

Reliable Asymmetric Path Capacity Measurement

Edmond W. W. Chan[‡], Xiapu Luo[§], and Rocky K. C. Chang[‡]

Department of Computing[‡]

The Hong Kong Polytechnic University

{cswwchan|csrchang}@comp.polyu.edu.hk

College of Computing[§]

Georgia Institute of Technology

csxpluo@cc.gatech.edu

ABSTRACT

Existing non-cooperative methods for network capacity measurement are quite restrictive for capacity-asymmetric paths, and their results could be affected by adverse network conditions (e.g., packet delay, packet loss, and packet reordering). In this paper, we propose to use two types of packet pairs—round-trip packet pair and two-way packet pair—and the minimum-delay-difference method for reliable capacity measurement. Our preliminary results show that the new proposal is correct and achieves accurate results even for highly capacity-asymmetric paths.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Measurement Techniques

General Terms

Measurement, Experimentation, Performance

Keywords

Network capacity, Non-cooperative, Packet pair

1 Introduction

A number of non-cooperative capacity measurement methods [5, 2] have been proposed in the past. Due to the packet size limitations, however, these tools cannot obtain accurate results for capacity-asymmetric paths. For example, DSLprobe [2] can only be used for measuring remote ADSL endpoints. Moreover, their measurement accuracy could be further impaired by packet delay, loss, and reordering events on the path.

In this paper, we explain the inadequacy of the existing non-cooperative methods and propose a new tool for reliable asymmetric capacity measurement. The new

tool employs TCP data packets and can detect whether the probe and response packets are lost or reordered. Second, it implements a new minimum delay difference technique (MDDIF) [1] that is more resilient to the cross-traffic interference than the traditional packet-pair dispersion method. Third, it can measure highly asymmetric paths from either endpoint of a path. Since the first two points are already covered in [4, 1], we will focus on the last point below.

2 Background

We consider a capacity-asymmetric round-trip path between a measuring node and a remote node with *forward-path capacity* C_f and *reverse-path capacity* C_r , where $C_f \neq C_r$. We let $C_{f/r} = C_f/C_r$ and further classify a capacity-asymmetric path into either a *Fast-forward-Slow-reverse* (FS) path if $C_{f/r} > 1$ or a *Slow-forward-Fast-reverse* (SF) path if $C_{f/r} < 1$. Moreover, the degree of capacity asymmetry decreases with $C_{f/r}$ for a SF path but increases with $C_{f/r}$ for a FS path.

2.1 The packet-pair dispersion method

Many existing non-cooperative tools (e.g., [5, 2]) are based on the packet-pair dispersion (PPD) method. A measuring node dispatches a sequence of two back-to-back probe packets (packet pairs). Each probe packet (with a packet size of S_f) elicits a response packet (with a packet size of S_r) from the remote node. An unbiased (round-trip) packet-pair dispersion (PPD) for a packet pair $P = \{p_1, p_2\}$ is given by [3]

$$\delta = \max \left\{ \frac{S_f}{C_f}, \frac{S_r}{C_r} \right\}, \quad (1)$$

where S_f/C_f and S_r/C_r are the *forward-path PPD* and *reverse-path PPD*, respectively.

Based on Eqn. (1), for measuring C_f , the measuring node could set $S_f \gg S_r$, such that $S_{f/r} > C_{f/r}$; and for measuring C_r , it sets $S_r \gg S_f$. However, this approach suffers from the packet size limitations imposed by the MTU and the measuring node's time resolution. For reverse-path measurement, for instance, using $S_{f/r} = 260/1500$ can only measure asymmetric paths with a degree of capacity asymmetry of 0.173 or higher.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CoNEXT Student Workshop'09, December 1, 2009, Rome, Italy.

Copyright 2009 ACM 978-1-60558-751-6/09/12 ...\$5.00.

2.2 Existing non-cooperative tools

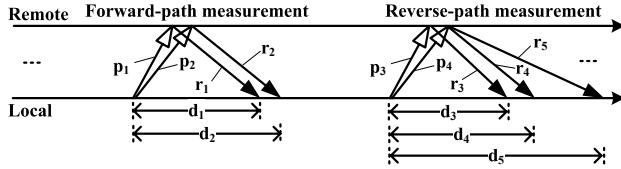


Figure 1: Existing capacity measurement methods.

Sprobe [5], based on the PPD method, can ensure $S_{f/r} > C_{f/r}$ for the forward-path measurement for both FS and SF paths by using 1500-byte TCP SYN probe packets and 40-byte TCP RST response packets. As shown in Fig. 1, $C_f^{\text{Sprobe}} = \frac{S_f}{d_2-d_1}$, where S_f is the size of p_1 . However, due to the packet size restriction, the forward-path estimate for FS paths could be affected by the reverse-path PPD. On the other hand, its reverse-path measurement is immune to the interference of the forward-path PPD, because it triggers two back-to-back TCP data packets for measurement. As shown in Fig. 1, $C_r^{\text{Sprobe}} = \frac{S_r}{d_5-d_4}$, where S_r is the size of r_4 .

DSLprobe [2] conducts asymmetric capacity measurement only for remote ADSL users (i.e., FS paths). For the forward-path measurement, similar to Sprobe, DSLprobe dispatches a sequence of large TCP data packets (> 40 bytes) to elicit a sequence of small TCP RST packets (40 bytes) from the remote node. For the reverse-path measurement, it uses the same types of packets but with the same size of 40 bytes. As shown in Fig. 1, $C_r^{\text{DSL}} = \frac{S_r}{d_4-d_3}$, which can underestimate the reverse-path capacity for SF paths.

3 A new capacity measurement tool

The new tool is distinguished from Sprobe and DSLprobe by the probing technique, the measurement data obtained for the capacity estimation, and the techniques for removing distorted samples. It exploits two types of packet pairs for asymmetric capacity measurement. Unlike classical packet pairs, these packet pairs eliminate the packet size restriction and do not require different probe and response packet sizes.

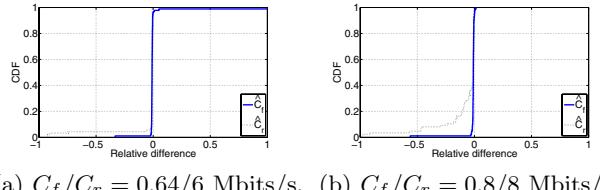
To avoid the reverse-path PPD interference, the new tool uses the minimum delay difference method (MD-DIF) for measuring the forward-path PPD. MDDIF obtains the minimal possible delay (minDelay) for a first probe packet and a second probe packet, both of which do not necessarily belong to the same packet pair. The difference of the two minDelays can be used to obtain the path capacity [1]. The new tool applies similar technique to eliminate the forward-path interference during the reverse-path measurement.

We have constructed the two types of packet pairs, such that they can be integrated with the probes used in OneProbe [4, 1]. Therefore, the new tool can identify

loss and reordering of the probe and response packets based on the OneProbe's probing technique. For example, it will receive a different TCP response packet pattern if a TCP probe packet of either type is lost. As a result, the new tool filters all packet pairs that do not trigger the expected TCP response packets to ensure that the capacity measurement come from lossless and order-preserved probe and response packets.

3.1 Preliminary results

We launched the new tool with 1420 bytes probe/response packet size in our campus to measure the capacity of the paths to 100 PlanetLab nodes located in 26 geographical locations. The measuring node was connected with the Click modular router to emulate two types of asymmetric link capacity scenarios: (i) $C_f/C_r = 0.64/6$ Mbits/s, and (ii) $C_f/C_r = 0.8/8$ Mbits/s. Figs. 2(a) and 2(b) report the CDFs for the relative differences between \hat{C}_f and C_f , and between \hat{C}_r and C_r for the two scenarios. From these figures, more than 95% of the forward-path measurement deviated less than 3% from the true capacity. However, notice that the new tool's reverse-path capacity measurement was less accurate than its forward-path measurement, especially when $C_r = 8$ Mbits/s, which could be caused by the PlanetLab node's uplink capacity being less than C_r .



(a) $C_f/C_r = 0.64/6$ Mbits/s. (b) $C_f/C_r = 0.8/8$ Mbits/s.

Figure 2: CDF of the relative differences between \hat{C}_f and C_f , and between \hat{C}_r and C_r using the new measurement tool.

4 References

- [1] E. Chan, X. Luo, and R. Chang. A minimum-delay-difference method for mitigating cross-traffic impact on capacity measurement. In *Proc. ACM CoNEXT*, 2009.
- [2] D. Croce, T. En-Najjary, G. Urvoy-Keller, and E. Biersack. Capacity estimation of ADSL links. In *Proc. ACM CoNEXT*, 2008.
- [3] L. Chen, T. Sun, G. Yang, M. Sanadidi, and M. Gerla. End-to-end asymmetric link capacity estimation. In *Proc. IFIP Networking*, 2005.
- [4] X. Luo, E. Chan, and R. Chang. Design and implementation of TCP data probes for reliable and metric-rich network path monitoring. In *Proc. USENIX Annual Tech. Conf.*, 2009.
- [5] S. Saroiu, P. Gummadi, and S. Gribble. Sprobe: A fast technique for measuring bottleneck bandwidth in uncooperative environments. In *Proc. IEEE INFOCOM*, 2002.