

Efficiency of Loss Differentiation Algorithms in 802.11 Wireless Networks

Stephane Lohier¹ and Yacine Ghamri Doudane²

¹ LIP6 - University of Paris VI, 8, rue du Capitaine Scott, 75015 Paris – France

² IIE-CNAM, 18, allée Jean Rostand, 91025 Evry Cedex – France
stephane.lohier@lip6.fr, ghamri@iie.cnam.fr

Abstract. Loss Differentiation Algorithms (LDA) are currently used to determine the cause of packet losses with an aim of improving TCP performance over wireless networks. In this work, we are interested in the distinction between losses due to interferences on 802.11 links from those due to congestions. To this end, we compare different LDA schemes existing in the literature and a proposal of a cross-layer LDA which use MAC parameters. The efficiency of our solution is then demonstrated through simulations.

1 Introduction

Due to frequent signal failures, the performance of TCP in the 802.11 networks can be significantly degraded, particularly in SOHO (Small Office Home Office) environments or in public points of distribution (Hot Spot) with a wireless last link. Indeed, TCP is conceived for wired networks and is not adapted to react to losses on wireless links. Loss Differentiation Algorithm (LDA) are thus necessary to determine if the segment loss is due to short signal failures caused by interferences in the vicinity of the wireless station and its AP (Access Point) or by congestions due to heavy traffic. In this work, we compare, in the selected context, different LDA schemes and a proposal of an alternative LDA based on a MAC layer parameter which is specifically adapted to the 802.11 links. A simulation set makes it possible to evaluate the effectiveness of the various solutions.

2 The different LDA schemes

As proposed in the literature [1-3], the differentiation decision can be obtained based on TCP parameters, namely packets inter-arrival times, *cwnd* (congestion window) and *RTT* (Round Trip Time). The LDA using packet inter-arrival times at the receiver [1] are unsuited here as they consider that the resolution is always completed at the TCP layer and thus do not take into account the MAC layer loss-recovery present in the 802.11 standard. A set of preliminary measurements (Fig. 1) confirms that, with a first loss-recovery level at MAC-layer, the inter-arrival times are not significant to differentiate these two type of loss. Even for important interference or congestion rates (greater than 70%), at the saturation limit of the wireless link, the obtained values remain relatively close.

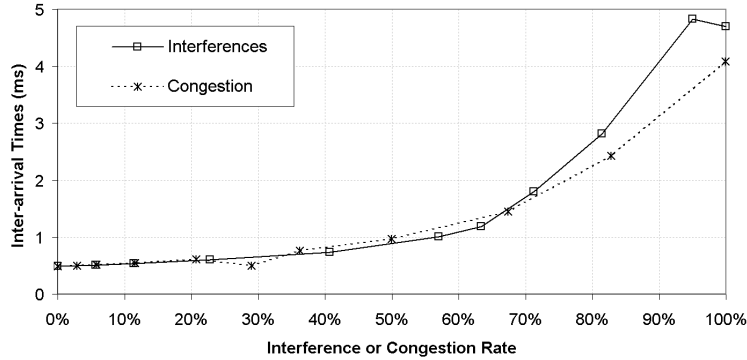


Fig. 1. Efficiency of the Inter-arrival Time LDA on an 802.11 link

For losses caused by relatively important interferences and unsolved by MAC retransmissions, TCP reacts as for congestions: *cwnd* reduced following the reception of three duplicate ACKs. As the *cwnd* variable is updated each time a segment-loss occurs, it became highly variable. The simulation study realized in [2] showed that, due to its variability, *cwnd* is not an accurate and selective variable for differentiation. Finally, in a context with a wireless last hop, *RTT* is the variable presenting the most significant variations in the event of short signal losses compared to congestions. We thus selected three TCP-layer LDA schemes based on comparisons of current *RTT* values with different thresholds and on filters giving more or less weight to the recent samples: the *Vegas Predictor* scheme [2], the *Spike* scheme [1] and the *Flip Flop Filter* [3]. Then, rather than using only TCP-layer parameters which do not take into account the 802.11 specificities, we suggest to use a cross-layer approach as an alternative to conventional TCP-Layer LDA schemes.

The idea of our alternative algorithm is to count the number of MAC retransmissions for each of the n segments composing the current TCP window when the TCP layer is alerted by the reception of three duplicated ACKs. As described in Table 1, if for one of these segments at least, the number of MAC retransmissions (*RetryCount*) is equal to the *Retry Limit* threshold (its default value is 6 for all the 802.11 equipments), we consider that the loss is due to interferences and not to TCP congestion. Indeed, in the case of congestion, the surplus of segments is eliminated from the queue of the concerned node and MAC retransmissions are theoretically not used; inversely, in case of persistent interferences, the segment is dropped by the MAC layer after reaching the *Retry Limit* threshold. This algorithm assumes that for all the not acknowledged TCP segments, the value of *RetryCount* is stored. The *ACKFailureCount* counter available in the 802.11 *Management Information Base* (MIB) gives the number of times that an expected ACK is not received and consequently the value of *RetryCount*.

Note that while the TCP sender is not a wireless host and that the TCP flow is forwarded to the wireless receiver through an AP, an additional stage is necessary. The *LDA_Estimator* is first set at the AP's MAC layer. Then this latter informs the TCP sender by setting consequently the *ELN* (*Explicit Loss Notification*) bit of the TCP header in the ACK segments (i.e. $ELN=LDA_Estimator=1$ in case of

interferences). The loss differentiation is finally made at the TCP sender when receiving three duplicated ACKs.

Table 1. Cross-layer LDA

if (3 dup ack) then	// loss indication in TCP algorithm
LDA_Estimator = 0	// initial value for congestion
for (i = 0 ; i ≤ n ; i ++)	// for all the not acknowledged segments
if (RetryCount = RetryLimit) then	// segment is dropped, probably a short loss
LDA_Estimator = 1	// set value for interferences
end if	
end for	
end if.	

3 Simulation Results

In order to realize a comparative study among the 4 LDA schemes, a set of simulations with a wireless last link undergoing congestions or interferences have been realized. The TCP flow is thus established between a server and a wireless station through its AP. Interferences are caused by the transmission on the same channel of a CBR/UDP flow between two other wireless stations out of the AP coverage and interferences areas. As we deactivated the RTS/CTS mechanism for both transmissions, the AP will not detect CBR transmissions and will thus transmit its TCP segments toward the receiver which is located in the interference area. The duration and the frequency of the interferences will vary according to the size of the packets and the rate of the CBR source (with 1000Bytes packets sent at 1800packets/s the wireless link is completely saturated, which correspond to 100% on the curves). Let us note that the simulated interferences and so the packets losses are carried out in a scenario close to reality (short losses are often caused by transmissions in the same frequency band) and not with a theoretical packet error rate as inaccurately used in most studies. Another CBR/UDP flow is established between the server and a fourth wireless station in order to saturate the AP and induce congestions (for this flow, a rate of 3500packets/s with 1000Bytes per packet is necessary to saturate the link). The simulation results presented in Fig. 2 and Fig. 3 show the accuracy (the percentage of correctly classified losses) of the four LDA schemes according to the interference or congestion rates. For the *Vegas predictor* scheme, we observe that the losses due to low interference rates or high congestion rates are badly classified (less than 60%). This is especially due to the evolution of *cwnd*, which is in these cases inadequately used in conjunction with *RTT* to compute the *Vegas predictor*. The *Spike* scheme, only based on *RTT* variations, gives slightly better results: accuracy higher than 80% in the majority of the cases. The badly classified losses are more random and are mainly due to the calculation mode of the Spike's thresholds. The *Flip Flop* filter is not very efficient, particularly for losses due to interferences. The used algorithm employs many parameters difficult to regulate. Finally, the proposed cross-layer LDA scheme gives the best results. For congestions, there are almost no MAC retransmissions and the *Retry Limit* threshold is never reached, which gives 100% of correctly classified losses. For interferences, some losses are badly classified when the segment is finally received with the last attempt. However accuracy remains in all the cases higher than 90%.




Fig. 2. Accuracy of the 4 LDA schemes with Interferences

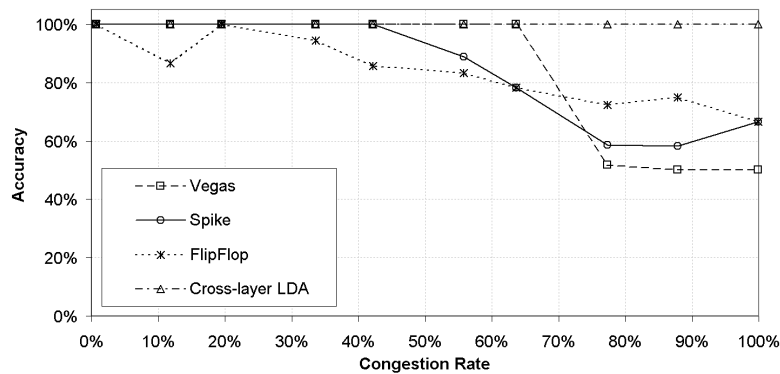


Fig. 3. Accuracy of the 4 LDA schemes with Congestions

4 Conclusion

In this paper, we demonstrated, through comparative simulations, the efficiency of our cross-layer LDA to distinguish congestions from short losses due to interferences in an 802.11 network. This differentiation is a first step to a complete solution integrating an improvement of the TCP behavior. Indeed, in case of interference, the TCP loss recovery mechanism does not have to be triggered which allows the source to achieve higher transmission rates when the MAC channel is rapidly restored.

References

1. S. Cen, P. C. Cosman, and G. M. Voelker, "End-to-End differentiation of congestion and wireless losses," *IEEE/ACM Transactions on Networking (TON)*, Volume 11, Issue 5, October 2003.
2. S. Biaz and N. H. Vaidya, "Distinguishing Congestion Losses from Wireless Transmission Losses: a Negative Result," *Proc. of IEEE 7th International Conference on Computer Communications and Networks, IC3N'98*, October 1998.
3. D. Barman and I. Matta, "Effectiveness of Loss Labeling in Improving TCP Performance in Wired/Wireless Networks," *Boston University Technical Report*, 2002.