Available online at www.sciencedirect.com

ScienceDirect

Computer Networks 51 (2007) 2788–2804

Computer Networks

www.elsevier.com/locate/comnet

Enhancing TCP for networks with guaranteed bandwidth services

Yujie Zhu *, Aravind Velayutham, Oyebamiji Oladeji, Raghupathy Sivakumar

School of Electrical and Computer Engineering, Georgia Institute of Technology, 330605 Georgia Tech Station, Atlanta, GA 30332, United States

Received 18 April 2006; received in revised form 22 October 2006; accepted 29 November 2006
Available online 12 January 2007

Responsible Editor: Giacomo Morabito

Abstract

In this paper, we consider TCP based applications with bandwidth guarantees, but can also benefit from any additional best-effort service offered by the network. Through simulations we show that default TCP cannot offer such applications the ideal throughput – the aggregate throughput of the reserved bandwidth and the best effort bandwidth. To illustrate the reasons for its degraded performance, we study TCP's congestion window adaptation and self-clocking mechanisms in detail. Based on the insights obtained from the study, we propose an adaptation of TCP called GTCP that employs changes to TCP's congestion control mechanisms to provide applications the optimal aggregate throughput of best-effort and reserved bandwidth. Compared with TCP, GTCP does not involve any additional implementation overhead, and only the sender need to be changed (the receiver remains to be a default TCP implementation). Through simulations and experiments over the Internet we show that GTCP achieves significantly better performance than default TCP in the target environment.

Keywords: TCP; GTCP; Diffserv; Intserv; Congestion control; Bandwidth reservation

1. Introduction

The Internet is increasingly being seen as an infrastructure that will lead to *digital convergence*, supporting a diverse set of applications including streaming media, video conferencing, digital telephony, and plain data transfer. This has in turn resulted in the focus on providing applications with quality of service (QoS) assurances that cannot be provided in the current best-effort service model of the Internet. The integrated services (*intserv*) [1] and differentiated services (*diffserv*) architectures standardized by the Internet Engineering Task Force (IETF) are examples of approaches that can provide QoS assurances to applications. The *intserv* approach provides fine-grained per-flow QoS assurances through its guaranteed and controlled load classes of service. On the other hand, the *diffserv* approach provides coarse-grained QoS assurances through its expedited forwarding and assured forwarding services.

* Corresponding author. Tel.: +1 404 385 3012.
E-mail addresses: zhuyujie@ece.gatech.edu (Y. Zhu), vel@ece.gatech.edu (A. Velayutham), siva@ece.gatech.edu (R. Sivakumar).

1389-1286/$ - see front matter
doi:10.1016/j.comnet.2006.11.026

In this paper, we consider applications that subscribe to bandwidth reservations in networks that are QoS enabled through either of the above-mentioned architectures. More specifically, we target applications that would still require the reliable sequenced delivery services of the ubiquitous TCP transport protocol. Examples of such applications include HTTP based web applications, file transfer applications, video download and playback applications, email, remote terminal access, integrated messaging, etc. We contend that such applications, while utilizing their reserved bandwidths, can further benefit from any best-effort bandwidth provided by the network.[1]

Some approaches that enable such applications to gain the joint benefit of the reserved bandwidths and the best-effort data rates include: (a) application layer striping [2,3] across two connections, one with the reserved bandwidth, and the other with the best effort service, (b) using an unaware transport layer such as TCP over a network that provides both the guaranteed bandwidth and the best effort service, and (c) using a transport layer protocol that is aware of the two types of data service and provides the application with the aggregated throughput. However, application layer striping approaches cannot provide the optimal aggregate throughput performance. In fact, such approaches can be shown to be severely limited by the lower of the bandwidths enjoyed by the two connections over which the data is striped [4]. Similarly, it has been shown (and we demonstrate later in the paper) that using an unaware protocol such as TCP will not result in the ideal aggregation of throughput [5].

In this paper, we propose an enhanced version of the TCP transport layer protocol called *GTCP*. GTCP successfully delivers to the application the aggregated data rates offered by the guaranteed and best-effort services. It adapts the slow-start, congestion avoidance, congestion control, and timeout reaction mechanisms in TCP to achieve the improved performance. The properties of GTCP include (a) changes only in the congestion control algorithm of TCP, (b) no additional timers required when compared to TCP, (c) TCP friendly best-effort service performance, and (d) no changes needed at the receiver. While similar attempts at an enhanced transport layer protocol have been made in related work [5,6], such approaches either rely on high implementation overheads (in the form of additional per-flow timers), or do not achieve the optimal aggregate throughput performance. We discuss related work in detail in Section 3.

Through a comprehensive set of simulations, we evaluate the performance of GTCP and show that it achieves the desired performance over a variety of scenarios, and scales well across a variety of factors. The contributions of this work are twofold:

- We profile the performance of the default TCP protocol over networks that provide both guaranteed and best-effort bandwidth services, and identify the key reasons for its non-optimal performance.
- We propose an adaptation of TCP called GTCP that achieves close to ideal aggregation of bandwidths provided by the guaranteed and best-effort services, through changes to TCP's congestion control mechanisms and without incurring any additional implementation overheads.

The rest of the paper is organized as follows: Section 2 identifies the key drawbacks of TCP and motivates the design of GTCP. Section 3 discusses related work, and Section 4 presents the design of GTCP in detail. Sections 5 and 6 evaluates the GTCP protocol under various network conditions. And finally, Section 7 discusses related research issues and concludes the paper.

2. TCP over networks with guaranteed bandwidth

2.1. Factors contributing to performance degradation

In this section we identify the key elements in TCP's congestion control mechanisms that contribute to its non-optimal performance when operating over networks with bandwidth guarantees. We classify these elements under the following two categories:

- *Congestion window*: TCP uses a window based congestion control algorithm, where the congestion window *cwnd* represents the instantaneous rate enjoyed by a connection normalized to its round-trip time. In other words,

$$cwnd = rate * rtt,$$

[1] Note that providing such applications with their best-effort share of bandwidth is also fair if we assume realistically that the basic best-effort service level agreement with the network service provider exists.