



Simulation analysis of RED with short lived TCP connections

Eitan Altman ^{a,*}, Tania Jiménez ^b

^a INRIA, BP93, 2004 Route des Lucioles, 06902 Sophia Antipolis Cedex, France

^b CESIMO, Facultad de Ingeniera, Univ. de Los Andes, 5101 Mérida, Venezuela

Received 5 October 2003; received in revised form 6 October 2003; accepted 8 October 2003

Responsible Editor: S. Low

Abstract

Several objectives have been identified in developing the random early drop (RED): decreasing queueing delay, increasing throughput, and increasing fairness between short and long lived connections. It has been believed that indeed the drop probability of a packet in RED does not depend on the size of the file to which it belongs. In this paper we study the fairness properties of RED where fairness is taken with respect to the size of the transferred file. We focus on short lived TCP sessions. Our findings are that (i) in terms of loss probabilities, RED is unfair: it favors short sessions, (ii) RED is fairer in terms of the average throughput of a session (as a function of its size) than in terms of loss probabilities. We study various loading regimes, with various versions of RED.

© 2003 Elsevier B.V. All rights reserved.

Keywords: RED; TCP; Buffer management; Simulations; Fairness

1. Introduction

One of the objectives of random early drop (RED) has been to get rid of the bias of drop-tail buffers against bursty traffic [3]. It has been confirmed in [4, Section IX] that indeed RED gets rid of this bias and is fair in terms of fraction of lost packets. This has been demonstrated in a framework where both bursty as well as “smooth” traffic were taken as permanent FTP connections. The higher burstiness was obtained by using smaller windows and longer round-trip times. In a recent paper [7], the authors show that conclusions drawn

from simulating permanent TCP connections can be qualitatively quite different than those obtained from simulations of transfers that have large time scale variability. The latter is obtained by replacing the infinite source model by ones in which transferred files have heavy-tailed distributions. This motivated us to question the qualitative conclusions drawn in [4] and restudy the fairness issue using Pareto distributed file sizes. We are interested in the following questions: (i) how does the drop probability of a packet depend on the file size? (ii) how does the average throughput experienced by a file depend on its size?

If we identify arrival of packets from a file as a “batch” (which is justified by our simulations), then the fairness in terms of the file size also answer the question of fairness in terms of batch sizes.

* Corresponding author. Tel.: +33-4-92-38-77-86.

E-mail address: eitan.altman@sophia.inria.fr (E. Altman).

Our main findings are that in terms of loss probabilities, RED is biased against long files. Furthermore, we show that this bias increases as the workload decreases. We study this phenomenon and provide an explanation for it. We further show that RED is fairer in terms of throughput than it is in terms of loss probabilities. When comparing to drop-tail buffers, we see that for all loads, RED has smaller loss probabilities for almost all transfer sizes; the drop-tail buffer is more fair in terms of drop probabilities (as a function of the transfer size) since it has many more drops for short file transfers. In terms of throughput, RED is slightly fairer than the drop-tail buffer. Our conclusions similar for four different variants of RED that we tested.

We briefly mention some related work. Fair RED (FRED) is proposed in [10] which needs however to keep some per-active-flow states.¹ Its performance is analyzed by simulating permanent TCP connections. In [14], stabilized RED (SRED) is proposed, requiring per-flow state information too. Unlike RED, the drop probabilities depend only on the instantaneous buffer occupation and on the estimated number of active flows. Fairness according to the transfer size is not considered there. [2] analyzes biasness with respect to bursty traffic. Both smooth as well as bursty traffic are modeled as Poisson processes, but in the smooth case each arrival corresponds to a single packet, where as in the bursty traffic case each arrival brings a batch of B packets. The conclusions of the paper are that indeed RED decreases the bias against bursty traffic. The traffic models used are not closed loop (they do not adapt to congestion), and as we learn from [1] and [7, Section 4], qualitative behavior of closed loop traffic can be quite different from open loop. Although there are also some TCP simulations of RED in [2], these use permanent connections. Other papers studied RED using long lived TCP connections [6,9,11–13,15–17] as well as short lived connections [8,17]. But the fairness issue is not examined in these references.

¹ It is argued in [10] that dropping “fairly” packets does not guarantee fair bandwidth sharing. Interestingly, our simulations show that bandwidth sharing is fairer in RED than could be expected because RED is unfair in dropping packets.

We study the standard as well as the adaptive RED version, both with and without the gentle option of RED.² We introduce the model and simulation setup in Section 2, present some preliminary simulation results in Section 3, then present the fairness results in terms of the loss probabilities in Section 4 and in terms of the throughput in Section 5. The analysis of these results and explanation of the causes for the observed behavior are given in Section 6, and we end with a concluding section.

2. Model and simulation setup

2.1. Traffic model and topology

There are X source nodes (X is determined below) connected to a bottleneck link N through which the packets are forwarded to a common destination D . We consider transfer of files from each source node whose distribution is Pareto with an average size of 30 kbytes and with shape parameter of 1.1.³ From each node, the time between beginning of transmissions of files has an exponential distribution with average of 0.9 s. Thus a source can be sending simultaneously packets belonging to more than one connection. The bottleneck link $N - D$ has a delay of 1 ms and a queue of size 50 packets. The rate of the link is 1.8 Mbps (it will later be increased up to 2.5 Mbps in order to study the dependence of the performance in the load of the system). The other links, i.e. between the sources and N , have all 100 Mbps bandwidth and a delay of 1 ms. The network is depicted in Fig. 1. The average input rate of data information into the bottleneck link from each node is $30 \times 10^3 \times 8 / 0.9 \text{ s} = 266.67 \text{ kbps}$. We took packet sizes of 500 bytes, so with an additional overhead of 40 bytes, their size is 540 bytes. Taking this into account, the total transmission rate of bits (without counting the retransmissions) from each source is 288 kbps. Thus the system can

² <http://www.icir.org/floyd/red/gentle.html>.

³ Recall that for Pareto distribution, $\Pr(\text{size} > s) = (k/s)^\beta$ and $E[\text{size}] = \beta k / (\beta - 1)$ where β is the shape parameter; thus the parameter k equals to 2727.27.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات