Drop Behavior of RED for Bursty and Smooth Traffic

Thomas Bonald and Martin May
INRIA Sophia-Antipolis, France
{tbonald,mmay}@sophia.inria.fr

February 1999

Abstract

We prove analytically that, as claimed in Jacobson and Floyd's initial paper on RED [2], RED gateways "avoid the bias against bursty traffic", meaning that bursty traffic suffers more losses than smooth traffic with Tail Drop (TD) gateways, but both types of traffic suffer equally with RED gateways. We also show, again as claimed by Jacobson and Floyd, that RED gateways mark or drop packets from a connection at a rate proportional to that connection's arrival rate. Interestingly, both properties above *only* hold when the bursty traffic may be modeled by Poisson arrivals of bursts; we exhibit non Poisson cases where these properties do not hold at all. We conclude the paper with a discussion on the implication of these results on traffic management in the Internet.

1 Introduction

The Random Early Detection (RED) buffer management scheme is expected to provide interesting performance improvements for congestion avoidance protocols like TCP. As a result, the end-to-end research task force has recently advocated the widespread deployment of RED in the Internet. While some of the advantages of RED are clear, for example the control of the average queue length, others have only been illustrated with simulations, for example the claim that RED avoids the loss bias against bursty traffic. Of course, there is nothing wrong with simulation. However, it is difficult with simulation to determine exactly under which conditions a property holds, and to quantify the impact of various parameters on the result. Hence the analytic approach taken in this paper. Specifically, we examine, and prove correct, a property claimed by Jacobson and Floyd in the initial RED paper [2], namely RED gateways i) "avoid the bias against bursty traffic" and ii) mark or drop packets from a connection at a rate proportional to that connection's arrival rate. But we'll show that these two properties *only* hold when the bursty traffic may be modeled by Poisson arrivals of bursts (they require PASTA: Poisson Arrivals See Time Averages).

2 RED and bursty traffic

Consider a gateway with a buffer of K packets. We define the *drop strategy* of the gateway as an increasing function $\alpha : \{0, \ldots, K\} \to [0, 1]$, with $\alpha(0) = 0$ and $\alpha(K) = 1$. Specifically, $\alpha(k)$ is the drop probability of an arriving packet when k packets are buffered in the queue. The Tail Drop (TD) strategy consists in choosing $\alpha(k) = 0$ for any $k \neq K$. In the numerical examples below, we take $K = 40$ packets and define the Random Early Detection (RED) strategy as $\alpha(k) = 0$ for $k \leq K_{th}$ and

$$\alpha(k) = \frac{k - K_{th}}{K - K_{th}} \quad k > K_{th},$$

where K_{th} is a threshold typically set to $K/2$.

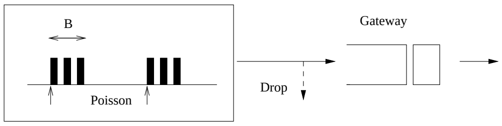

Figure 1: Traffic model

Assume that packets arrive in bursts of B packets, according to a Poisson process of rate λ. The processing times of the packets in the gateway are assumed to be exponentially distributed with mean μ^{-1}. We define the *offered load* by $\rho = B\lambda/\mu$. The number of packets buffered in the queue defines a Markov chain, the stationary distribution of which can be easily computed. Denoting by π this stationary distribution, we obtain using the PASTA property the drop probability of a packet in a TD gateway:

$$P_{TD} = \pi(K) + \pi(K-1)\frac{B-1}{B} + \ldots + \pi(K-B+1)\frac{1}{B}.$$

Concerning RED, we use the following approximation:

$$P_{RED} \sim \pi(K) + \pi(K-1)\alpha(K-1) + \ldots + \pi(1)\alpha(1).$$

Note that the stationary distribution π is different from that obtained with TD. It turns out that this approximation (which is actually a lower bound) is very accurate, as illustrated in Figure 2.

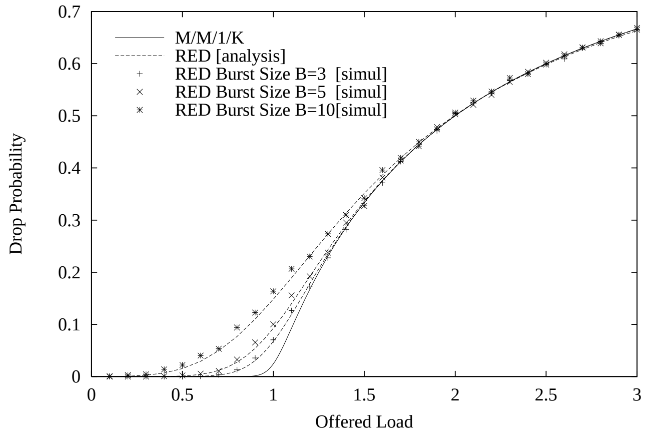

Figure 2: Drop probability vs. offered load for bursty traffic

In addition, when the offered load is high, that is during a period of congestion, the drop probability is very closed to that suffered by a (smooth) Poisson traffic in a TD gateway, namely

$$P_{M/M/1/K} = 1 - \frac{1-\rho^K}{1-\rho^{K+1}} \sim 1 - \frac{1}{\rho}.$$

Since $P_{M/M/1/K} \leq P_{TD} \leq P_{RED}$, the drop probability is then similar with RED and TD, whatever the burst size B. In fact, the behavior of RED and TD differ essentially when both smooth and bursty traffic share the same gateway, as we shall see now.

3 RED and mixed bursty/smooth traffic

Consider two flows sharing the same gateway, one bursty with Poisson arrivals of bursts (we take $B = 3$ in the examples below) and another smooth with Poisson arrivals. Denote by $\rho(b)$ and $\rho(s)$ the load of the bursty and the smooth traffic, and $\rho = \rho(b) + \rho(s)$ the total offered load.

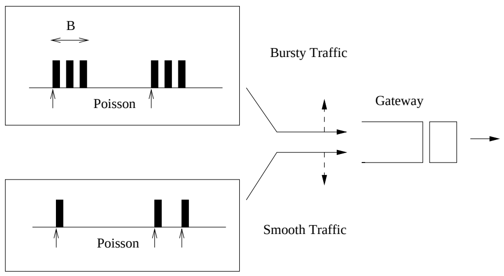

Figure 3: Model for the traffic mix

Let π be the stationary distribution of the number of packets buffered in the queue. Using the PASTA property, we obtain the drop probability of a packet for the bursty flow and the smooth flow in a TD gateway:

$$P_{TD}(b) = \pi(K) + \pi(K-1)\frac{B-1}{B} + \ldots + \pi(K-B+1)\frac{1}{B},$$

and

$$P_{TD}(s) = \pi(K).$$

On the other hand, using PASTA and the same approximation as in §2, we get for a RED gateway,

$$P_{TD}(b) \sim \sum_{k=1}^{K} \pi(k)\alpha(k) = P_{TD}(s).$$

Thus the drop probability is the same for bursty and smooth traffic in RED gateways, whereas it is is much higher for bursty traffic than for smooth traffic in TD gateways (see Figure 4 below). Note, however, that RED decreases significantly the drop probability for bursty traffic *only* when the fraction of smooth traffic is high. Otherwise, RED behaves exactly as TD with respect to the bursty traffic, but increases significantly the loss rate suffered by smooth traffic.

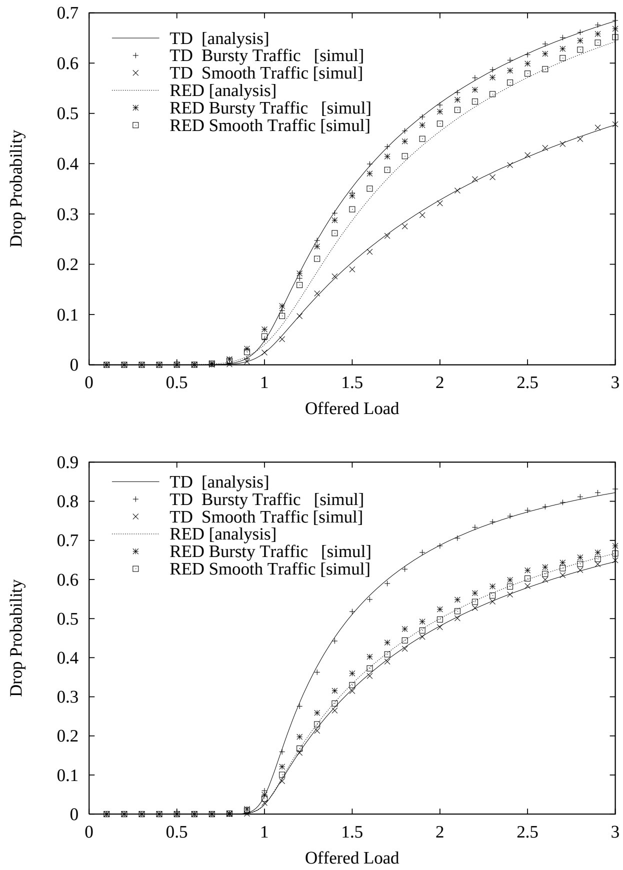

Figure 4: Drop probability vs. offered load for low (10%) and high (90%) fraction of smooth traffic

Figure 5 below shows the drop probability obtained for an offered load of $\rho = 2$ with respect to the fraction of smooth traffic. As expected in view of previous results, the drop probability does *not* depend on the fraction of smooth traffic in RED gateways, as opposed to TD gateways. Hence the properties i) and ii).

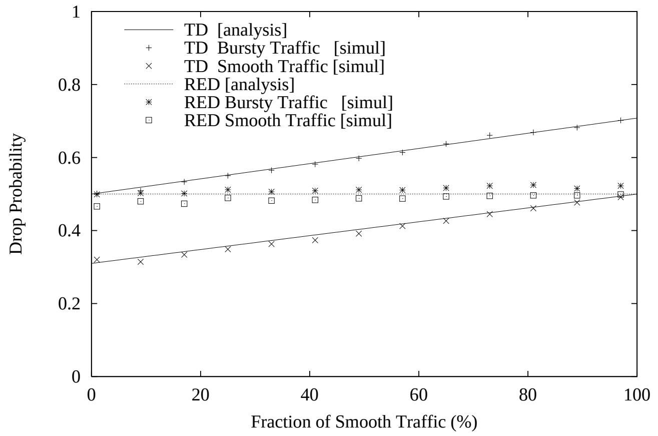

Figure 5: Drop probability vs. fraction of smooth traffic for an offered load of $\rho = 2$

4 The Poisson hypothesis

Previous analysis depends in a crucial way on the Poisson hypothesis, because proving the two properties above requires the PASTA property. In general, it is indeed not true that the stationary distribution of the number of packets k buffered in the queue immediately *before* the arrival of a packet or a burst of packets (that is under the Palm probability [1]) coincides with π, the continuous–time stationary distribution of k.

Figure 6 shows the drop probabilities obtained when a bursty traffic with Pareto inter arrival times of bursts is mixed with a Poisson (smooth) traffic. The Pareto coefficient is 1.4, the fraction of smooth traffic is 10%.

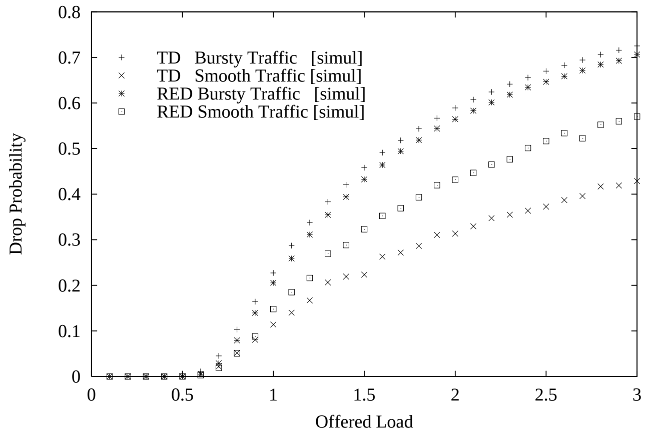

Figure 6: Drop probability vs. offered load for non Poisson arrivals of bursts

Clearly, the drop probability for bursty traffic is different from the drop probability for smooth traffic even in the presence of a RED gateway.

5 Discussion

We conclude this paper by discussing some implications of the upper results. In particular:

- Does it help to not bias against bursty traffic ?

  When a buffer management scheme avoids the bias against bursty traffic as RED does, it will encourage the use of bursty traffic. Clearly this is not a desirable goal of a network provider. ISPs are more interested in a smooth and predictable traffic.

  In addition our results show that RED is efficient only if a high fraction of the traffic is smooth. In the current Internet where most of the traffic is generated by TCP, this is probably not the case. Our results then show that while bursty traffic will not experiment any significant improvement, smooth CBR/UDP traffic like voice, even rate controlled and TCP-friendly, will suffer if RED gateways are deployed.

- The fact that packets from a connection are dropped by a RED gateway at a rate proportional to that connection's arrival rate is interesting for modeling purposes. For example one key question with linear increase/multiplicative decrease algorithms such as those in TCP is how their efficiency and fairness properties will be affected when deploying RED in place of TD. One can answer this question using models such as that in Jain and Chiu [3] or Hurley and LeBoudec [4]; the key there is to determine the rate at which each source receives feedback information. Our result shows that this rate is proportional to the source's sending rate only for Poisson arrivals.

In general, we believe it is important to quantify the performance of RED gateways and the impact of RED on rate-controlled connections. Much of the work done so far has relied on simulation and experimental results. However, a full quantitative understanding of RED is still missing. Our contribution is a first step in that direction.

References

[1] F. Baccelli and P. Brémaud, *Elements of Queueing Theory*, Springer-Verlag, 1994.

[2] S. Floyd and V. Jacobson, Random Early Detection gateways for congestion avoidance, IEEE/ACM Trans. on Networking 1 (1993) 397–413.

[3] D. Chiu and R. Jain, Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks, Computer Networks 17 (1989) 1–14.

[4] P. Hurley and J.Y. Le Boudec, A Note on the Fairness of Additive Increase and Multiplicative Decrease, *Workshop on TCP*, Paris, 1998.