Journal of Network and Computer Applications 75 (2016) 169–180

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

journal homepage: www.elsevier.com/locate/jnca

A new explicit congestion notification scheme for satellite IP networks

Satoshi Utsumi [a,*], Salahuddin Muhammad Salim Zabir [b], Sumet Prabhavat [c]

[a] Fukushima University, Faculty of Symbiotic Systems Science, 1 Kanayagawa, Fukushima, Fukushima, Japan
[b] National Institute of Technology, Tsuruoka College, Department of Creative Engineering, 104 Sawada, Inooka, Tsuruoka, Yamagata, Japan
[c] Faculty of Information Technology, King Mongkut's Institute of Technology Ladkrabang, No. 1, Chalongkrung 1 Soi, Ladkrabang, Bangkok, Thailand

ARTICLE INFO

Article history:
Received 6 January 2015
Received in revised form
15 June 2016
Accepted 29 August 2016
Available online 30 August 2016

Keywords:
TCP (Transmission Control Protocol)
ECN (Explicit Congestion Notification)
Congestion control
Satellite IP network
Analytical model.

ABSTRACT

Longer transmission delay and higher link errors affect data communication performance over satellite links. The problem gets compounded in case a congestion occurs in the corresponding data path. Taking preventive measures before a congestion actually occurs can help in avoiding such situations. Irrespective of the network characteristics, Explicit Congestion Notification (ECN) either through ICMP messages or through marking packets aims to achieve this objective. However, use of ICMP messages for explicit congestion notification leads to vulnerability to attacks by malicious hosts. On the other hand, use of conventional marking schemes that mark data packets in the forward direction can be slow as they have to traverse the satellite links. In this manuscript, we propose a new mechanism of ECN for an even faster notification of an incipient congestion over satellite IP networks. Our proposal, Mark-Reverse ECN, marks packets corresponding to acknowledgement (ACK) segments in the reverse direction to the sender. This leads to a reduction of the latency between detection and actual notification. As such, our mechanism leads to a more accurate and effective approach to congestion control. In addition, our proposed mechanism is free from vulnerability to malicious attacks. We also develop a new analytical model for Mark-Reverse ECN on TCP Reno for performance prediction and validation of simulation outcomes. Simulation results show that our proposed mechanism (i) leads to up to 25% improvement of TCP performance, (ii) reduces packet loss at the congested router by down to about 0% and (iii) helps to improve the performance of conventional TCP when deployed together.

1. Introduction

Longer transmission delay and higher link errors affect data communication performance over satellite links. Longer transmission delay results in slower data communication rate and slower response to changes in the network conditions. Higher link error, on the other hand, leads to loss of data while being transmitted. Both result in reduction of throughput in the corresponding data path. The problem gets compounded in case a congestion occurs in the corresponding data path. Taking preventive measures before a congestion actually occurs can be effective in avoiding such situations. Irrespective of the network characteristics, Explicit Congestion Notification (ECN) provides effective means to handle an incipient network congestion by notifying the corresponding TCP senders in advance. Such notifications can be achieved through ICMP messages or marking packets on the congested routes. However, both the schemes have their inherent weakness under certain conditions.

For example, in Akujobi et al. (2002, 2003), the authors examined use of ICMP Source Quench (ISQ) messages for Backward Explicit Congestion Notification (BECN) to inform the sender of the congestion situation in the network. However, such a mechanism allows the malicious hosts the opportunity to easily launch fake notification of congestion which would reduce the throughput of the target connection unnecessarily (Gont, 2010).

On the other hand, use of packet marking has been proposed as a superior option to packet dropping for notifying an incipient congestion to the sender (Akujobi et al., 2002; Kuhlewind et al., 2013). Usually, such marking is performed on Random Early Detection (RED) Gateways (Floyd and Jacobson, 1993). It has been found that ECN leads to better TCP performance than conventional buffer overflow dropping or RED based early dropping schemes in terms of throughput (Floyd and Jacobson, 1993; Almasri et al., 2012; Kuhlewind and Scheffenegger, 2012) or time delays for interactive traffic. However, if such marked data packets traverse through long delay links like a satellite link, the notification to the sender would take some time to reach the sender. During this time, the sender might continue sending data packets leading to worsening the congestion situation of the network and dropping of data packets if an overflow limit is reached.

* Corresponding author.
E-mail address: u.satoshi0215@gmail.com (S. Utsumi).

http://dx.doi.org/10.1016/j.jnca.2016.08.027
1084-8045/

To overcome the above problems, in this manuscript, we propose a new mechanism of ECN for even faster notification of an incipient congestion over satellite IP networks. We name this as Mark-Reverse ECN. Our proposed scheme marks packets corresponding to acknowledgement (ACK) segments in the reverse direction to the sender. As a result, the latency between the detection of an imminent congestion and its notification to the end host is reduced. Since this allows the end host/sender to take actions in a more timely manner, our mechanism leads to an even more accurate and effective approach to congestion control. Moreover, our proposed mechanism is free from vulnerabilities to throughput-reduction attacks like the ones with ISQ because of the difficulty in obtaining connection information including the current TCP Acknowledgement sequence number etc.

We also develop a new analytical model for Mark-Reverse ECN on TCP Reno for performance prediction and validation of simulation outcomes. Simulation results show that our proposed mechanism (i) leads to up to 25% improvement of TCP performance, (ii) reduces packet loss at the congested router by down to about 0% and (iii) helps to improve the performance of conventional TCP when deployed together.

In our previous research (Utsumi and Zabir, 2014), we devised a mechanism called Wireless Friendly Congestion Control (WFCC), to handle issues resulting from link errors over satellite channels. WFCC yields a significantly higher (up to 250%) throughput than standard TCP and is friendly to co-existing standard TCP flows. In this manuscript, we also show that when Mark-Reverse ECN is applied with WFCC over satellite links (which is equivalent to handling both the problems of link errors and long propagation delay over satellite link), the resulting throughput surpasses the same from other existing schemes.

The rest of this paper is organized as follows. We present a brief outline of relevant TCP congestion control schemes in Section 2. In Section 3, we describe a conventional ECN scheme, i.e., Mark-Tail ECN. In Section 4, we introduce our newly proposed Mark-Reverse ECN scheme. Section 5 describes a new analytical model for Mark-Reverse ECN. In Section 6, we evaluate our scheme comparing with the conventional scheme through simulations. Finally, we conclude this paper in Section 7.

2. Related works

In this section, we explain the related works for improving TCP performance over satellite networks.

Conventional TCP mechanisms of Slow Start, Congestion Avoidance and Fast Recovery have been designed to be well mannered over the traditional Internet. As a result, on links where the bandwidth-delay product is large, TCP is significantly slow in discovering and utilizing the available bandwidth resource optimally. Satellite links belong to this category. Several research efforts address these issues from different perspectives.

HighSpeed TCP (Floyd, 2003) provides a mechanism to improve TCP performance where the Congestion Window is adjusted for networks with large bandwidth delay products. The basic idea behind HighSpeed TCP is to modify the AIMD (Additive Increase and Multiplicative Decrease) as a function of cwnd. However, this scheme, as is not specialized for satellite links, considers link error losses as congestion loss. Also the moderate loss assumption where this is most effective, is often considerably less than practical for satellite links with relatively higher error rate. Hence despite some betterment than conventional TCP, it is not the optimal solution for use over satellite links.

BIC TCP (Xu et al., 2004) and CUBIC TCP (Rhee and Xu, 2005) are also for networks with large bandwidth delay products. BIC TCP uses two congestion window size control policies called additive increase and binary search increase. BIC TCP is better than HighSpeed TCP in terms of bandwidth scalability and RTT (Round Trip Time) fairness (Xu et al., 2004). CUBIC TCP is an improvement of BIC TCP in term of RTT fairness (Rhee and Xu, 2005).

TCP Hybla (Caini and Firrincieli, 2004) has been proposed specifically for high RTT connections like satellite links. The basic idea of TCP Hybla is to obtain for long RTT connections the same instantaneous transmission rate of a comparatively fast reference TCP connection (e.g. wired ones). Although we find TCP Hybla to be quite effective even in cases of high link error rates, we observe that there is still room for proposing a more robust solution to yield better TCP performance over satellite links.

In a different thread, TCP-Peach, a unique approach for satellite links to solve the prohibitively slow inflation of the initial *cwnd* in Slow Start phase and recovering from a segment loss due to link errors has been proposed in Akyildiz et al. (2001). TCP-Peach (Akyildiz et al., 2001; Morabito et al., 2001) is a congestion control scheme based on probing segments, namely, dummy segments, for satellite networks with long propagation delays and relatively high link error rates. TCP-Peach replaces Slow Start and Fast Recovery in TCP Reno or in TCP NewReno with Sudden Start and Rapid Recovery, respectively. In Sudden Start and Rapid Recovery, the sender sends dummy segments, which are low-priority data segments carrying duplicate data blocks, to detect the available bandwidth in the network and tries to use the detected bandwidth for the connection.

TCP-Peach+ (Akyildiz et al., 2002) is an improvement of TCP-Peach. TCP-Peach+ replaces Sudden Start and Rapid Recovery algorithms of TCP-Peach by Jump Start and Quick Recovery respectively. Jump Start and Quick Recovery use NIL segments for probing the available bandwidth in the network, instead of dummy segments. Because NIL segments carry unacknowledged data blocks, they can be used to recover lost segments by the receivers. At high link error rates, NIL segments in some scenario, may be more efficient than dummy segments, which carry duplicate data blocks.

TCP-Cherry (Utsumi et al., 2008) is an improvement of TCP-Peach+. TCP-Cherry replaces Jump Start and Quick Recovery algorithms of TCP-Peach+ by Fast-Forward Start and First-Aid Recovery respectively. Fast-Forward Start and First-Aid Recovery use supplement segments for carrying new data blocks along with probing the available bandwidth in the network. At high link error rates, supplement segments in some scenario, may be more efficient than NIL segments, which carry unacknowledged data blocks.

As HighSpeed TCP, BIC TCP and CUBIC TCP are quite aggressive, standard TCP experiences starvation of throughput when deployed together with these aggressive TCP versions (Trinh et al., 2004; Munir et al., 2007).

TCP Hybla is also more aggressive than standard TCP over satellite networks. Hence it affects standard TCP performance in such networks (Marcondes et al., 2008).

TCP-Peach, TCP-Peach+ and TCP-Cherry need for all intermediate nodes in the path to have a priority queueing system. It is difficult in a realistic Internet environment.

Various research works consider Explicit Congestion Notification (ECN) as a means for congestion control. A few of them, like Akujobi et al. (2002, 2003), consider use of Backward ECN (BECN) in the form of ICMP Source Quench Messages (ISQ) for the purpose. In addition, the authors of Akujobi et al. (2003) propose two thresholds for RED gateway for ECN and ISQ based BECN. However, use of ISQ leads to vulnerability to attacks by aggressive hosts and has thus been considered unfair as well as ineffective. As such, in 1995, support of ISQ has been deprecated at the router level as outlined in RFC 1812 (Baker, 1995). Then, the use of ISQ has also been deprecated at the transport level in 2012 (RFC 6633 (Gont, 2012)).