$, $P(s)$, and $Q(s)$. Hence, it is necessary to find one more equation to solve the problem. This can be achieved by bridging $P(s)$ and $Q(s)$ through AQM. As discussed earlier, our proposed LRED is a proportional controller; we thus focus on the proportional AQM control in our analysis. Consider an AQM proportional control mechanism that determines p according to the instantaneous queue length

q. Its general control equation can be formulated as:

$$\delta p = H_c * \delta q. \tag{11}$$

The corresponding system equation can be written as:

$$P(s) = H_c * Q(s), \tag{12}$$

where $H_c > 0$ is a pre-defined parameter.

4.2 General Properties of Proportional AQM Control

From Eqs. (9), (10), and (12), we can obtain the characteristic equation for such a system as:

$$s^2 + K_{11}s + K_cH_ce^{-Rs} = 0, \tag{13}$$

where $K_c = K_{12}K_{21} = C^2/\eta N$.

The stability of the above system depends on whether the root of Eq. (13), $s = \sigma + j\omega$, lies in the left half-complex plane. To facilitate its stability analysis, we have made the following observations on Eq. (13).

First, if the root of Eq. (13) strictly lies in the left half-complex plane, the combined system, defined by network parameters (N, C, R) and AQM control parameter H_c, is stable. Hence, given (N, C, R), we can choose H_c to satisfy this condition. On the other hand, given control parameter H_c, only some of the system (N, C, R) can be stably controlled. Specifically, when $R = 0$, the root of Eq. (13) is:

$$s = s_0 = (-K_{11} \pm \sqrt{K_{11}^2 - 4K_cH_c})/2, \tag{14}$$

which strictly lies in the left half-complex plane irrespective of $K_{11}^2 \ge 4K_cH_c$ or $K_{11}^2 < 4K_cH_c$. Therefore, the system with zero delay is stable.

When $R > 0$ and letting $s = \sigma + j\omega$, the root of Eq. (13) can be calculated as follows:

$$\sigma^2 - \omega^2 + K_{11}\sigma + K_cH_ce^{-R\sigma}\cos(R\omega) = 0, \tag{15}$$

$$2\sigma\omega + K_{11}\omega - K_cH_ce^{-R\sigma}\sin(R\omega) = 0. \tag{16}$$

Clearly, the imaginary part (ω) of root s is nonzero; otherwise, Eq. (15) will be invalid. Also note that s changes continuously in the complex plane when N, C, R, or H_c changes continuously. Assume the first time that s meets the imaginary axis is at $R = R^+$. When $0 \le R < R^+$, s should strictly lie in the left half-complex plane, and the system is thus stable. The remaining problem therefore is to find the value of R^+.

We first consider the absolute imaginary root $(s = j\omega)$ with $\omega > 0$ (the case of $\omega < 0$ is symmetric). When $R > 0$, Eq. (13) can be re-written as:

$$T(s) = (e^{-Rs}K_cH_c)/(s(s + K_{11})) = -1. \tag{17}$$

Given Eq. (17), the following conditions on magnitude and angles must be met,

$$|T(j\omega)| = 1, \quad \angle T(j\omega) = (2k+1)\pi, \quad k = 0, \pm 1, \pm 2,$$

and ω can be calculated as:

$$\omega(N, C, R, H_c) = \omega = \sqrt{y(N, C, R, H_c)/2}, \tag{18}$$

$$y(N, C, R, H_c) = \sqrt{K_{11}^4 + 4K_c^2H_c^2} - K_{11}^2, \tag{19}$$

$$R\omega + \arctan(\frac{\omega}{K_{11}}) + \frac{\pi}{2} = (2k+1)\pi, \quad k = 0, 1, 2. \tag{20}$$

Since K_{11} is a decreasing function of R, and K_c is independent of R, we have that $\omega(N, C, R, H_c)$ is an increasing function of R, if H_c is independent of or an increasing function of R. Therefore, R^+ corresponds to the smallest R satisfying Eqs. (13) and (17). Now the problem is to de-

termine the value of k that yields the smallest R, which can be solved through the following Lemma.

Lemma 1. *When $k = 0$, Eqs. (18)-(20) yield the smallest value of R and ω, if H_c is independent of or an increasing function of R. Also Eq. (20) can be simplified to:*

$$R\omega + \arctan(\omega/K_{11}) = \pi/2. \tag{21}$$

Proof. We prove this by contradiction. Assume that $k \to (R_k, \omega_k)$, $k > 0$, $R_k > 0$, and $\omega_k > 0$. According to Eq. (20), we have:

$$R_k\omega_k - R_0\omega_0 = 2k\pi + [\arctan(\frac{\omega_0}{K_{11}}) - \arctan(\frac{\omega_k}{K_{11}})]$$
$$> 2k\pi - \pi/2 > 0.$$

Assume $R_k \le R_0$, we have $\omega_k \le \omega_0$ according to Eqs. (18) and (19), if H_c is independent of or an increasing function of R. Since $R_k \le R_0$ and $\omega_k \le \omega_0$, we have $R_k\omega_k \le R_0\omega_0$, which contradicts to that $R_k\omega_k - R_0\omega_0 > 0$. Hence, $R_k > R_0$, or equivalently, Eqs. (18)-(20) yield the smallest value of R and ω when $k = 0$. $\square$

Lemma 2. *Given network parameters (N, C) and AQM control parameter H_c. Assume that R^+ satisfies:*

$$R^+\omega + \arctan(\omega/K_{11}) = \pi/2, \quad R^+ > 0, \tag{22}$$

where ω is the solution to Eqs. (18) and (19). If function $y(N, C, R, H_c)$ in Eq. (19) is an increasing function of R, then the system is stable for all $R < R^+$, and R^+ is unique.

Proof. Since $y(N, C, R, H_c)$ is an increasing function of R, according to Lemma 1, R^+ is the smallest R that satisfy Eqs. (13) and (14), and is unique. It also means that the first time the root meets the imaginary axis is at $R = R^+$. Therefore the system is stable for all $R < R^+$. $\square$

Lemma 3. *Given network parameters (C, R) and AQM control parameter H_c. Assume that N^- satisfies:*

$$R\omega + \arctan(\omega/K_{11}) = \pi/2, \quad N^- > 0, \tag{23}$$

where ω is the solution to Eqs. (18) and (19). If function $y(N, C, R, H_c)$ in Eq. (18) is an increasing function of R and a decreasing function of N, then the system is stable for $N > N^-$.

Proof. Let $R(N)$ be the solution to Eq. (13) with $N > N^-$. If $R < R(N)$, from Lemma 2, the system is stable for all $N > N^-$. Since $y(N, C, R, H_c)$ in Eq. (18) is a decreasing function of N, we have:

$$\omega(N) = \omega(N, C, R, H_c) < \omega(N^-, C, R, H_c) = \omega(N^-).$$

Moreover, K_{11} is an increasing function of N, which implies:

$$R\omega(N) + \arctan(\omega(N)/K_{11})$$
$$< R\omega(N^-) + \arctan(\omega(N^-)/K_{11}) = \pi/2.$$

Since $R(N)\omega(N) + \arctan[\omega(N)/K_{11}] = 0.5\pi$, we have $R(N) > R$, and hence, for all $N > N^-$, the system is stable. $\square$

Lemma 4. *Given network parameters (N, C, R), and assume that H_c^* satisfies:*

$$R\omega + \arctan(\omega / K_{11}) = \pi / 2 , \ H_c^* > 0 , \qquad (24)$$

where ω is given by Eqs. (18) and (19). If function $y(N,C,R,H_c)$ in Eq. (19) is an increasing function of both R and H_c, then the system is stable for all $H_c < H_c^*$.

Proof. Similar to that of Lemma 3. □

Theorem 1. Let the network parameters be (N^-, C, R^+), and assume that H_c^+ satisfies:

$$R^+\omega + \arctan(\omega / K_{11}) = \pi / 2, \ H_c^+ > 0, \qquad (25)$$

where ω is given in Eqs. (18) and (19). If function $y(N,C,R,H_c)$ in Eq. (19) is an increasing function of R, a decreasing function of N, and an increasing function of H_c, then the system is stable for $H_c < H_c^+$, $N > N^-$, and $R < R^+$.

Proof. Directly follows Lemmas 2 through 4. □

5 ANALYSIS OF LRED AND PARAMETER SETTINGS

5.1 Stability Analysis of LRED

We now investigate the stability of LRED and discuss the settings of several important parameters, in particular, β. Since the packet loss ratio is close to the stable packet drop probability p_0 in LRED, i.e., $\bar{l}(k) \approx p_0$, we can approximately rewrite Eq. (3) as:

$$p = p_0 + \beta\sqrt{p_0}(q - q_0), \qquad (26)$$

where, according to Eq. (6), $p_0 = \eta N^2 / (R^2 C^2)$. It follows that:

$$\delta p = \beta\sqrt{p_0}\delta q, \qquad (27)$$

or

$$P(s) = \beta\sqrt{p_0}\, Q(s), \qquad (28)$$

and the system transfer function of LRED (see Eq. (12)) is thus given by:

$$H_c = P(s) / Q(s) = \beta\sqrt{p_0}. \qquad (29)$$

Substituting Eq. (29) for H_c in Eqs. (18) and (19), we have the following Lemma.

Lemma 5. *For LRED, function $y(N,C,R,H_c)$ in Eq. (19) is a decreasing function of N, and an increasing function of β. Moreover, if $\beta < \sqrt{2\eta}(2N)^2 /(R^3 C^3)$, it is an increasing function of R.*

Proof. For LRED, $y(N,C,R,H_c)$ in Eq. (18) can be calculated as:

$$y(N,C,R,H_c) = \sqrt{(\frac{2N}{R^2C})^4 + \frac{4\beta^2 C^2}{\eta R^2}} - (\frac{2N}{R^2C})^2. \qquad (30)$$

It can be easily observed that $y(N,C,R,H_c)$ is a decreasing function of N and an increasing function of β. We now consider its relation with R. The derivative of function y with respect to R is:

$$\frac{\partial y}{\partial R} = \frac{-\dfrac{8(2N)^4}{R^9 C^4} - \dfrac{8\beta^2 C^2}{\eta R^3}}{2\sqrt{(\dfrac{2N}{R^2 C})^4 + \dfrac{4\beta^2 C^2}{\eta R^2}}} + \frac{4(2N)^2}{R^5 C^2} .$$

$$= \frac{-f_1(R)}{f_2(R)} + f_3(R)$$

Note that y is an increasing function of R if $\partial y / \partial R > 0$, which is equivalent to:

$$[f_2(R) \cdot f_3(R)]^2 - f_1^2(R) = \frac{64\beta^2[2\eta(2N)^4 - \beta^2 R^6 C^6]}{\eta^2 R^{12} C^2} > 0 .$$

It follows that $\beta < \sqrt{2\eta}(2N)^2 /(R^3 C^3)$. □

Theorem 2. *Given network parameters (N^-, C, R^+), and assume that β_0 satisfies:*

$$R^+\omega + \arctan(\omega / K_{11}) = \pi / 2, \ \ \beta_0 > 0 . \qquad (31)$$

where ω is defined in Eqs. (18) and (19), and $H_c = \beta\sqrt{p_0}$ in LRED. If $\beta < \beta^+ = \min(\beta_0, \sqrt{2\eta}(2N^-)^4 /(R^{+3} C^3))$, the system is stable for any $N > N^-$ and $R < R^+$.

Proof. Directly follows Lemma 5 and Theorem 1. □

Given Theorem 2, it is easy to find β that guarantees the stability of the system. For example, consider a network in which the mean packet size = 500 bytes, $C = 2500$ packets/sec (or equivalently, 10 Mbps), $N^- = 300$, and $R^+ = 0.35$ seconds. The AQM parameter β^+ is thus 0.001, and, for any $\beta < \beta^+$, the system is stable with all $N > N^-$ and $R < R^+$.

5.2 Response Time Analysis and Comparison

Since enhancing stability and minimizing response time often conflict with each other, existing algorithms such as PI [12] and REM [13] have tried to find a tradeoff between them. If network parameters, in particular, N and R, are known *a priori*, these algorithms can achieve a stable control with minimized response time. In a dynamic network, however, it is difficult to accesses these parameters precisely. Hence, they generally resort to a conservative design that guarantees stability, but may sacrifice the corresponding response time. For example, the default parameter for PI in NS2 Simulator [16] is set based on a small N and large R. When N increases or R decreases, it will yield a long response time, though the system remains stable.

In this subsection, we provide a simple analysis on the response times of the typical AQM schemes. We focus on a highly dynamic scenario: at time $t = 0$, N TCP flows become active simultaneously, where N is large enough such that the buffer is fully filled before the system converges to a steady state with packet drop probability p_0 and expected queue length q_0.

We first consider PI [12], which periodically updates its packet drop probability with a sampling frequency f_{PI} (Hz). Each update is as follows:

$$p(k) = p(k-1) + a[q(k) - q_0] - b[q(k-1) - q_0], \qquad (32)$$

where $a > 0$ and $b > 0$ are two constants [12]. We can assume that $p(0) = 0$ and $q(k) \approx Q$ before reaching a steady-state, where Q is the maximal buffer size. It follows

that:

$$p(k) \approx k(a - b)(Q - q_0). \tag{33}$$

Denote $p(k_0) = p_0$. The lower bound of the response time of PI (RT^-_{PI}) is thus:

$$RT^-_{PI} \geq k_0 / f_{PI} = p_0 / ((Q - q_0)(a - b)f_{PI}). \tag{34}$$

In REM [13], packet drop probability is also periodically updated with a sampling frequency f_{REM} (Hz), as follows [13]:

$$p(k) = 1 - \phi^{-u(k)}, \tag{35}$$

$$u(k) = u(k-1) + \gamma[q(k) - (1 - \alpha)q(k-1) - \alpha q_0)], \tag{36}$$

$$u(k) = \max(0, u(k)), \tag{37}$$

where $\phi > 1$, $\gamma > 0$, and $\alpha > 0$ are three constants, and their optimal values are derived in [13]. Similarly, we can also assume that $u(0) = 0$, $u(k) << 1$, and $q(k) \approx Q$ before reaching the steady-state, which follows that:

$$u(k) \approx k\gamma\alpha(Q - q_0), \tag{38}$$

$$p(k) = 1 - \phi^{-u(k)} = 1 - \sum_{i=0}^{+\infty} \{([-u(k)\ln\phi]^i) / (i!)\}$$
$$\approx u(k)\ln\phi = k\gamma\alpha(Q - q_0)\ln\phi \tag{39}$$

and the lower bound of the response time for REM (RT^-_{REM}) is thus:

$$RT^-_{REM} \geq k_0 / f_{REM} = p_0 / (f_{REM}\gamma\alpha(Q - q_0)\ln\phi). \tag{40}$$

From Eqs. (34) and (40), we can see that the response times of PI or REM mainly depend on the following parameters: buffer size Q, expected queue length q_0, and packet drop probability p_0 (which is an increasing function of TCP flow number N, and a decreasing function of round-trip time R and link capacity C). Consequently, under heavy congestion or with a high drop probability, PI and REM suffer from long response times. When buffer size Q is small, the responsiveness of PI and REM becomes worse as well.

In LRED, the response time is mainly influenced by the period (t_m) to measure the packet loss ratio. In particular, under heavy traffic (e.g., when N is large and/or R is small), packets will be dropped frequently and the packet loss ratio can be accurately estimated within a couple of measurements. As a result, LRED is very responsive in this case, while PI and REM perform poorly given that p_0 is high. More importantly, when network conditions change dramatically, LRED can quickly converge to new steady states. When the traffic load is light (e.g., when N is small and/or R is large), there are few packet losses, and more rounds of measurement are thus needed for accurately estimating the stable packet loss ratio. In this case, the response time of LRED would slightly increase, but remain comparable to that of PI and REM, as will be shown in our simulations. In summary, LRED decouples the response time and packet drop probability, making its response time almost independent of the congestion levels.

6 SIMULATION RESULTS

In this section, we examine the performance of LRED through NS2 [16] simulations. We also compare it with existing AQM schemes, in particular, PI [12] and REM [13].
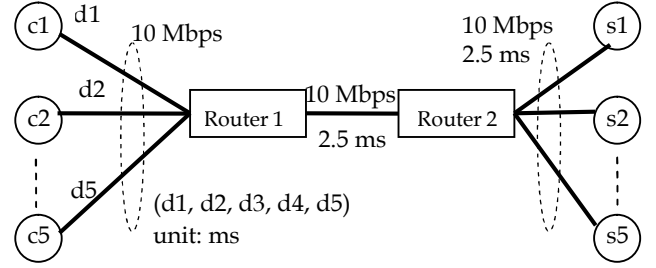
Fig. 4. Network topology for simulations.

Fig. 5. Experiment 1: Response times for the AQM schemes.

For loss ratio measurement in LRED, we set w_m =0.1, t_m =1.0 second, and $M = 4$; β is set to 0.001 according to Theorem 2. The network topology for the simulation is the commonly used dumb-bell topology (see Fig. 4). In this topology, 5 clients are linked to router 1, and 5 servers are behind router 2. The capacity of each link is 10 Mbps, and the link between router 1 and router 2 thus becomes a bottleneck. The link delay between router 2 and any server is 2.5 ms, and the delays between the clients (c1 through c5) and router 1 are heterogeneous, denoted by a 5-tuple $(d_1, d_2, d_3, d_4, d_5)$. All the flows in the network are uniformly distributed among the pairs of client c(i) and server s(i). The default packet size is 500 bytes. The buffer size of each router is 200 packets unless another value is explicitly configured such as in Experiment 6. We run each simulation for 200 seconds, which is long enough to observe both transient and steady-state behaviors of an AQM scheme.

The following parameter settings are used in our simulations. 1) For REM [13], $\phi = 1.001$, $\alpha = 0.1$, $\gamma = 0.001$, and the sampling interval is 2 ms. These values are adapted from [13] and [16]. 2) For PI [12], we use two settings: default and optimal, respectively denoted by PI and PI*. The default values for PI are a = 0.00001822, b = 0.00001816, and the sampling frequency is 170Hz, which are adapted from [16]; the optimal values for PI* are derived according to the design rules in [12], which depend on network parameters (N^-, C, R^+). Here, N^- is the minimum number of the TCP flows and R^+ is the maximal round-trip time. Hence, the optimal values for PI* are not fixed for the experiments.

In our study, we focus on the following key performance

(a) light congestion ($p_0 = 0.0025$)   (b) heavy congestion ($p_0 = 0.165$)
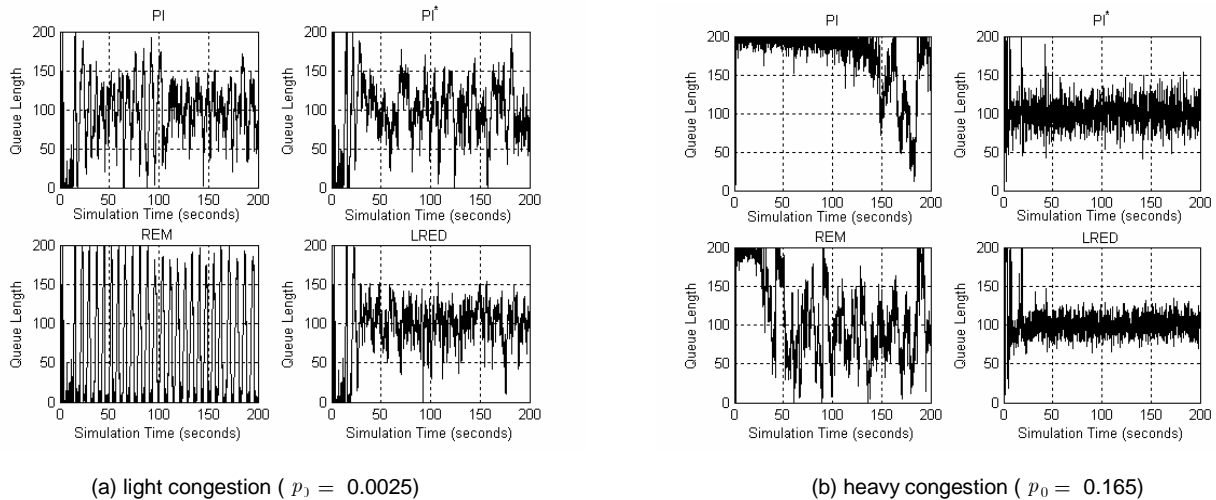
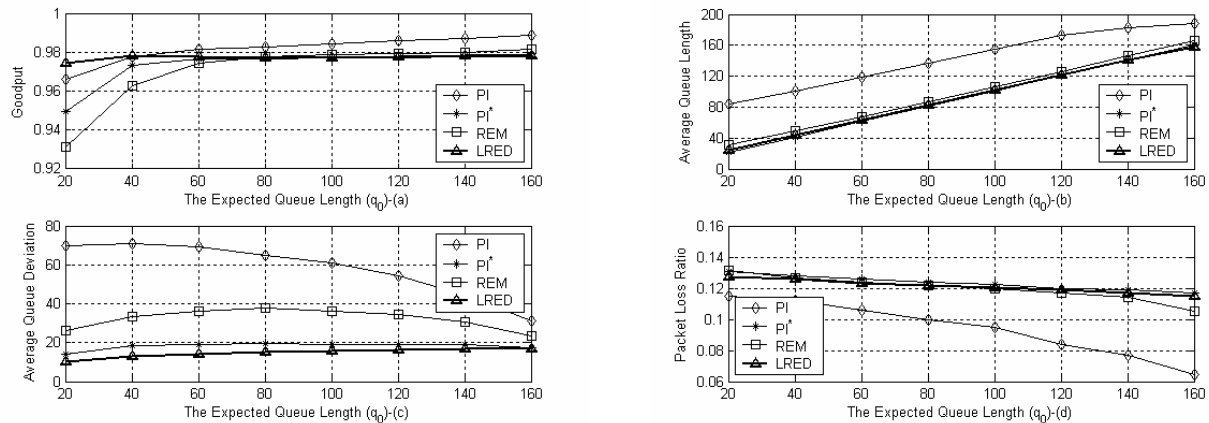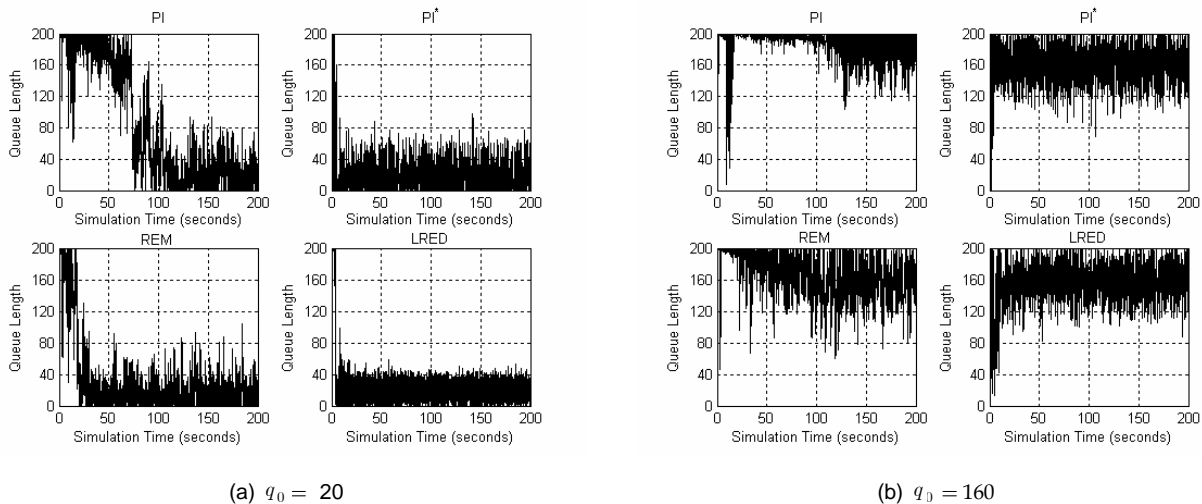Fig. 6. Experiment 2: Queue length under two extreme cases.
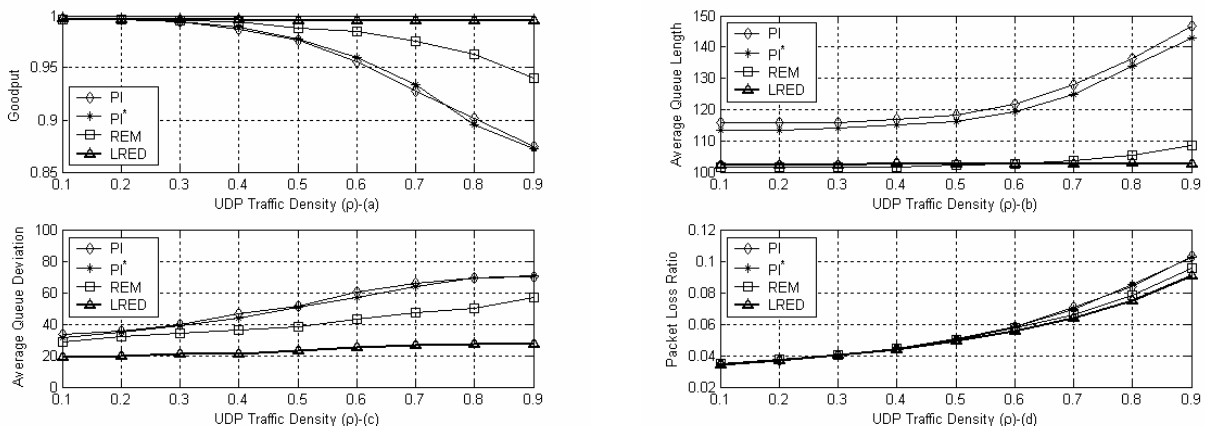

Fig. 7. Experiment 3: Comparisons of (a) goodput, (b) average queue length, (c) average queue deviation, and (d) packet loss ratio.

metrics: goodput, average queue length, average queue deviation, and packet loss ratio. The goodput is the overall throughput of the system excluding retransmissions; the average queue length is calculated as the arithmetic mean of the instantaneous queue lengths; and the average queue deviation is the average of the absolute deviations of the instantaneous queue lengths from the mean.

6.1 Homogenous Traffic: Long-lived TCP Flows Only

In this set of experiments, we focus on a network with homogeneous traffic, i.e., all are persistent TCP flows.

Experiment 1: Response time under various congestion degrees

We first investigate the response times for the AQM schemes to reach a steady state. The expected queue length q_0 is set to 100 packets, and the vector for link delays between Clients and Router 1, $(d_1, d_2, d_3, d_4, d_5)$, is (10, 50, 100, 150, 200). The total number of TCP flows, N, varies from 100 to 1000, which leads to 10 scenarios with various degrees of congestion. For PI*, we set N^- from 100 to 1000 for the 10 scenarios, R^+ is set to 450 ms for $N^- \leq 500$, and 500 ms to 900 ms for N^- from 600 to 1000, which guaran-

tee that $N^- < (R^+ * C)/2$ according to the design rule of PI in [12].

The response times as a function of the loss ratio are presented in Fig. 5. It can be seen that when the packet loss ratio increases (which is a result of an increase of the number of TCP flows), the response times of PI or REM increase, for both simulated and analytical results. The mismatches between the simulated and analytical results for PI and REM are relatively small, especially with high loss ratios. PI* however has a big mismatch between its simulated and analytical results. The main reason is that we assume the buffer is always full before reaching a steady-state when analyzing the lower bound of the response time in the Section 5. Since PI* uses the optimal parameter according to the design rules in [12], it has faster response than PI using the default parameters. Nonetheless, the response time of LRED is generally lower than other schemes, and it is nearly independent of the packet loss ratio. As a result, the gaps between LRED and other schemes under a high loss ratio are pretty significant.

Experiment 2: Stability under extreme conditions

In this experiment, we examine the stability of the AQM

(a) $q_0 = 20$ (b) $q_0 = 160$

Fig. 8. Experiment 3: Queue length for AQM schemes.

Fig. 9. Experiment 4: Comparisons of (a) goodput, (b) average queue length, (c) average queue deviation, and (d) packet loss ratio.

schemes under two extreme cases: 1) light congestion with a small number of TCP flows N and large round-trip time R, and 2) heavy congestion with a large N and small R. Fig. 6 presents the instantaneous queue lengths under such two cases: 1) In the first, we set N=80 and $d_1 = d_2 = d_3 = d_4 = d_5 = 250\ ms$. Therefore, we have N^-=80 and R^+=550 ms for PI*. 2) In the second, N=800 and $d_1 = d_2 = d_3 = d_4 = d_5 = 10\ ms$, and accordingly, we have N^-=800 and R^+=100 ms for PI*. Other parameters are the same as those in Experiment 1. We can see that LRED and PI* have similar response times, but LRED has less overshoots and smaller queue deviations. It can still be seen that, under heavy congestion, LRED achieves a shorter response time and better stability than PI and REM. Under light congestion, the queue length of REM drastically oscillates and almost out of control. On the contrary, LRED and PI can still stably regulate the queue, and LRED is relatively better. Note that the response time of LRED slightly increases in the light congestion case because it needs more rounds to have a stable loss ratio estimate. Nevertheless, its response time is still comparable to that to PI or PI*.

Experiment 3: Varying the expected queue length

This experiment is used to study the performance of

AQM schemes when the expected queue length varies. We fix the number of persistent TCP flows to N=400. The expected queue length q_0 varies from 20 to 40, 60, 80, 100, 120, 140, and 160. The other parameters are the same as that in the Experiment 1. Accordingly, for PI*, we set N^-=400 and R^+=450 ms. We can see from Fig. 7 that PI achieves higher throughput and lower loss ratio. The reason is that its slow response leads to longer queue length than the other three schemes (as shown in Fig. 8), which in turn reduces the loss ratio and increases the goodput. On the contrary, LRED, PI*, and REM effectively realize the expected average queue length and thus have almost the same packet loss ratio. Compared to PI* and REM, LRED has a higher goodput when q_0 is smaller than 60, as shown in Fig. 7 (a). Moreover, LRED has the smallest queue deviation. Fig. 8 presents the instantaneous queue lengths for the AQM schemes when q_0 equals 20 and 160. It can be seen that LRED achieves better tradeoff between the average queue length and other QoS performance measures, including goodput and loss ratio.

6.2 Non-Homogenous Traffic: Hybrid Flows

Experiment 4: Adding unresponsive UDP flows

(a) UDP traffic density (ρ) is 0.5   (b) UDP traffic density (ρ) is 0.9

Fig. 10. Experiment 4: Queue lengths for AQM schemes, where ρ is UDP traffic density.
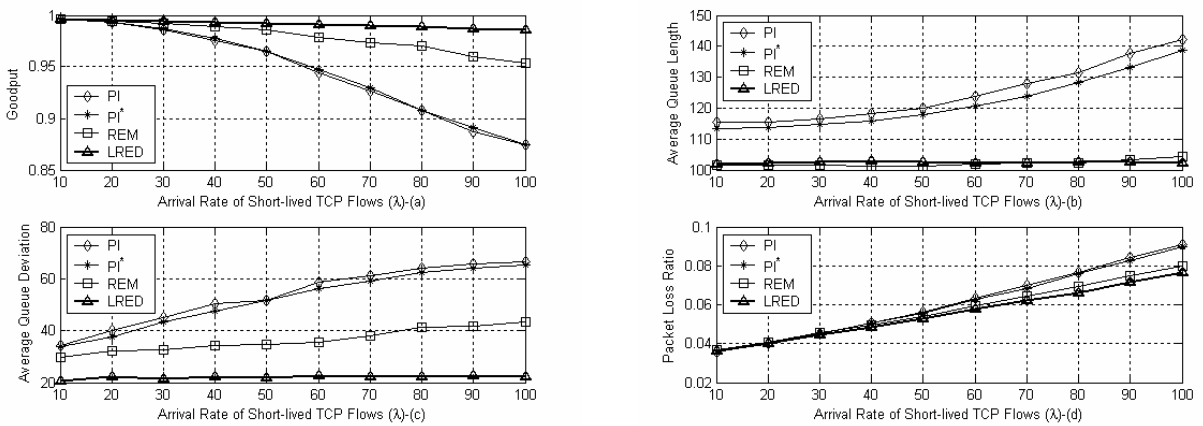
Fig. 11. Experiment 5: Comparisons of (a) goodput, (b) average queue lengths, (c) average queue deviations, and (d) packet loss ratios.

In this experiment, we investigate the performance of the AQM schemes with the existence of unresponsive UDP flows. In addition to the 100 persistent TCP flows, we introduce 100 UDP flows arriving in interval [50 sec, 150 sec]. Each is an ON/OFF flow, where the durations of the ON and OFF states are exponentially distributed with a mean of 1.0 second. The density of the UDP traffic over the total traffic, ρ, ranges from 0.1 to 0.9, and the rate of each UDP flow is $r = \rho \times 10 \text{ Mbps} / 100$ during the ON period. Other settings are the same as those in Experiment 1. Accordingly, we choose $N^- = 100$ and $R^+ = 450$ ms for PI*.

Fig. 9 presents results of the goodput, average queue length, average queue deviation, and packet loss ratio, as functions of the UDP traffic density. It can be seen that LRED generally outperforms PI, PI*, and REM in all these performance measures, especially when the UDP traffic density is high. Note that REM achieves better performance than PI in this experiment; however, it is stable with quite restricted network conditions only, as shown in Experiment 2.

We also present the instantaneous queue lengths for the AQM schemes in Fig. 10. There are two interesting observations. First, when ρ increases, LRED can still regulate the queue length to the expected value with much smaller un-

der- or over-shoots than PI, PI* and REM. Second, the buffers of PI, PI* and REM are overflowed (or empty) for a long time when the UDP flows start arriving from 50 seconds (or stops after 150 seconds), especially if ρ is large, i.e., 0.9. This is due to the slow responsiveness of PI, PI* and REM, which result in low goodput and high loss ratio. On the contrary, there is only a short-term increase (or decrease) at time 50 seconds (or 150 seconds), which implies that LRED has a much shorter response time.

Experiment 5: Adding short-lived TCP flows

Besides unresponsive UDP flows, short-lived TCP flows can also affect the performance of an AQM scheme. In this set of experiments, we introduce short-lived TCP flows, which arrive in intervals [50 sec, 150 sec] following to a Poisson process. The mean arrival rate λ varies from 10 flows/second to 100 flows/second, and the length of each short-lived TCP flow is uniformly distributed in [1.0 sec, 2.0 sec]. Other parameters are the same as those in Experiment 4 and $N^- = 100$ and $R^+ = 450$ ms are still set for PI*.

The average queue lengths, average absolute queue deviations, goodputs, and packet loss ratios for the three AQM schemes in this experiment are compared in Figs. 11. Clearly, LRED outperforms PI, PI*, and REM in all these

(a) λ =30 flows/second
(b) λ =100 flows/second
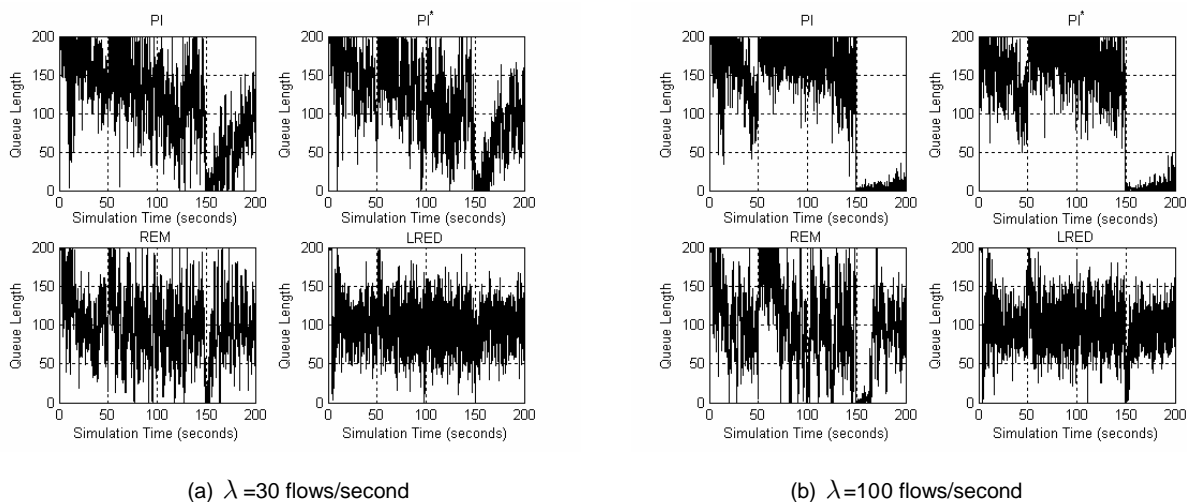
Fig. 12. Experiment 5: Queue lengths for the AQM schemes, where λ is the short-lived TCP arrival rate.
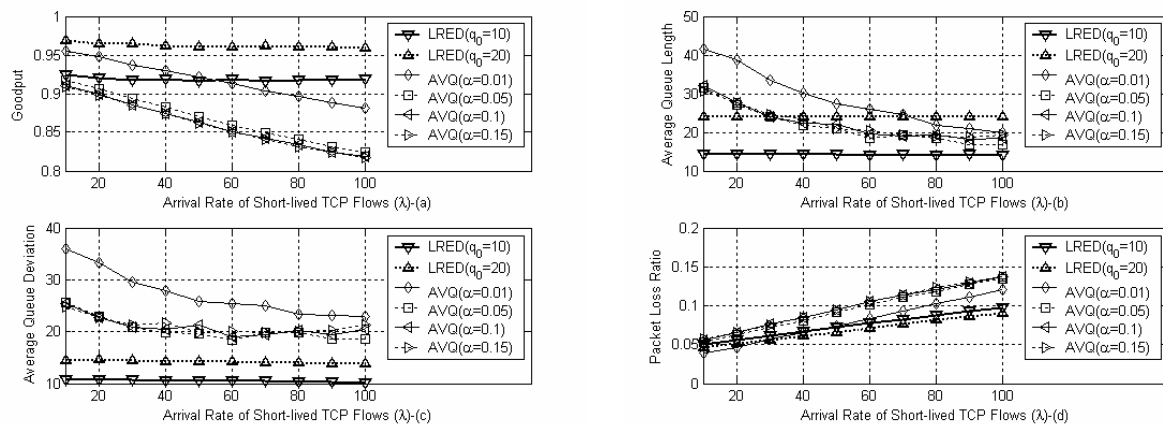
Fig. 13. Experiment 5: Comparisons between LRED and AVQ. (a) goodput, (b) average queue lengths, (c) average queue deviations, and (d) packet loss ratios.

measures. In addition, Fig. 12 shows the instantaneous queue length for PI, PI*, REM, and LRED for λ=30 and 100. LRED is again more stable due to its good responsiveness.

We also compare LRED with AVQ [10] in such a heterogeneous traffic environment. AVQ is known to be effective in regulating the queue length and achieving high link utilization [10]. In the design rules of AVQ in [10], $\alpha < 0.2$ is required to guarantee stability for this scenario ($N^- $=100, $R^+ = 450$ ms, and C=2500 packets/seconds). Hence, we let parameter α vary from 0.01 to 0.15 and set the expected utility of AVQ, γ, to 0.98. To make a fair comparison, we also set q_0 in LRED to small values (10 and 20), which match the average queue length in AVQ. The performance measures are presented in Fig. 13. When α decreases, AVQ achieves higher goodput and lower loss ratio but larger average queue length as well as queue deviation. When $\lambda \leq 60$, LRED with q_0=20 is better than AVQ. When $\lambda > 60$, LRED with q_0=10 outperforms AVQ. It implies that LRED achieves better performance than AVQ, if assigned with a small q_0.

According to [17] and [18], the Pareto distributions of file sizes and durations could contribute to the self-similar characteristics of the Internet traffic. Hence, we have also

conducted experiments with the packet inter-arrival time and flow's duration being set to Pareto distributions, with the same means in the previous experiment. The results are

Fig. 14. Experiment 5: Queue lengths for the AQM schemes when the short-lived TCP flows follow a Pareto process (λ =100 flows/second).

(a) The reverse direction (Short-lived TCPs)   (b) The forward direction (Long-lived TCPs)
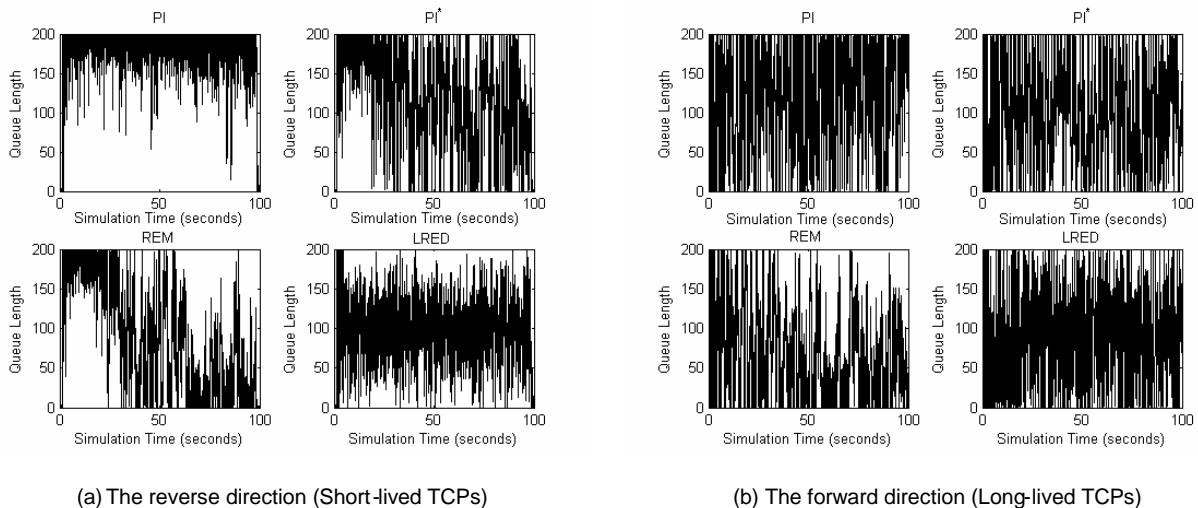
Fig. 15. Experiment 6: Queue lengths for the AQM schemes (q_0 =100, Q =200).

(a) The reverse direction (Short-lived TCPs)   (b) The forward direction (Long-lived TCPs)
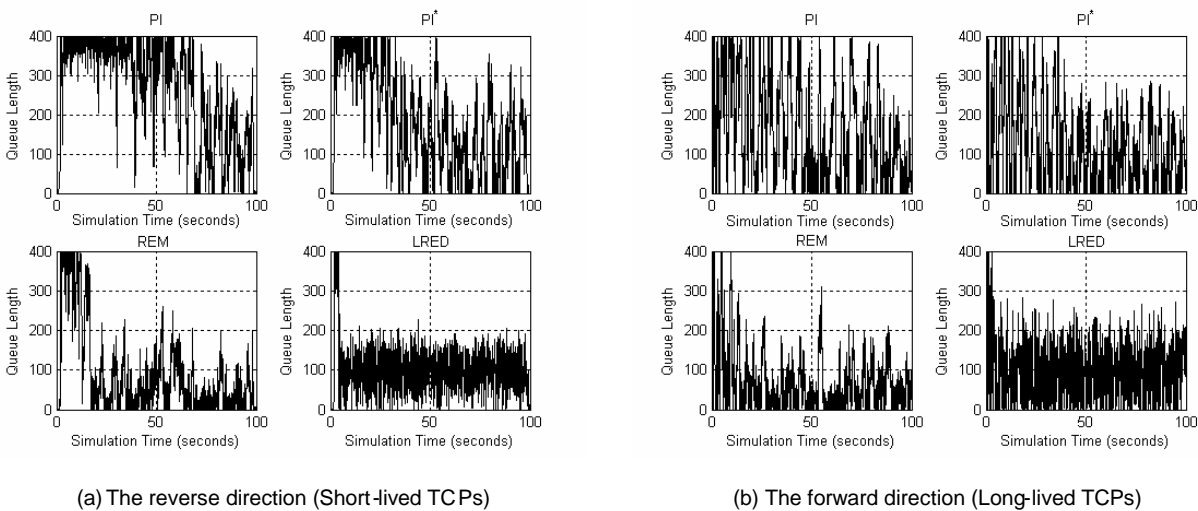
Fig. 16. Experiment 6: Queue lengths for the AQM schemes (q_0 =100, Q =400).

presented in Fig. 14, for PI, PI*, REM, and LRED, respectively. It shows that LRED still has faster response and better performance, and the difference between the results of Poisson and Pareto distributions is generally insignificant. However, it is worth noting that the Pareto distribution is not necessary the best Internet traffic model, particularly considering that the Internet traffic has always been changing with such emerging new technologies and applications as peer-to-peer communications. We expect that, in our future study, more results can be obtained using up-to-date Internet traces or advanced traffic models, e.g., the Fractional Gaussian Noise (FGN) distribution [19].

6.3 Two-Way Traffic: Forward and Reverse Direction

Experiment 6: Two-way traffic and two-way congestion

In this experiment, we introduce two-way traffic: 1) In the reverse direction, namely, from the clients to the servers (see Fig. 4), only short-lived TCP flows with a Poisson arrival process of a 200 flows/second arrival rate are configured. The other parameters related to short-lived TCP flows are the same as that in the Experiment 5; 2) In the forward

direction, or from the servers to the clients, there are 300 persistent TCP flows. We set N^- =200 and R^+ =600 ms for PI*. Note that congestion will occur in both directions of the link connecting Routers 1 and 2 in Fig. 4. We set the simulation time to 100 seconds, which enables 20000 short-lived TCP flows, and is long enough to discover the performance difference among the AQM schemes (see Figs. 15 and 16).

We collect the queue length for the AQM schemes in both directions of the congested link. In Fig. 15 (q_0 =100 and Q =200), we can see that, all the AQM schemes have noticeable overshoots in this case with bi-directional congestion on the same link; yet LRED controls the queue length better than others. The forward traffic that consists of persistent TCP flows is influenced more noticeably than the short-lived TCP flows in the reverse direction (Fig. 15 (b)). To avoid buffer overflow (see in Fig. 15(b)), we increase the buffer size from Q =200 to Q =400. The results are presented in Fig. 16, where both PI and PI* still show slower response and bigger overshoot. Although REM responds more quickly, its queue length is often below the

expected value (100), leading to lower goodput. On the contrary, LRED still regulates queue length around the expected value (100) and achieves better control effect than other schemes.

6.4 Summary

Our simulation results suggest that LRED achieves fast response even under heavy congestion and its performance is quite good in terms of goodput, queue length and queue deviation, and packet loss ratio. PI and REM however have slower response, which leads to performance degradation under dynamic network environments. Specifically, their packet dropping probability is iteratively computed (see Eqs. (32) and (35)); if p_0 is high (e.g., when the RTT R is low and the number of TCP flows N is large), PI and REM need a long time for the drop probability p to converge to p_0. The convergence rate and response time of LRED are almost independent of p_0, as shown in Fig. 5, but mainly influenced by the measurement period. More importantly, when the network is highly dynamic, LRED can quickly converge to a new stable state through its multi-granular update.

LRED relies on the measured packet loss ratio and therefore the parameters involved in measurement should be carefully configured. Specifically, the measurement weight (w_m) in Eq. (2) should be set to a small value in order to capture the latest packet loss. Regarding measurement period (t_m), if it is too small, the measurement could be inaccurate; but if it is too big, the response time can be longer. Our experience shows that w_m =0.1 and t_m =1.0 are reasonable choices in most scenarios.

Another parameter in LRED is β, whose guideline is given by Theorem 2. Note that, if β is too big, the packet drop probability calculated by Eq. (3) will be either bigger than 1 (when $q > q_0$) or smaller than 0 (when $q < q_0$). In this case, LRED behaves like a virtual Tail-Drop with a virtual buffer size of q_0. Fig. 17 presents the simulated results for such a scenario, where β =10. It can be seen that the packet drop probability almost equals 1 or 0, and the queue length stays below 100, like in a traditional Tail-Drop buffer.

7 CONCLUSIONS AND FUTURE WORKS

In this paper, we have proposed a novel AQM algorithm, LRED, which incorporates packet loss ratio as a complement to queue length for congestion estimation. In LRED, the packet drop probability is updated over multiple grains: on a fine grain, LRED uses the instantaneous queue length mismatch to update the drop probability upon each packet arrival; on a coarse grain, LRED adjusts the drop probability according to the packet loss ratio, which has never been considered in existing AQM algorithms. We have developed an analytical model for LRED, which suggests that this multi-granular update improves not only the stability of the system, but also its responsiveness. Such observations have been validated by our simulation results under various configurations. We have also compared LRED with existing AQM algorithms, including PI, REM, and AVQ. Our results have showed that LRED remains quite stable when the number of TCP flows and round-trip
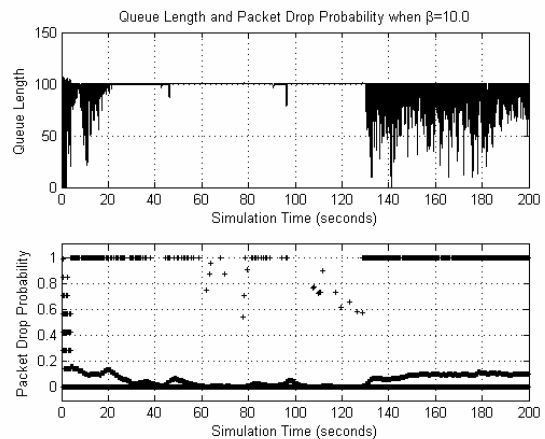
Fig. 17. LRED dynamics with a large β (N =400, q_0 =100).

times vary significantly, or when many short-lived flows or unresponsive UDP flows exist in the network. Moreover, it can effectively control the queue to the expected length, and achieves a better tradeoff between the goodput and queue length. Finally, LRED achieves reasonably good performance under two-way traffic.

There are many possible future works for enhancing the LRED algorithm. We are particularly interested in extending LRED to support differentiated QoS, where packets with different priority may have non-uniform dropping probabilities. Meanwhile a quantitative analysis of LRED's response time is very important to more precisely evaluate LRED's performance. The effect of substituting packet dropping with packet marking is also worth investigating. Another challenging work is to model LRED's performance for short-lived TCP flows, which will complement our existing analytical results with long-lived flows. Finally, we are interested in examining the performance of LRED under more realistic Internet traffic traces, or more sophisticated traffic models that reflects the recent Internet development, e.g., Fractional Gaussian Noise (FGN) distribution [19].

REFERENCES

[1] R. Gurin and V. Peris, "Quality-of-Service in Packet Networks: Basic Mechanisms and Directions," *Computer Networks*, vol. 31, no. 3, pp. 169-179, Feb. 1999.

[2] Y.-T. Hou, D. Wu, B. Li, T. Hamada, I. Ahmad, and H. J. Chao, "A Differentiated Services Architecture for Multimedia Streaming in Next Generation Internet," *Computer Networks*, vol. 32, no. 2, pp. 185-209, Feb. 2000.

sor Networks (2nd Quarter of 2005). In addition, He has been involved in organizing over 40 conferences, esp. IEEE Infocom since 1996. He was the Co-TPC Chair for IEEE Infocom 2004.

Kazem Sohraby received B.S. (highest distinction), M.S., and PhD, all in Electrical Engineering, took graduate courses in computer science at the George Washington University, and received an MBA from the Wharton School, University of Pennsylvania. He is a Professor and head of department of computer science and computer engineering at University of Arkansas at Fayetteville. He also serves as a Principal Consultant on contracts with the Defense Information Systems Agency (DISA). Before coming to Arkansas, he served as Chair of interdisciplinary Telecommunications Management Department at Stevens Institute of Technology (1997-2002), and as adjunct Professor of Electrical and Computer Engineering, and Computer Science (1987-2002). While with Bell Labs, Lucent Technologies and AT&T Bell Labs (1983-2003), he was a Distinguished Member of Technical Staff, with the Performance Analysis Department in the Advanced Communications Technologies Center, with the Mathematics of Networks and Systems Department in the Mathematical Sciences Research Center, and forward looking center on switch design. His areas of interest include computer networking, signaling, switching, performance analysis, and traffic theory. Prior to joining Bell Labs, Dr. Sohraby served as the Head of Network Design and Planning Department at Computer Sciences Corporation, Systems Division (1978-1983). He has total of over 20 patent applications (14 granted) on computer protocols, wireless and optical systems, circuit and packet switching, and Optical Internet. He has several publications, including a book on the performance and control of computer communications networks. He is a Distinguished Lecturer of the IEEE Communications Society, and Director of IEEE Communications Society (on-line content), and served as its president's representative on the Committee on Communications and Information Policy (CCIP). He served as chair of several conferences in both IEEE and ACM, ACM Mobicom special interest group, and is vice chair of the IEEE Infocom executive committee. He also served on the education committee of the IEEE Communications Society, is on the editorial boards of several publications, and panelist and reviewer with the National Science Foundation. He served as grant referee and reviewer with the US Army, and the Natural Sciences and Engineering Research Council of Canada.

Y. Thomas Hou (S'91-M'98-SM'04) obtained his B.E. degree from the City College of New York in 1991, the M.S. degree from Columbia University in 1993, and the Ph.D. degree from Polytechnic University, Brooklyn, New York, in 1998, all in Electrical Engineering. From 1997 to 2002, He was a researcher at Fujitsu Laboratories of America, IP Networking Research Department, Sunnyvale, California. Since fall 2002, he has been an Assistant Professor at Virginia Tech, the Bradley Department of Electrical and Computer Engineering, Blacksburg, Virginia. His current research interests include resource (spectrum) management and networking issues for SDR-enabled wireless networks, optimization and algorithm design for wireless ad hoc and sensor networks, and video communications over dynamic ad hoc networks. In recent past, he also had work on scalable architectures, protocols, and implementations for differentiated services Internet, service overlay networking, video streaming, and network bandwidth allocation policies and distributed flow control algorithms. He has published over 100 journal and conference papers in the above areas and is a recipient of the 2004 IEEE Communications Society Multimedia Communications Best Paper Award, the 2002 IEEE International Conference on Network Protocols (ICNP) Best Paper Award, and the 2001 IEEE Transactions on Circuits and Systems for Video Technology (CSVT) Best Paper Award. He is a member of ACM and a senior member of IEEE.