Available online at www.sciencedirect.com

ScienceDirect

Computer Communications 30 (2007) 479–495

computer communications

www.elsevier.com/locate/comcom

Analytical and experimental evaluation of TCP with an Additive Increase Smooth Decrease (AISD) strategy

Laura M. Voicu, Steven Bassi, Miguel A. Labrador *

Department of Computer Science and Engineering, University of South Florida, Tampa, FL 33620, USA

Received 21 June 2006; received in revised form 9 September 2006; accepted 20 September 2006
Available online 13 October 2006

Abstract

This paper presents an analytical model to study the throughput performance of TCP with an Additive Increase Smooth Decrease (AISD) strategy. The new AISD-based TCP protocol produces smooth transfer rates and performs the flow and congestion control functions of TCP. As a result, it is suitable for streaming applications and solves the unfairness problem that occurs when TCP and UDP share the same bottleneck link. The AISD strategy is very simple to implement, only modifies the multiplicative part of the Additive Increase Multiplicative Decrease strategy of TCP, and can be implemented in any TCP version. The smooth part of the strategy is implemented using a low pass filter that considers history in the calculation of the congestion window. The modeling of this new strategy raises new challenges compared to the classical TCP modeling in two ways: first, it needs to be adapted to a more complex dynamism of the congestion window, and second, the model needs to incorporate a scheduler that periodically updates the value of the congestion window. The model is progressively built to finally characterize the steady-state send rate and throughput of a flow as a function of the loss probability, the round-trip time (RTT), the time-out interval, and the scheduler interval. The performance of this AISD-based TCP is compared analytically and experimentally with TCP Reno, and its superior performance is demonstrated.

Keywords: TCP modeling; Low pass filter; Scheduler; Web100; Dummynet

1. Introduction

The amount of streaming traffic over the Internet continues to grow. Many applications not available a few years ago are considered main stream today. This is the case of videoconferencing and voice, which were traditionally transported over circuit-switched networks. The Internet, dedicated to carry traffic from data-oriented applications, has now to handle both types of traffic, while providing good performance. Unfortunately, TCP and UDP, the current transport layer protocols used to carry data and streaming applications, do not work together very well. It is well known that during congestion events, when TCP and UDP share the same bottleneck link, UDP is not fair to TCP, obtaining a disproportionate share of the bandwidth [5]. In [25], the need for end-to-end congestion control for streaming applications has been recently emphasized as the solution to this fairness problem and the congestion collapse problem in the Internet [5].

In this paper, we present an analytical model to study the steady-state throughput performance of a TCP protocol with an Additive Increase Smooth Decrease (AISD) strategy, like the one proposed in [1]. The new AISD-based TCP protocol changes the Multiplicative decrease part of current TCP versions for a Smooth decrease strategy that makes the protocol appropriate for real-time applications. In addition, the protocol retains all the other mechanisms and properties of TCP. Thus, the reliability of the protocol is as good as TCP's for data applications, and the flow and congestion control mechanisms eliminate or substantially

* Corresponding author. Tel.: +1 813 974 3260; fax: +1 813 974 5456.
E-mail addresses: voicu@cse.usf.edu (L.M. Voicu), sbassi2@cse.usf.edu (S. Bassi), labrador@cse.usf.edu (M.A. Labrador).

0140-3664/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.comcom.2006.09.010

reduce the unfairness problem. This strategy, which is implemented with a low pass filter that considers history in the calculation of the congestion window, has been shown to provide better throughput performance to data applications, smooth transfer rates to streaming applications, and complete fairness if data and streaming applications utilize the protocol [1].

The model presented here is very complete compared with current TCP models. As in other models, ours also captures the behavior of the protocol as a function of the channel loss rate and the round trip time (RTT) of the connection, and includes the fast retransmit mechanism and the effects of time-outs. In addition, the model indirectly includes the fast recovery and slow start processes, not usually considered in most models. These two mechanisms are embedded in the smooth decrease strategy, which governs the dynamics of the congestion window after loss events, and eliminates the need of the slow start phase. However, the modeling of the AISD strategy goes beyond current models because of the complexity of the processes that need to be included. First, the analysis of the dynamics of the congestion window is complicated by a recurrent formula that also depends on time and on the type of event that triggered the update. Second, the calculation of the congestion window does not occur only when a loss is detected, as in the case of TCP, but also, in the absence of such losses. The low pass filter that implements the AISD strategy runs a scheduler that triggers the update of the congestion window at periodic intervals. We believe that in addition to the performance evaluation of the protocol, this paper provides a guide for other researchers to model these other mechanisms, which are embedded in several other protocols.

The theoretical model of the AISD-based TCP protocol includes two processes: the dynamics of the congestion window and the packet loss process. These two processes are included in the model progressively. First, the model considers only losses detected by triple-duplicate acknowledgments, with the restriction that one such loss happens during each scheduler's interval. The model then incorporates losses detected via triple-duplicate acknowledgments, while eliminating the above restriction. Finally, the model includes losses detected via time-outs.

The rest of the paper is organized as follows. Section 2 provides a brief literature review in the area of TCP modeling. Section 3 describes the AISD-based TCP protocol and the filter utilized to calculate the congestion window. Section 4 presents the analytical characterization of the steady-state send rate and throughput of the protocol as a function of the loss rate, round trip time, and the duration of the scheduler's interval. Performance results that illustrate the behavior of the proposed protocol compared with TCP Reno are included here. In Section 5, a Linux-based Web100-Dummynet testbed is utilized to provide experimental results comparing the current SACK version of TCP with the AISD-based TCP protocol. Finally, Section 6 concludes the paper.

2. Related work

Since its inception several years ago, the Transmission Control Protocol (TCP) has grown to be the most important and widely used communication protocol. Over all these years, many algorithms and mechanisms have been embedded into TCP to improve its performance. For instance, TCP Tahoe [8], TCP Reno [8], TCP Newreno [9], TCP SACK [4,15], TCP Westwood [18] and TCP Vegas [2] are some of the important versions of TCP that have emerged as a result of these enhancements. These mechanisms and algorithms are very well known and, therefore, not described here. Similarly, several protocols and algorithms have been proposed to solve the TCP-UDP unfairness problem, which can be broadly classified in router-based or end-to-end solutions. A survey on this topic can be found in [17] and the references therein. One of the most important problems facing all these solutions is the lack of practical implementation. This is due to performance issues, because the algorithms are too heavy to be implemented in routers or do not perform well enough, or due to the fact that they constitute complete new solutions, difficult to replace in practice. The AISD-based TCP protocol proposed here has the potential to overcome these problems, as it is TCP-based and very easy to implement in practice.

Several analytical models and performance evaluations of the most important TCP versions have appeared in the literature. These works have concentrated on different aspects of the protocol [3,7,10,11,14–16]. While some concentrated on modeling the throughput of infinite TCP connections as a function of the round-trip time and the packet loss rate [10,11,14–16], others concentrated on modeling the latency of finite connections as a function of the transfer size, round-trip time, and packet loss rate [3,15]. Since the proposed protocol is TCP-based, studying the classical TCP analytical models is the starting point in understanding the process of construction of its analytical model.

Analytical models for TCP Tahoe, Reno, and SACK are included in [15] to estimate both the latency and the steady-state throughput of the protocols. Based on the analytical models, the three versions of TCP are then compared. The throughput of TCP Vegas is modeled in [14], as a function of the average round trip time, minimum round trip time, and the loss rate of the transfer. The model includes the slow-start, congestion avoidance, and congestion recovery mechanisms. The latency of a TCP connection is modeled in [3] as a function of the transfer size, round trip time, and packet loss rate. The model includes both the connection establishment and data transfer phases. This approach is intended to predict the performance of both short and long TCP flows under different packet loss conditions. Perhaps the most well-known model that characterizes the steady-state send rate of a bulk transfer TCP flow is the one presented in [11]. The model captures both the fast retransmit and the time-out mechanisms of TCP Reno and calculates the send rate and throughput of a TCP flow as a function of the loss rate and the round trip