Spatio-Temporal Available Bandwidth Estimation with STAB

Vinay J. Ribeiro[‡], Rudolf H. Riedi[†] and Richard G. Baraniuk^{‡*}

[‡]Department of Electrical and Computer Engineering
6100 Main Street, Rice University
Houston, TX 77005

[†]Department of Statistics
6100 Main Street, Rice University
Houston, TX 77005

ABSTRACT

We study the problem of locating in space and over time a network path's *tight* link, that is the link with the least available bandwidth on the path. Tight link localization benefits network-aware applications, provides insight into the causes of network congestion and ways to circumvent it, and aids network operations. We present *STAB*, a light-weight probing tool to locate tight links. STAB combines the probing concepts of self-induced congestion, tailgating, and packet chirps in a novel fashion. We demonstrate its capabilities through experiments on the Internet and verify our results using router MRTG data.

Categories and Subject Descriptors

C [Computer Systems Organization]: C.2 Computer-Communication Networks C.2.5 Local and Wide Area Networks

General Terms

Measurements, Performance

Keywords

Probing, bandwidth, available bandwidth, estimation, tailgating, bottleneck, tight link, chirps

1. INTRODUCTION

Locating the link with the minimum available bandwidth on a path, the *tight* link, has many important uses. First, it can provide insight into the causes of congestion and suggest ways of circumventing it. While intuition suggests that tight links are located at the very edge of the network, recent studies indicate that tight links can also occur at intra-ISP and peering links [1]. Second, it can enhance certain applications that benefit from knowing whether paths share a common tight link or not [2]. Third, real-time tight link location information can aid network operators in adjusting traffic routes and detecting network anomalies.

We develop a new tool to locate the tight link of a path in space and over time called the *spatio-temporal available bandwidth* tool (STAB). STAB uses a novel probing scheme that combines the powerful concepts of *self-induced congestion* and *packet tailgating*. A significant advantage of STAB over existing tools such as

*This work was supported by NSF grants ANI-0099148 and ANI-0338856, DoE SciDAC grant DE-FC02-01ER25462, the DARPA/AFRL grant F30602-00-2-0557, and the Texas Instruments Leadership University program.
Email: {vinay, riedi, richb}@rice.edu, URL: spin.rice.edu.

Figure 1: Available bandwidth estimation of path segments using tailgating trains. Large packets congest the network up to node l after which they vanish due to TTL expiry. Small packets carry timing information to the receiver.

Figure 2: Chirp probe train. Packet interarrival times decrease exponentially.

BFind [1] and *TReno* [3] is that it introduces a light probing load onto the network and so does not influence TCP cross-traffic significantly. Crucial to the STAB's efficiency are special probe trains called *chirps* [4]. We validate STAB through an Internet experiment.

2. STAB

This section provides a brief overview of STAB. We describe the key concepts of self-induced congestion, packet tailgating, and chirp probe trains, and then describe STAB's tight link localization technique. We begin with the definition of available bandwidth.

Available bandwidth: Numbering the links on a path $1, 2, \dots, N$ starting from the source (see Figure 1), we define the available bandwidth of link l , A_l , as its average unused capacity. Terming the set of links $h, \dots, k$ as *segment* (h, k) , we define $A(h, l)$ as the minimum available bandwidth among all links comprising (h, k) . The available bandwidth of the path, A , is simply $A(1, N)$.

Self-induced congestion: The principle of self-induced congestion allows a straightforward technique for estimating A . It relies on the following heuristic: if the probing bit-rate R exceeds A then the probe packets become queued at some router, resulting in an increased transfer time. On the other hand, if $R < A$, then the packets face no extra delay. We thus estimate A simply as the probing rate at the onset of congestion.

Chirp trains: In a chirp probing train the interarrival time between successive packets decreases exponentially (see Figure 2). As a re-

Figure 3: (a) Experiment topology. Available bandwidth from the sender to intermediate hops l , $A(1,l)$, for path (b) SLAC–Rice and (c) UIUC–Rice measured from a 1 hour interval.

sult, chirps rapidly sweep through a wide range of probing bit-rates using few packets. This allows an efficient available bandwidth estimation scheme based on the self-induced congestion principle [4].

Packet Tailgating: Packet-tailgating is a powerful technique that provides local information about segments of network paths. It uses special probe trains consisting of large packets interleaved with small tailgating packets (see Figure 1). The large packets exit the path midway due to limited TTLs¹ but the small packets travel to the destination while capturing important timing information.

Tight link localization: STAB employs packet-tailgating chirps to locate the tight link of a path. A tailgating chirp is a chirp as depicted in Figure 2 except that each packet is replaced by a large packet closely followed by a small tailgating one. STAB keeps the TTL of the large packets within each chirp fixed. Chirps with large packet TTL set to l provide estimates of $A(1,l)$. By varying the large packet TTL from chirp to chirp we thus obtain estimates of $A(1,l)$ for $l = 1, 2, \dots, N$.

STAB exploits a fundamental property of A_l to locate the tight link, namely that $A(1,l)$ decreases with l up to the tight link after which it remains constant. We thus estimate the tight link as the link l after which estimates of $A(1,l)$ become more or less constant.

3. INTERNET EXPERIMENTS

This section describes an Internet experiment with STAB. The experiment uses end-hosts at 3 geographically diverse US institutions, the University of Illinois at Urbana Champaign (b), the Stanford Linear Accelerator Center (SLAC), and Rice University (see Figure 3 (a)). We locate the tight link across paths SLAC–Rice and UIUC–Rice simultaneously and address the following questions. Where do the tight links occur on the paths? Are they near the edge or in the interior? Do the observations correlate between the two paths? Do the results compare favorably with available bandwidth data obtained from routers?

STAB estimates link 13 and link 15 as the tight links for paths SLAC–Rice and UIUC–Rice respectively. We note two interesting facts. First, both paths share the same tight link. From router MRTG data we ascertain that this link is a 100Mbps link with available bandwidth 73Mbps. Second, the tight link is near the edge of the path as we would intuitively expect. We now compare the STAB available bandwidth estimates to router MRTG data. Since we were not able to obtain MRTG data from all routers on either path we can only compute an upper bound for $A(1,l)$. This upper bound is the minimum available bandwidth of all known links in the segment $(1,l)$.

¹Each IP router decrements the packet time-to-live (TTL) field by 1. If the TTL value decrements to zero, the packet is dropped.

From Figure 3(b) we make several interesting observations about the path from SLAC to Rice. First, the estimates of $A(1,l)$ roughly decrease with l as expected. Second, for $l = 1, 2, \dots, 10$ the STAB estimates of $A(1,l)$ are about 400Mbps. This is in accordance with the fact that segment $(3,10)$ has available bandwidth in excess of 800Mbps and that the end-host at SLAC is on a Gigabit network. Third, the STAB estimates are close to the MRTG upper bound for $l = 13, \dots, 15$. Fourth, the STAB estimates are larger than the MRTG upper bound for $l = 11$ and $l = 12$ but within a factor of 2. This makes STAB locate link 13 as the tight link while the MRTG data indicates that the tight link is link 12. We note however that links 12 and 13 have comparable available bandwidths of 63Mbps and 73Mbps respectively and thus a small error in available bandwidth measurement can lead to STAB locating the tight link at link 13.

The results for path UIUC–Rice are similar to that of path SLAC–Rice as we see from Figure 3(c). Notice however that all estimates are below 140Mbps unlike for path SLAC–Rice. This is due to the fact that the first link on the path is a 100Mbps link. We also observe that the STAB estimates are close to the MRTG upper bound values.

4. CONCLUSIONS

In this paper we have studied the problem of identifying tight links in space and over time. We described STAB which is to the best of our knowledge the first light-weight probing tool to locate tight links. We presented promising results from an Internet experiment. STAB has the potential to contribute to several applications including network management, load balancing, and anomaly detection. The code is publicly available at [5].

5. REFERENCES

- [1] A. Akella, S. Seshan, and A. Shaikh, “An empirical evaluation of wide-area Internet bottlenecks,” *Internet Measurement Conference*, 2003.
- [2] D. Rubenstein, J. Kurose, and D. Towsley, “Detecting shared congestion of flows via end-to-end measurement,” *ACM SIGMETRICS*, 2000.
- [3] M. Mathis and J. Mahdavi, “Diagnosing Internet congestion with a transport layer performance tool,” *INET*, 1996.
- [4] V. Ribeiro, R. Riedi, R. Baraniuk, J. Navratil, and L. Cottrell, “pathchirp: Efficient available bandwidth estimation for network paths,” *Passive and Active Measurement Workshop*, 2003.
- [5] <http://www.spin.rice.edu/Software/>.