Computer Communications 31 (2008) 3883–3893

Contents lists available at ScienceDirect

Computer Communications

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

PBProbe: A capacity estimation tool for high speed networks ☆

Ling-Jyh Chen [a,*], Tony Sun [b], Bo-Chun Wang [a], M.Y. Sanadidi [c], Mario Gerla [c]

[a] Institute of Information Science, Academia Sinica, 128, Sec. 2, Academia Road, Taipei 11529, Taiwan
[b] PacketMotion Inc., 110 Baytech Drive, 2nd Floor San Jose, CA 95134, USA
[c] Department of Computer Science, University of California at Los Angeles 4732 Boelter Hall, Los Angeles, CA 90095, USA

ARTICLE INFO

Article history:
Received 17 August 2007
Received in revised form 19 May 2008
Accepted 19 May 2008
Available online 14 August 2008

Keywords:
Analysis
High speed networks
Link capacity estimation
Simulation
Experiments

ABSTRACT

Knowledge about the bottleneck capacity of an Internet path is critical for efficient network design, management, and usage. In this paper, we propose a new technique, called PBProbe, for estimating high speed links rapidly and accurately. Although it is based on CapProbe, instead of relying solely on packet pairs, PBProbe employs the concept of "Packet Bulk" and adapts the bulk length to compensate for the well known problem with packet pair-based approaches, namely the lack of accurate timer granularity. As a result, PBProbe not only preserves the simplicity and speed of CapProbe, but also correctly estimates link capacities over a much larger range. Using analysis, we evaluate PBProbe with various bulk lengths, network configurations, and traffic models. We then perform a set of experiments to evaluate the accuracy of PBProbe on the Internet over wired and wireless links. Finally, we perform emulation and Internet experiments to verify the accuracy and speed of PBProbe on high speed links (e.g., the Gigabit Ethernet connection). The results show that PBProbe is consistently fast and accurate in the majority of test cases.

1. Introduction

Estimating the bottleneck capacity of an Internet path is a fundamental research problem in computer networking, since knowledge of the capacity is critical for efficient network design, management and usage. In recent years, with the growing popularity of emerging technologies, such as overlay, peer-to-peer (P2P), sensor, grid, and mobile networks, it has become increasingly important to have a simple, fast and accurate tool for capacity estimation and monitoring. To accommodate the diversity of network arrangements, a capacity estimation tool should also be scalable and applicable to a variety of network configurations.

A number of techniques have been proposed for capacity estimation of generic Internet paths [3,7,14,16,18,20,29]. Among them, Pathrate [14] and CapProbe [18] are widely recognized as fast and accurate tools for generic network scenarios. However, CapProbe is a round-trip estimation scheme that only works well on paths with a symmetric bottleneck link, while Pathrate is based on histograms and may converge slowly if the initial dispersion measurements are not unimodal. As a result, CapProbe has difficulty estimating the capacity of asymmetric links [12], and Pathrate performs poorly on wireless links [18].

To address the above problems, specialized capacity estimation tools have been proposed for specific and emerging network scenarios. For instance, ALBP [24] and AsymProbe [12] are intended for capacity estimation on asymmetric links, and AdHoc Probe [11] tries to estimate the end-to-end path capacity of wireless networks. However, recent studies have shown that, for emerging high speed network links (i.e., gigabit links), estimating high speed link capacity remains a major challenge [17,19]; hence, a simple, fast, and accurate technique is desirable.

In this paper, we propose a capacity estimation tool, called PBProbe, for high speed network links. Although PBProbe is inspired by CapProbe, instead of relying solely on one pair of packets, it employs the concept of "Packet Bulk" to adapt the number of probe packets in each sample in accordance with the dispersion measurement. More specifically, when the bottleneck link capacity is expected to be low, PBProbe uses one pair of packets as usual (i.e., the bulk length is 1). For paths with high bottleneck capacities, PBProbe increases the bulk length and sends several packets together. This enlarges the dispersion between the first and last packets, and resolves the well known timer granularity problem. Timer granularity is the main challenge in estimating the capacity of high capacity links [17,19]. This issue is discussed in detail later in the paper.

We study the performance of PBProbe under different Poisson cross traffic loads via analytical models; and evaluate the technique's accuracy and speed through testbed experiments, a laboratory testbed, and on high speed Internet paths. We also compare it

---

☆ A preliminary version of this paper was published in the Proceedings of the 2006 IFIP Networking Conference [10]. This version extends the analysis section in [10] with additional traffic models. Furthermore, it contains a more comprehensive set of testbed experiments than those in the previous paper, and demonstrates PBProbe's compatibility with general networks. Hence, this journal submission is a much more thorough and authoritative presentation of PBProbe.

[*] Corresponding author. Tel.: +886 2 2788 3799x1702; fax: +886 2 2782 4814. *E-mail address:* clljj@iis.sinica.edu.tw (L.-J. Chen).

0140-3664/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.comcom.2008.05.047

with Pathrate. The results show that PBProbe can accurately and rapidly estimate the link capacities in all tested scenarios, thereby outperforming Pathrate in most experiments.

The remainder of this paper is organized as follows. In Section 2, we summarize related work on capacity estimation. In Section 3, we describe PBProbe. Section 4 contains an analysis of PBProbe and an evaluation of its speed and accuracy with Poisson cross traffic, Deterministic cross traffic, and Pareto ON/OFF cross Traffic. In Section 5, we describe PBProbe experiments for general network scenarios, and demonstrate the approach's compatibility with general Internet links. In Section 6, we evaluate PBProbe on high speed links (e.g., Gigabit Ethernet) in our emulator testbed as well as on the Internet. Then, in Section 7, we summarize our conclusions.

2. Background and overview

2.1. Related work

Previous research on capacity estimation relied on delay variations among probe packets, as illustrated by Pathchar [16], or on dispersion among probe packets, as in Nettimer [20] and Pathrate [14]. Pathchar-like tools (such as pchar [7] and Clink [3]) are limited in terms of accuracy and speed, as reported in [18,21]. Moreover, they evaluate the capacity of a link based on the estimates of previous links along the path; thus, estimation errors accumulate and amplify with each link measured [29].

Dispersion-based techniques suffer from other problems. In particular, Dovrolis' analysis in [14] shows that the dispersion distribution can be multi-modal due to cross traffic, and that the strongest mode of such a distribution may correspond to (1) the capacity of the path; (2) a "compressed" dispersion, resulting in capacity over-estimation; or (3) the Average Dispersion Rate (ADR), which is always lower than the capacity. Another dispersion-based tool, SProbe [29], exploits the SYN and RST packets of the TCP protocol to estimate the downstream link capacity, and employs two heuristics to filter out samples affected by cross traffic. However, SProbe does not work efficiently when the network is heavily utilized [23].

Unlike the above approaches, CapProbe [18] uses both dispersion measurements and end-to-end delay measurements to filter out packet pair samples distorted by cross traffic. The method has been shown to be both fast and accurate in a variety of scenarios. The original implementation of CapProbe uses ICMP packets as probe packets, and measures the bottleneck capacity on a round-trip basis. As a result, the capacity estimate does not reflect the higher capacity link when the path is asymmetric. Other difficulties are encountered when intermediate nodes block ICMP packets [25], or employ priority schemes to delay ICMP packet forwarding (e.g., the Solaris operating system limits the rate of ICMP responses, and is thus likely to perturb CapProbe measurements) [30].

Recent capacity estimation studies have extended the target network scenarios to more diverse environments. For instance, Lakshminarayanan et al. evaluated tools for estimating the capacity and available bandwidth of emerging broadband access networks [22], while Chen et al. extended CapProbe to estimate the *effective path capacity* in ad hoc wireless networks [11]. In addition, ABLP [24] and AsymProbe [12] have been developed for capacity estimation of the increasingly popular asymmetric links, such as DSL and satellite links.

Nonetheless, capacity estimation on high speed links remains a challenge. Though recent studies verified the accuracy of Pathrate in estimating gigabit links [28], the evaluations were conducted on an emulator-based testbed, which cannot represent realistic Internet dynamics [15]. Thus, an experimental evaluation of capacity estimation on high speed links is still required.

In this paper, we propose a novel packet bulk technique called PBProbe for estimating the capacity of high speed links. PBProbe is based on CapProbe, but it probes the bottleneck link capacity using UDP packets (instead of the ICMP packets used by CapProbe). We summarize the CapProbe algorithm in the next subsection.

2.2. CapProbe: an overview

CapProbe is a packet pair-based capacity estimation technique that has proven both fast and accurate in several scenarios [18]. Conceptually, when a back-to-back packet pair is transmitted over a network, and assuming both packets arrive at the bottleneck link unperturbed (i.e., back-to-back) by cross traffic on previous links, they are always dispersed at that link according to the bottleneck capacity. Although the dispersion will accurately reflect the bottleneck capacity (as shown in Fig. 1c), if either packet in a pair has been queued due to cross traffic, the dispersion of the sample might be expanded or compressed; "expansion" leads to under-estimation and "compression" leads to over-estimation of the capacity (as shown in Fig. 1a and b).

CapProbe combines dispersion measurements and end-to-end delay measurements to filter out packet pair samples distorted by cross traffic. The basic idea is that if a sample has not been queued by cross traffic, it can be used to estimate the bottleneck capacity correctly. For such "good" samples, the sum of the delays of the packet pairs, called the *delay sum*, does not include cross-traffic induced queuing delay, and is indeed smaller than the delay sum of the samples distorted by cross traffic. Thus, these "good" samples can be identified easily because the delay sum will be the minimum of the delay sums of all packet pair samples. We call this delay sum the *minimum delay sum*. In this way, the capacity can be estimated by the equation:

$$C = \frac{P}{T}, \tag{1}$$

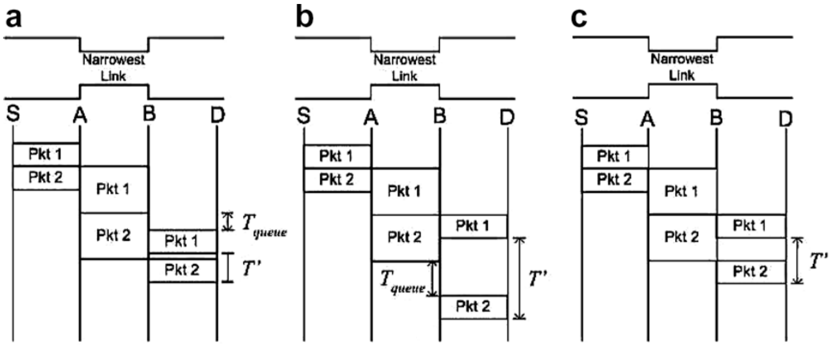

Fig. 1. (a) Over-estimation caused by "compression", (b) under-estimation caused by "expansion", (c) the ideal case.