International Journal of Computing and Business Research (IJCBR)

ISSN (Online) : 2229-6166

Volume 3 Issue 2 May 2012

ACTIVE QUEUE MANAGEMENT FOR FORWARD ERROR CORRECTION

Sushil Kumar Sharma
Pursuing Ph.D in Computer Science
E-mail: shilu_online@rediffmail.com

Abstract

This paper studies the interaction of a forward error correction (FEC) code with queue management schemes like Drop Tail (DT) and RED. Since RED spreads randomly packet drops, it reduces consecutive losses. This property makes RED compatible with the use of FEC at the packet level. The FEC combined with RED may indeed be more efficient than FEC combined with DT. This however depends on several parameters like the burstiness of the background traffic, the FEC block size and the amount of redundancy in a FEC block. We conclude generally that using FEC is more efficient with RED than with DT when the loss rate is small, a relatively important amount of redundancy and at most a moderate FEC block size is used. This paper complement these observations with a simple model, which is able to capture the tradeoff between the locality and the frequency of losses.

Introduction

The Internet traffic suffers from heavy losses due to network congestion caused by the limited capacity of queue in the routers. There exist two end-to-end error control techniques to repair these losses: ARQ (Automatic Repeat reQuest) which consist in retransmitting dropped packets upon the destination's request, and FEC (Forward Error Correction) which consists in sending redundant packets to the destination, allowing it to repair losses without requiring packet retransmission. Because of retransmissions, ARQ is not appropriate for real-time applications. FEC is also used in bulk data transfer applications such as Digital Fountain. Several drawbacks are attached to the use of FEC. First, FEC cannot recover all lost packets. In addition, the transmission of redundant packets increases the overall network load. Finally, the electiveness of FEC is known to depend on the way packet drops are distributed in the data stream. FEC is more efficient when packets losses are independent, and much less when they occur in groups. In conjunction to end-to-end error control techniques, there exist queue management schemes operating inside routers that control network congestion. The “queue management” scheme traditionally used in the current Internet is Drop Tail (DT), which consists in discarding arriving packets when the buffer of the router overflows. Active queue management schemes, in particular the Random Early Detection (RED) scheme have been recommended recently by the IETF as an alternative, aimed at eliminating deficiencies of DropTail. The RED scheme basically discards packets earlier so that incipient stages of congestion can be detected. The aim of this paper is to study the interaction of FEC with RED (RED/FEC) and to compare the obtained results with those obtained from the combination of FEC with Drop Tail (DT/FEC).

2.0 FEC and Queue Management

Properties of codes used for Forward Error Correction has k data packets bearing the relevant information, the encoder (based for instance on a Reed-Solomon Erasure code generates h redundant packets useful for the recovery of the lost data packets. The concatenation of the k data packets and the h redundancy packets is called a FEC block of size n = k+h. If the total number of lost (data and redundant) packets is at most h, the decoder at the destination can retrieve successfully all lost packets. As a result all the relevant information is saved. Otherwise, if the total loss exceeds h packets, it is impossible to recover the lost packets. The queue management schemes studied here are Drop Tail, the principle of which is straightforward, and RED. With RED, the router maintains an estimate of the average queue length, using an exponential moving average. Based on this value, it accepts or rejects incoming packets with a certain probability. The rejection probability function is a parameter of the mechanism.

3.0 Experimental Setup

3.1 Network topology

The topology incorporated in it may have traffic generated by nodes S0 to SN is multiplexed on a 10 Mbps bottleneck link between nodes R1 and R2 with a propagation delay of 30 ms. The bottleneck link is provided with a DropTail or a RED queue of limited capacity of 35 packets. The other links located between nodes S0,..., SN and node R1 have a capacity of 100Mbps. These links have different propagation delays uniformly distributed from 20ms to 100ms. We have studied a traffic mix of TCP and UDP flows, with UDP representing a minority of the traffic. Our purpose in doing so is that we wish to study the behavior of the FEC technique under bursty conditions. A background traffic generated by TCP sources is appropriate in that respect since TCP traffic is usually quite bursty. Alternately, a simulation with variable-rate UDP sources could have been used.

3.2 Performance metrics

The metrics used for a specific flow are:

- The packet loss rate before correction (PLRBC) that is the ratio of the average number of lost packets in a FEC block before correction to the size of the FEC block.
- The packet loss rate after correction (PLR) that is the ratio of the average number of lost packets in a FEC block after correction to the size of the FEC block.
- The loss run length that is the number of packets of a particular flow that are lost consecutively. This is a random variable which gives an insight into the packet loss process. Studying this metric allows to investigate which queue management is more efficient when used with FEC.

4.0 Performance Measures

4.1 Influence of the number of TCP flows

The evolution of the packet loss rate before correction (PLRBC) and after correction (PLR) for the UDP source as a function of the number of TCP flows and the number of redundancy h. When the amount of redundancy is increased, the PLR of DT/FEC decreases for this network configuration, i.e. for a configuration where the number of sources implementing FEC is small. This observation is maintained for the PLR of RED/FEC in this case. The observation is that the PLRBC for RED is greater than the PLRBC for Drop Tail since RED starts dropping packets earlier without reaching the buffer capacity of the queue. Nevertheless, after the correction of lost packets, RED/FEC out-performs DT/FEC and gives a lower PLR for a small number of TCP flows. For a larger number of flows, the situation is reversed: RED yields a worst performance. Indeed, for one packet of redundancy (h = 1) and k = 16 data packets, i.e. for an addition of 6% of load, Drop Tail can divide the PLRBC by about 1.9 for 10 TCP flows. RED does better by dividing the PLRBC by about 3.4. In the case of a 25% load increase, for h = 4 and k = 16, Drop Tail can divide the PLRBC by about 33.4 for 10 TCP flows. In this case, the correction rate of RED is more significant since it is able to divide the PLRBC by 66.5 for 10 TCP flows and by 79.6 for 30 TCP flows. These results show that the number of TCP flows under which RED experiences an improvement on the PLR as compared to Drop Tail depends on the amount of redundancy. Hence, this threshold number increases as the amount of redundancy increases. For instance, this number is: 45 flows for h = 1 and 80 flows for h = 4. But the increase becomes slower as h grows. We have therefore shown that even if RED increases the UDP packet loss rate (which is in accordance with previous studies, it becomes possible to reduce its PLR by using FEC and to obtain a PLR less than the PLR for Drop Tail under certain conditions (precisely for a certain number of TCP flows and a certain amount of redundancy)

4.2 Loss run length

It is also studied that the distribution of the loss run length, for both queue management mechanisms and in function of the cross traffic. The results fully reported in showed that for k = 16 and h = 1, under the RED scheme, the probability to have a loss run length of size 1 packet is much larger than under Drop Tail (about 90% against 60%, respectively). This holds whatever the number of flows: the distribution does not appear to depend much on the cross traffic. This last observation shows that the situation is not as simple as initially thought. The starting assumption was: if RED shows a smaller loss run length, then FEC will perform better with RED than with DropTail concerning the capacity of repairing lost packets. This turns out not to be valid for a large cross traffic, although the loss run length of RED is small throughout the range of experiments.

4.3 Influence of redundancy and FEC block size

In this experiment, the size of the FEC block vary, while maintaining the rateh/k constant, in order to obtain the same load increase for the UDP flow and therefore maintain the same overall network load. This way we can directly observe the influence of the FEC block size without the interference of the network load. The load of the system being constant in this case, the PLRBC is fixed for both RED and Drop Tail .However, the PLR decreases when k increases for both queue management schemes. This means that for this configuration of the system, it is interesting to increase the FEC block size. The RED/FEC appears to be more advantageous than DT/FEC concerning the PLR under certain conditions depending on the number of TCP flows, the FEC block size and the amount of redundancy. This is the case even though the PLRBC of RED/FEC is larger than the one obtained for DT/FEC. Indeed, RED queue management is able to give a slight performance improvement.
To summarize, all these results suggest that in case of a constant UDP offered load, the larger the block size, the better the correction rate is for both RED and Drop Tail. It is more interesting to use FEC with RED instead of Drop Tail in case of small number of flows, moderate FEC block size, and a relatively important amount of redundancy.

6.0 Conclusions and Perspectives

This paper studied the effect of a forward error correction (FEC) code on queue management schemes like DropTail and RED. It has been shown in literature that RED losses are spread as compared to Drop Tail. For this reason, one can assume that FEC would be more efficient combined with RED than with Drop Tail. But this study have also shown that RED/FEC does not always perform better than DT/FEC. This turns out to depend on certain parameters, in particular on the number of TCP flows that constitute the background traffic, the FEC block size and the amount of redundancy in a FEC block. It suggest that RED/FEC performs better than DT/FEC when the loss rate is not too large, when there is a relatively important amount of redundancy in the FEC block, and a moderate FEC block size is used.

References

1. Byers, J., Luby, M., Mitzenmacher, M., Rege, A.: A digital fountain approach to reliable distribution of bulk data. In: Proc. ACM SIGCOMM, Vancouver, Canada (1998) 56–67

2. Cidon, I., Khamisy, A., Sidi, M.: Analysis of packet loss processes in high-speed networks. IEEE Transactions on Information Theory39 (1993) 98–108

3. Floyd, S., Jacobson, V.: Random early detection gateways for congestion avoidance. IEEE/ACM Transactions on Networking1 (1993) 397–413

4. Braden, B., et al.: Recommendations on queue management and congestion avoidance in the Internet (RFC 2309). IETF. (1998)

5. McAuley,A.: Reliable broadband communications using a burst era sure correcting code. In: Proc. ACM SIGCOMM'90, Philadelphia, PA, USA (1990) 297–306

6. Floyd, S.: Discussions on setting REDparameters (1997)http://www.aciri.org/ floyd/red.html.

7. NS: Simulator homepage (2004)http://www.isi.edu/nsnam/ns.

8. Brandauer, C., Iannaccone, G., Diot, C., Ziegler, T., Fdida, S., May, M.: Comparison of tail drop and active queue management performance for bulk-data and web-like Internet traffic. In: Proc. ISCC'01, Hammamet, Tunisia (2001)

9. Sanneck, H., Carle, G.: A framework model for packet loss metrics based on run lengths. In: Proc. Multimedia Computing and Networking Conference (MMCM), San Jose, CA, USA (2000) 177–187

10. Biersack, E.: A simulation study of forward error correction in ATM networks. Computer Communication Review22 (1992) 36–47

11. Bonald, T., May, M.: Drop behavior of RED for bursty and smooth traffic. In: Proc. IEEE/IFIP IWQoS'99. (1999)

12. May,M., Bonald, T., Bolot, J.: Analytic evaluation of RED performance. In:Proc.INFOCOM'00. (2000)

13. Alemu,T.: Performance evaluation of quality of service mechanism in the Internet. PhD thesis, University of Montpellier II (2004)