Available online at www.sciencedirect.com
ScienceDirect
Computer Networks 51 (2007) 1800–1814
Computer Networks
www.elsevier.com/locate/comnet

A novel high speed transport protocol based on explicit virtual load feedback

Xiaomeng Huang *, Chuang Lin, Fengyuan Ren

Department of Computer Science and Technology, Tsinghua University Beijing 100084, China

Available online 28 November 2006

Abstract

Since traditional TCP congestion control is ill-suited for high speed networks, designing a high speed replacement for TCP has become a challenge. From the simulations of some existing high speed protocols, we observe that these high speed protocols make the round-trip time bias problem and the multiple-bottleneck bias problem more serious than for standard TCP. To address these problems, we apply the population ecology theory to design a novel congestion control algorithm. By analogy, we treat the network flows as the species in nature, the sending rates of the flows as the population number, and the bottleneck bandwidth as the food resources. Then we extend the construction method of population ecology models to design a control model, and implement the corresponding end-to-end protocol with virtual load factor feedback, which is called Explicit Virtual Load Feedback TCP (EVLF-TCP). The virtual load factor is computed based on the information of the bandwidth, the aggregate incoming traffic and the queue length in routers, and then senders adjust the sending rate based on the virtual load factor. Theoretical analysis and simulation results validate that EVLF-TCP achieves high utilization, fair bandwidth allocation independent of round-trip time, and near-zero packet drops. These characteristics are desirable for high speed networks.

Keywords: Congestion control; High bandwidth-delay product networks; Population ecology models

1. Introduction

With the rapid advances in the deployment of high speed and long distance links in the Internet, it is natural for people to use TCP to transfer data in high Bandwidth-Delay Product (BDP) networks. However, many recent studies report that the classical Additive Increase Multiplicative Decrease (AIMD) mechanism [1,2] used in traditional TCP congestion control does not perform well in high BDP networks since increase by one packet per round-trip time (RTT) is too slow, and decrease by halving the window per loss event is too drastic.

This problem motivates the proposal of several new high speed transport protocols. It is possible to classify these protocols into two types: implicit congestion feedback and explicit congestion feedback [3]. Implicit congestion feedback schemes, such as HSTCP [4], STCP [5], BICTCP [6], H-TCP [7], CUBIC [8], Layered TCP [9], FAST [10] and many more, use the loss or delay as the congestion signal which is incidental to the delivery of the packets.

* Corresponding author. Tel.: +86 10 62772487.
E-mail address: xmhuang@csnet1.cs.tsinghua.edu.cn (X. Huang).

1389-1286/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2006.11.003

Explicit congestion feedback schemes, such as XCP [11], CADPC [12], EMKC [13] and VCP [14], use direct communication from routers to tell the end hosts the state of the network, accomplished by sending special packets or by changing some fields in packets as they traverse the routers. The use of explicit congestion feedback is likely to result in superior congestion control protocols that converge faster and keep a lower packet loss rate than protocols using implicit congestion feedback.

Although these high speed TCP variants have demonstrated better performance than traditional TCP in several limited network experiments, as we observed, there are still two limitations of most of them in more complex network environments.

Firstly, it is well-known that traditional TCP has a bias against flows passing through multiple bottlenecks [16–18], and a bias against flows with longer round-trip time (RTT). The multiple-bottleneck bias problem is that the TCP flow passing through multiple bottlenecks will have a lower rate than the TCP flow passing through one bottleneck, and the RTT bias problem is that the TCP flow with longer RTT will have a lower rate than the TCP flow with shorter RTT. Most of the implicit congestion feedback schemes mainly depend on increasing the congestion window faster and decreasing less than traditional TCP to achieve high utilization. Unfortunately, these protocols make the multiple-bottleneck bias problem and the RTT bias problem more serious than in the traditional TCP case. See in Section 5 for detailed simulation results.

Secondly, how to choose an operation point of congestion control is still an open issue. For the most of implicit loss-based approaches such as HSTCP, STCP and BICTCP, full utilization, even if achievable, often comes at the cost of severe oscillations and potentially large queuing delays. On the contrary, some explicit congestion feedback such as XCP and VCP and some implicit delay-based approaches such as FAST, choose the zero queue length as the operation point of congestion control. Thus a near-zero queue length and near-zero packet loss rate is easy to achieve at the cost of a little spare bandwidth. Their main difference is that one works in a slight overload state, and the other one works in a slight underload state. In our opinion, an ideal transport protocol should function such that the average aggregate traffic should be equal to the bottleneck capacity, and the buffer should not empty in the steady state, yielding full utilization and small queuing delay.

In this paper we propose a novel explicit feedback based solution. Our main contribution is that we apply population ecology theory to design a novel congestion control model based on network information. Population ecology, which is the branch of ecology that studies the dynamics of populations, is analogous to congestion control. In essence, the major problem in both cases is resource allocation. For instance, in population ecology, different species in nature compete for food resources, and in congestion control, network flows have to compete for bandwidth resources. In order to study the rules of population evolution as time goes on, biologists have proposed many classical and practical models, such as the Malthus Model [19], the Logistic Model [20] and the Lotka–Volterra Model [21]. If we treat the network flows as the species in nature, the rate of the flows as the population number, and the bottleneck bandwidth as the food resources, these mature population ecology models would be helpful in designing a novel congestion control algorithm. Therefore, we extend the construction method of population ecology models to design a congestion control model, and implement the corresponding protocol. In routers, we combine the information of the bandwidth, the aggregate incoming traffic rate and the instantaneous queue length to compute a virtual load factor, and store it in an extension packet header. Then senders adjust the sending rate based on the virtual load factor. We called this protocol Explicit Virtual Load Feedback TCP (EVLF-TCP). To gain insight into the behavior of EVLF-TCP, we prove the local asymptotic stability and analyze the efficiency and fairness convergence speed of EVLF-TCP, then determine the preset parameters of EVLF-TCP. Finally, we evaluate the performance of EVLF-TCP using ns2 simulations. We conclude that EVLF-TCP achieves high utilization, fair bandwidth allocation independent of round-trip time, and near-zero packet drops. However, one limitation of EVLF-TCP is that it takes a long time to converge to fair operation. It will be studied further.

The remainder of this paper is organized as follows. Section 2 introduces the fundamentals of population ecology theory, and discusses the basic relationship between the population ecology models and the congestion control models. Section 3 presents the control model and the implementation of EVLF-TCP. Section 4 evaluates its performance through theoretical analysis and simulation results are given in Section 5. Finally, Section 6 concludes the paper and describes future work.