Journal of Network and Computer Applications 68 (2016) 56–64

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

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

TCP congestion control algorithm for heterogeneous Internet

Zhiming Wang*, Xiaoping Zeng, Xue Liu, Man Xu, Ya Wen, Li Chen

College of Communication Engineering, Chongqing University, Chongqing 400030, China

ARTICLE INFO

Article history:
Received 29 August 2015
Received in revised form
28 October 2015
Accepted 28 March 2016
Available online 30 March 2016

Keywords:
Transmission control protocol (TCP)
Congestion control
Fairness
Bandwidth delay product (BDP)
Heterogeneity

ABSTRACT

The available bandwidth, round trip time (RTT) and packet loss rate can vary over many orders of magnitude, which characterizes the heterogeneity of the Internet. To cope with the heterogeneity in Internet congestion control, we propose a new TCP protocol known as INVS that contains three key components: (1) INVS employs an exponential-function-based growth function of the congestion window, which is more efficient than the cubic function in CUBIC; (2) INVS introduces an adaptive increase factor into the growth function to ensure that the window growth rate matches the path condition, and this increase factor measures the path condition using a custom function of the available bandwidth and minimum RTT; (3) INVS adopts an adaptive queue threshold in the loss classification scheme to improve the performance of TCP over lossy links. In addition, INVS requires modification only on the TCP sender and traces the path capacity to enable quick convergence of the congestion window. The performance analysis and evaluation show that INVS achieves good throughput, fairness, RTT-fairness and utilization in heterogeneous networks.

1. Introduction

The traditional communication network has evolved into a seamless global Internet that encompasses a variety of heterogeneous IP networks (wired, wireless and satellite networks) (Damnjanovic et al., 2011; Pan and Yuguang, 2012). The Internet Congestion Control Research Group has outlined a variety of distinguishing link and path characteristics for the Internet (Papadimitriou et al., 2011): (1) the available bandwidth can be either scarce over radio links or abundant over high-speed optical links; (2) the round trip time (RTT) ranges from much less than a millisecond (local interconnects) to notably large, i.e., equal or greater than 500 ms (satellite links); (3) the packet loss rate (PLR) ranges from notably low over optical fiber links (less than 10^{-6}) to notably high over certain wireless links (higher than 1%). Consequently, the available bandwidth, RTT and PLR of the Internet can vary over many orders of magnitude, which characterizes the heterogeneity of the Internet. The Transmission Control Protocol (TCP), which provides reliable end-to-end data delivery across the Internet, must cope with the different transmission media that travel as Internet traffic over the best-effort IP networks (Ros and Welzl, 2013). With the increasing heterogeneity of the Internet, this mission of TCP has become more difficult, and the heterogeneity of the Internet has been identified as one of the global challenges in Internet Congestion Control (Papadimitriou et al., 2011).

Many publications have addressed congestion control algorithm in terms of the heterogeneity of the Internet. Most of the literature focuses on one aspect of the heterogeneity or need for network aid and can be grouped into several classifications.

First, Hybla (Caini and Firrincieli, 2004) is proposed for satisfying RTT-fairness to cope with RTT difference, which is often considered as an RTT-unfairness problem. However, when passing through lossy links in which the packet loss is no longer the indication of congestion, Hybla flows back off upon packet loss, which is often caused by link error. As a result, Hybla flows cannot use the available bandwidth effectively, which is treated as performance degradation (Afanasyev et al., 2010).

Second, HS-TCP (Afanasyev et al., 2010), CUBIC (Ha et al., 2008), Compound TCP (CTCP) (Song et al., 2008), ER-TCP (Park et al., 2012) and Illinois (Liu et al., 2008) have been proposed to cope with the difference of bandwidth and RTT, especially the poor bandwidth utilization of traditional TCP in high-speed networks and long-delay networks (Alrshah et al., 2014), and CUBIC and CTCP have been widely deployed in the Internet. However, these proposals are not compatible with lossy wireless links and also suffer from unfairness problems (Leung and Li, 2006; Jingyuan et al., 2013). For example, when sharing bandwidth with TCP Reno or CTCP, CUBIC, STCP and HS-TCP (which employ highly aggressive

* Corresponding author.
E-mail addresses: jamewzm@163.com (Z. Wang), zxp@cqu.edu.cn (X. Zeng), maojiu2005@qq.com (X. Liu), 323109402@qq.com (M. Xu), 327704920@qq.com (Y. Wen), chenli@cqu.edu.cn (L. Chen).

http://dx.doi.org/10.1016/j.jnca.2016.03.018
1084-8045/

growth functions of the congestion window regardless of high-speed and long-delay networks or traditional networks) take up most of the total bandwidth, which is treated as a fairness problem (Jingyuan et al., 2013). Due to the high random packet loss over wireless links, especially for third generation (3G) communication links, mobile devices (smartphones, tablets, etc., most of which use CUBIC as the default TCP congestion control algorithm) suffer from significant TCP performance degradation (Wan et al., 2014; Chen et al., 2012).

To cope with the PLR difference, especially the performance degradation over lossy links, many proposals attempt to distinguish non-congestion losses from congestion losses by comparing the queue delay or queue length with a fixed threshold and apply different multiplicative decrease schemes accordingly (Afanasyev et al., 2010; Jingyuan et al., 2013; Leung and Li, 2006; Cheng Peng and Liew, 2003; Casetti et al., 2002; Caini et al., 2007; Ye et al., 2005; Sundararajan et al., 2011), e.g., Veno (Cheng Peng and Liew, 2003), TCP Westwood (Casetti et al., 2002), etc. With the increasing variety of wireless networks, the accuracy of the fixed threshold approaches in loss classification is decreasing. Certain other proposals use particular congestion control algorithms. PEPsal (Caini et al., 2007) and indirect-TCP (Ye et al., 2005) split the connection into several portions to shield wireless links from wired links. TCP-Jersey (Afanasyev et al., 2010) adopts a router-aided explicit congestion warning scheme to determine the cause of the loss. Network Coded TCP (Sundararajan et al., 2011) incorporates network coding into TCP to provide significant throughput gains in lossy networks. However, these particular algorithms were not widely accepted and deployed due to their special requirements.

Finally, to cope with differences in bandwidth, RTT and PLR simultaneously, the authors in Xie et al. (2010) proposed a congestion control scheme based on RTT, and PLR, but taking PLR as the congestion degree is not suitable for paths with lossy links. In Yuan-mei (2009), a link-type-based congestion control scheme is proposed that uses Vegas over satellite links, Westwood over wireless links, HS-TCP over high-speed links and NewReno over traditional links. Due to the demand of adding link type information in IP packets by routers, this method could not be deployed in the Internet. In Nan et al. (2015), TCP BRJ grades the congestion into five grades based on delay jitter to differentiate random losses from congestion losses. Because the delay jitter in the Internet varies widely, the reliability of grading according to the fixed delay jitter threshold must be further investigated.

As the diversity of communication links increases, the TCP congestion control algorithm must match the path condition for each connection. Although TCP has been extensively studied in recent past years, most of the proposals focus on particular environments, and the proposals for heterogeneous networks cannot be deployed in or not suitable for the heterogeneous Internet. Therefore, in this paper, we propose a new TCP congestion control algorithm, i.e., INVS, to cope with the multiple differences that exist in the heterogeneous Internet. This method incorporates the merits of CUBIC and Westwood and distinguishes itself from other proposals in three key aspects: (1) it employs an exponential-function-based growth function for the congestion window that is more efficient than the cubic function in CUBIC; (2) it introduces an adaptive increase factor in the growth function to scale the window growth rate to match the path condition. This increase factor measures the path condition using a custom function of the available bandwidth and the minimum RTT in which a parameter γ is also introduced to reflect the conservation of RTT influence; (3) it adopts an adaptive queue threshold in the loss classification scheme to improve the performance of TCP over lossy links. In addition, INVS requires modification of TCP sender only and traces the path capacity to enable quick convergence of the congestion window.

The remainder of the paper is organized as follows. Section 2 describes the proposed algorithm in detail. Section 3 provides performance analysis, and Section 4 presents the performance evaluation. The deployments of various TCP protocols result in Internet TCP traffic composed of mixed TCP flows, and the performance is investigated with mixed TCP variants. Section 5 concludes the paper.

2. Proposed congestion control algorithm: INVS

To cope with the heterogeneity of the Internet, we propose an exponential-function-based congestion window growth function to ensure efficiency, introduce an adaptive increase factor to match the window growth rate with the path condition, and adopt an adaptive queue-threshold scheme to improve the performance of loss classification.

2.1. Growth function of the congestion window

For efficiency in large BDP networks, INVS uses an exponential-function-based convex window growth function at the front of each congestion avoidance (CA) phase and a concave function for exploring available resources. With the convex function, the TCP sender can increase the congestion window quickly at the beginning of the CA phase to fully utilize the available bandwidth and provide slow increases thereafter. The size of the congestion window at time t in INVS at the front of each CA phase is given as follows:

$$cwnd(t) = cwnd_{sp}(1 - (1 - \beta)\alpha^t), \quad 0 < \alpha < 1 \tag{1}$$

where t is the time from the beginning of the congestion avoidance phase, $cwnd_{sp}$ is the congestion window at *saturationpoint* (which represents the state in which the available resources of the path are fully used), β is the multiplication decrease factor, and α is a parameter that controls the window growth rate.

INVS uses the convex growth and the concave probing through a congestion window update scheme as shown in formula (2). Fig. 1 shows the exponential-function-based congestion window growth functions and the cubic function of CUBIC. INVS1, INVS2 and INVS3 are the exponential-function-based growth functions when parameter k increases. Comparing INVS1 with CUBIC, we note that the exponential-function-based growth function is more efficient when the duration of the CA phase is the same. The

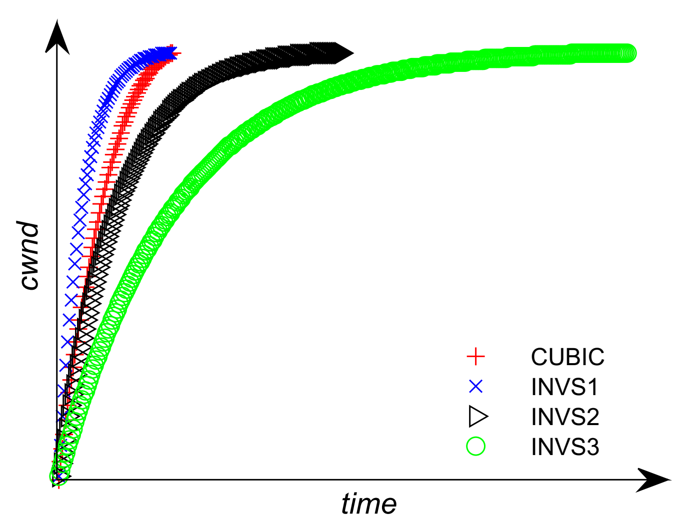

Fig. 1. Convex growth function.