A Simple Markovian Model of TCP Startup Behavior (Extended Abstract)

Stefano Giordano, Michele Pagano, Gregorio Procissi, Raffaello Secchi

Department of Information Engineering - University of Pisa
Via Caruso 16, 56122 Pisa, Italy
Phone: (+39) 050 2217 511
Fax: (+39) 050 2217 522
E-mail: {s.giordano, m.pagano, g.procissi, raffaello.secchi}@iet.unipi.it

Abstract

The paper proposes a Markovian approach to the performance evaluation of the ESSE (Early Slow Start Exit) modification of the TCP congestion control mechanism. ESSE takes advantage of estimations of the optimal pipesize at the sender side to properly select the initial slow start threshold. Previous simulative experiments have shown that ESSE allows to speed-up TCP connections and significantly reduce the packet drop rate at the bottleneck. This work makes a step further in understanding the ESSE behavior by developing a model of TCP source to evaluate the influence of different settings of slow start threshold on TCP performance. As confirmed by comparison with simulations, the model provides, in significant less time than simulations, accurate estimates of typical performance indicators such as the average completion time and average drop rate of short-lived TCP connections.

1 Introduction and Motivation

Recent studies [1, 2] have revealed that the majority of the Internet flows are short-lived (*mice*), while a smaller number of long-lived connections carry the most of the Internet traffic (*elephants*). The domination of *mice* over *elephants* is mainly due to the wide-spread diffusion of web browsing, which is characterized by small (26-32KB on average) and frequent data transfers. This phenomenon makes the purely loss-driven scheme of standard TCP implementations inadequate for short flows, since short connections are often terminated well before reaching the steady state behavior. Considering that the network steady state is determined by the superposition of many short-lived flows, the initial slow start phase has a heavy impact on the network stability and a great influence on the overall TCP performance.

For these reasons, a growing interest towards the proposal of novel approaches to improve TCP transient behavior is taking place and has recently led to [3, 4]. In particular, the ESSE algorithm (Early Slow Start Exit) [4] belongs to the class of measurements-driven congestion control schemes first adopted by TCP Vegas and Westwood [5] in that it uses estimations of the available optimal pipesize to initialize the slow start threshold (ssthresh) during the TCP start-up phase.

This paper presents a simple Markovian approach to the analysis and performance evaluation of the ESSE modification to TCP startup behavior. In the following, a brief description of the ESSE modification to TCP is given first. Next, the rationale of the analytic model is presented together with the numerical and simulative scenario in which it has been validated. In the last section, we conclude the paper with final remarks.

2 TCP Congestion Control Schemes

The slow start phase of TCP occurs when the cwnd is less than the ssthresh. In this phase, every received acknowledgment increases the cwnd by one segment, so that the cwnd exponentially increases at every RTT. Obviously, the correct selection of ssthresh is a critical issue for the performance of the whole system, and it should be chosen to guarantee fair allocation of resources.

A reasonable setting of ssthresh should be close to the bandwidth-delay product (BDP) and can be generally written as the ratio between the transmission delay and the spacing of packets at the bottleneck. In our research, four specific pipesize estimators are proposed and estimation results are thoroughly compared.

3 The Markovian Model

To evaluate the behavior of TCP startup modification, we develop an analytic model that, under a few common assumptions, takes into account the impact of the initial ssthresh setting on the overall TCP performance. By assuming as state variables of the system the triple (w, s, n), where w denotes the window size, s the ssthresh and n the numbers of packet to be delivered, the rationale of the model is that the state evolves according to a Discrete Time Markov Chain (DTMC) on the state space, where the time granularity is given by the RTT (Round Trip Time). The transition probability matrix is determined by the knowledge of the packet loss probability, assumed as independent of the size of the congestion window (a reasonable assumption in the case of short-lived connections).

Our model aim at evaluating the average *completion time* and average *drop rate* of TCP connections expressed as functions of the underlying states of Markov chain. For instance, the mean time $\mathcal{C}_{(w,s,n)}$ to complete the transmission of n packets when the initial cwnd is w and initial ssthresh is s can be expressed recursively as a function of the completion time associated to states that can be reached from (w, s, n).

Notice that the model can be applied to both standard TCP implementations, for which the initial ssthresh is not set (or equivalently $s = n$), as well as to the ESSE modification to TCP, for which the ssthresh is set to the current estimated pipesize. Indeed, the analytic model proves to be consistent with respect to traditional TCP analysis as it can be easily verified that, by assuming the first order approximation proposed in [6], it provides an accurate estimation of the TCP Newreno latency.

4 Model Analysis and Validation

The model has been validated by comparing its numerical solutions to the results of simulations carried out by aggregating short-lived TCP connections sharing a FIFO/RED bottleneck link in wide-area and metropolitan-area network scenarios. In particular, we compared the results of simulations for several load levels with the results provided by model with the ssthresh set to network BDP.

The ESSE modification to the TCP slow start phase have been implemented by adding the pipesize estimation module outlined in section 2 to ns-2[7] simulator that inherits from Newreno the standard scheme of loss recovery.

5 Conclusions

This paper proposes a simple Markovian model of TCP sources that provides accurate estimates of the average completion time of TCP connections as a function of RTT, loss probability and initial ssthresh setting. Since this model explicitly considers the ssthresh as a state variable, it allows to study the behavior of the TCP modifications,such as ESSE, which are based on the dynamic setting of ssthresh. The analytical model has been validated by means of detailed simulations in different network scenarios referring to a single bottleneck topology.

References

[1] L. Guo and I. Matta. The War Between Mice and Elephants. In *procedings of ICNP'01*, November,2001.

[2] B.A. Mah. An Empirical Model of HTTP Traffic. In *proceedings of INFOCOM*, April 1997.

[3] J. Hoe. Improving the Start-up Behavior of a Congestion Control Scheme for TCP. In *proceedings on Applications, technologies, architectures and protocols for computer communications*, pages 270–280. ACM Press, 1996.

[4] R. Lo Cigno, G. Procissi and M. Gerla. Sender-Side TCP Modifications: An Analytical Study. In *proceedings of NETWORKING*, pages 600–611, 2002.

[5] S. Mascolo, C. Casetti, M. Gerla, M. Sanadidi and R. Wang. TCP Westwood: Bandwidth Estimation for Enhanced Transport over Wireless Links. In *proceedings of the 7th annual international conference on Mobile computing and networking*, pages 287–297. ACM Press, 2001.

[6] Neal Cardwell, Stefan Savage, and Thomas Anderson. Modeling TCP latency. In *Proceedings of INFOCOM (3)*, pages 1742–1751, 2000.

[7] Information Sciences Institute. The network simulator - ns2. http://www.isi.edu/.