Detecting and Filtering Non-responsive Traffic in AQM Queues without Packet Header Examination

Yi Pan
Bren School of Information and Computer Science
University of California, Irvine
Email: ypan@ics.uci.edu

ABSTRACT

In the current Internet, long-term TCP flows and bursty short-term flows both exist. The mix between the aggregated long-term TCP traffic (responsive traffic) and the bursty traffic (non-responsive traffic) has great impact on the performance of Active Queue Management (AQM) algorithms in routers. We introduce a new solution to differentiate between responsive and unresponsive traffic in AQM queues without identifying individual flows through packet headers. The proposed solution is based on statistic measurement on the incoming traffic rate and the AQM packet loss rate. We show the design of our scheme and illustrate the effectiveness through simulation results.

1. INTRODUCTION

Internet congestion control remains an important issue in network research. The Internet congestion control has been described as a feedback control system [1] [2] with two components: end system's TCP protocol that adjusts its sending window size using an Additive Increase and Multiplicative Decrease (AIMD) method in response to packet losses in the network, and router's queue management algorithms that determine the packet loss rate in congestion. Although there are many variants, the de facto AQM algorithm is RED.

It has also been observed that there are two major classes of traffic flows in the Internet: the long-term TCP flows and short-term bursty flows. A long-term TCP flow is characterized by its long duration and usually has its transmission rate controlled by TCP AIMD congestion control mechanism (e.g. FTP bulk file transfer), while a short bursty flow is characterized by its short duration and bursty transmission rate (e.g. HTTP short transactions) without response to congestion in the network. Since most long-term TCP flows are responding to network congestion due to TCP AIMD congestion control, we refer to the long-term TCP traffic as *responsive traffic* while the short-term bursty flows are referred as *non-responsive traffic*.

The mix between the responsive and non-responsive traffic has great impact on the performance of AQM algorithms, since AQM algorithms are designed to detect and determine the congestion signal (i.e. packet loss/ECN mark) with the hope that the end systems will respond. Existence of non-responsive traffic in the input to AQM queues often results in performance degration in the throughput of responsive traffic. Earlier research has focused on differentiating TCP and non-TCP traffic by identifying individual flows and applying different policies to protect well-behaved TCP flows. However, flow identification requires packet header examination, which is a heavy load for high-speed links. Header based flow identification may also fail to identify non-responsive traffic since many non-responsive traffic also uses similar headers as TCP (e.g. HTTP traffic).

In this extended abstract, we propose a new scheme to differentiate responsive and non-responsive traffic without flow identification through packet headers. The proposed scheme is based on sampling traffic rate and AQM packet loss rate at fixed periods. In the following, we first give an overview of the proposed scheme, then present the simulation results showing the effectiveness of the proposed scheme.

2. OVERVIEW

From [2] we know that due to the round trip time delay and the AIMD congestion control mechanism, TCP controlled responsive traffic periodically overshoot the bandwidth and then backoff. According to the AIMD congestion control mechanism used in most responsive traffic, the average periodic cycle of rate fluctuation in responsive traffic is mainly determined by three factors: the round trip time, the number of TCP flows in the aggregated traffic, and the available bandwidth on the bottleneck link. Traffic bursts occurring at time scales other than the average periodic cycle are unlikely to be responsive traffic. Hence, we classify changes in incoming traffic rate as being responsive or non-responsive based on the time scale of those changes.

To implement this scheme, we first estimate the average periodic cycle of the responsive traffic using a traffic estimator. Then, based on the estimated responsive traffic, we apply a wavelet de-noising filter to remove the non-responsive traffic bursts before they enter the AQM queue. To avoid unnecessary drops of nonresponsive traffic, the wavelet filter also let the non-responsive traffic pass the router without being counted in the AQM queue size, when extra buffer space is available. The Fig. 1 shows the overview of the proposed scheme.

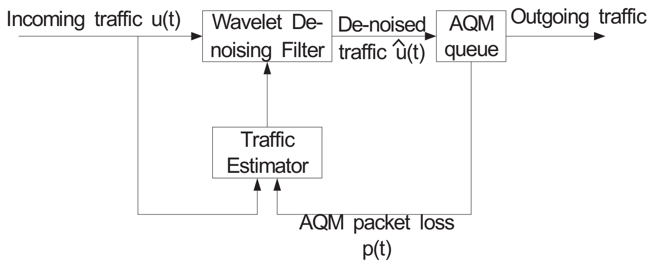

Figure 1: System overview

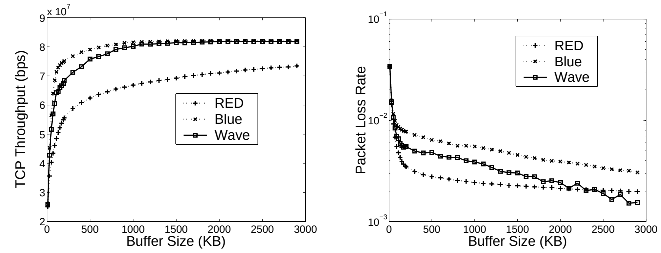

(a) TCP throughput (b) Average Packet Loss Rate

Figure 2: Performance vs Buffer Size

2.1 Traffic Estimator

Applying TCP fluid model in [2], we can derive the following theorem describing the relationship between the observed average periodic cycle and the corresponding AQM packet loss rate for the aggregated responsive traffic.

THEOREM 2.1. *Let the observed long-term average packet loss rate be p_{avg}, and the length of the average periodic cycle of responsive traffic be $Cycle$, in order to keep the aggregated responsive traffic lower than the available bandwidth, the minimum AQM packet loss rate p_{loss} is defined by the following equation:*

$$p_{loss} = \frac{r}{p_{avg}} \left(\frac{N_{loss}}{C * Cycle}\right)^2 \quad (1)$$

in which, r is the targeted percentage of queue busy time ($r \leq 1$), N_{loss} is the observed number of packet loss in the queue busy period and C is the link speed.

The proof is not included due to space limit.

Given the result from Theorem 2.1, we designed a sampling algorithm to achieve the estimation on $Cycle$ and p_{loss}. At first, samples of $Cycle$ are collected. Then, a p_{new} is calculated in each sampled cycle according to Eq. (1). A weighted averaging method is applied to achieve the final estimation on $Cycle$ and p_{loss}.

2.2 De-noising Filter

With the estimation on $Cycle$ and p_{loss}, we apply a wavelet de-noising filter to remove non-responsive traffic bursts at time scales different from $Cycle$. The de-noising filter uses a thresholding function to remove the high peaks of traffic changes at different time scales. Calculation of the de-noising thresholds is one key issue in designing our de-noising filter. Idealistically, the thresholds should be set such that the resultant de-noised traffic is exactly the same as the responsive traffic. We measure the accuracy of de-noising through the Mean Square Error (MSE) between the packet loss according to the de-noised traffic and the estimated packet loss p_{loss} according to the responsive traffic, in each periodic cycle $Cycle$. The goal is to find a set of threshold values to minimize the MSE. Note that we also use extra buffer space to allow non-responsive traffic passing the router, available buffer size puts an additional constraint to how much non-responsive traffic we can filter before packet drop happens. Therefore, we solve the following constraint-optimization problem to get the threshold values.

$$\begin{aligned} \min \quad & f(\vec{T}) = \\ & \sum_{l=1}^{\lceil \frac{n}{Cycle} \rceil} \Big(\sum_{t=(l-1)*Cycle+1}^{l*Cycle} p(t) - Cycle * p_{loss} \Big)^2 \\ \text{s.t.:} \quad & \forall i \in [1...n], \\ & \sum_{t=1}^{i} (u(t) - \hat{u}(t)) \leq B \end{aligned} \quad (2)$$

in which $p(t)$ is the corresponding AQM packet loss for the de-noised traffic, $\vec{T}$ is the threshold values, and B is the extra available buffer. We designed an online algorithm to find the solution for the above optimization problem.

3. PERFORMANCE EVALUATION

We evaluated the proposed scheme with detailed simulation experiments in ns-2. The simulation results presented in this abstract is from a dumbell network with 50 long-term TCP flows. We also included HTTP traffic sources to create a realistic mixture of the Internet traffic. In Fig. 2, we show that our proposed scheme successfully improves the throughput of responsive traffic without performing packet header examination, when buffer is available. Our scheme also achieves a lower packet loss rate with similar TCP throughput to BLUE, a representative AQM algorithm applying flow identification through packet header examination.

4. REFERENCES

[1] V. Hollot, Y. Liu, V. Misra, and D. Towsley, *Unresponsive Flows and AQM Performance*, in Proceedings of IEEE INFOCOM 2003, vol. 1, pp. 85-95, April 2003.

[2] S. Floyd, M. Handley, J. Padhye, and J. Widmer, *Equation-Based Congestion Control for Unicast Applications*, in the Proceedings of ACM SIGCOMM 2000, pp. 43-56, August 2000.