QuickProbe: Available Bandwidth Estimation in Two Roundtrips

George Kola and Mary K. Vernon
Computer Sciences Department, University of Wisconsin-Madison
Madison, WI
kola@cs.wisc.edu , vernon@cs.wisc.edu

Categories and Subject Descriptors: C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks

General Terms: Performance, Measurement.

Keywords: Available bandwidth estimation

1. INTRODUCTION

An accurate estimate of the available (or unused) bandwidth in a network path can be useful in many applications, including route selection in an overlay or multi-homed network, initial bit-rate selection for video streams or improving the slow start phase of existing TCP protocols.

Previously proposed methods for estimating available bandwidth include Pathload [2], PTR [1], Spruce [4], and references therein. The Pathload technique has been shown to be reasonably accurate under a wide range of conditions by independent researchers [3,5]. PTR is more efficient and has been shown to have accuracy similar to Pathload in a more limited set of experiments [1]. With an accurate measure of the bottleneck link capacity, Spruce has been found to be more accurate than Pathload when a new cross traffic stream with known rate is injected. The principal drawback of these techniques is that they require on the order of hundreds of probe packets, and several to several tens of seconds to obtain the available bandwidth estimate.

This paper develops a new bandwidth estimation technique, 'QuickProbe', that uses 19 probe packets to obtain a conservative estimate of the available bandwidth within a single round trip, and then uses 9-17 further probe packets in each subsequent round trip to refine the estimate. We have compared the QuickProbe and Pathload estimates for hundreds of Internet paths between PlanetLab and other nodes. The paths have measured round-trip times in the range of 20-800 milliseconds, capacities in the range of 0.1 - 600 Mb/s, and ratios of available bandwidth to capacity in the range of 5-95%. Over such paths, 'QuickProbe' obtains conservative available bandwidth estimates after two round trips that are a within a factor of 0.7 - 1.0 times the Pathload estimate.

2. QUICKPROBE TECHNIQUE

Similar to the Pathload and PTR techniques, the QuickProbe method sends a fixed-length train of maximum-size packets with fixed spacing, and determines the sending rate to be feasible (i.e., less than or equal to the available bandwidth) if the receive rate for the train of packets is within 10% of the send rate. QuickProbe uses two initial probe rates (for example, 6 Mb/s and 80 Mb/s), as well as two initial packet pairs to measure the capacity of the bottleneck link, and then uses the measured path capacity and the initial probe rate feasibility results to determine the next probe rate in order to perform a binary search for the maximum feasible transmission rate - i.e., the available bandwidth in the path.

One key difference between QuickProbe and previous methods is that we determine whether a given transmission rate is feasible using a single probe train of length 9 packets (or 17 packets if the probe rate is above 100 Mb/s). This number if probe packets was determined experimentally (over commercial Internet paths as well as high capacity research network paths) to yield results that are nearly as accurate as the Pathload estimates. An intuitive explanation for this is, briefly, as follows. Consider a probe train of 100 packets, as used in Pathload. During the transmission of this probe train, either the bottleneck queue occasionally empties or it doesn't. If it never empties then the probe rate is almost certainly infeasible.[1] In this case, it is highly likely that some queue build up occurs during every sequence of 9 or 17 out of the 100 equally spaced probe packets, and thus a 9-packet probe train would also likely obtain the same infeasibility result. On the other hand, if the queue does occasionally empty, then there is excess idle bandwidth and the feasibility result depends on the last n packets of the train that do not find the bottleneck queue empty. To be infeasible, these n packets need to encounter *excess* cross traffic packet transmissions having total duration equal to 10% of the total time to send the 100 probe packets, instead of excess idle time, which is unlikely. Each sequence of 9 probe packets may arrive when the unused bandwidth is lower than the average available bandwidth, and thus may obtain an "infeasible" result although the longer probe train obtains a "feasible" result; this is more likely as the probe rate approaches the actual available bandwidth, and leads to conservative QuickProbe estimates. These intuitive arguments explain why a single probe train of length 9 or 17 might be long enough to obtain results that are close to as accurate as the 100-packet probe trains in Pathload, but experimental evaluation is needed to determine how close the QuickProbe estimates are to the Pathload estimates.

[1]It is highly unlikely that the sum of the idle periods between cross traffic transmissions is exactly equal to the sum of the 100 probe packet transmission times.

SIGMetrics/Performance'06, June 26–30, 2006, Saint Malo, France.
ACM 1-59593-320-4/06/0006.

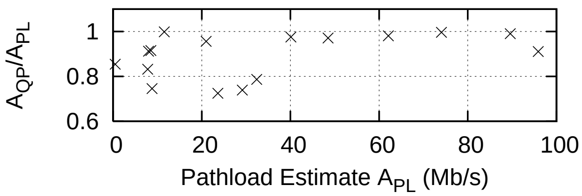

1: Representative $A_{QP,2}/A_{PL}$ for a wide range of Internet paths.

3. RESULTS

Figure 1 shows results for running QuickProbe and Pathload (version 1.2.0) back-to-back on over 16 different Internet paths. These results are representative of a much larger set of experiments over hundreds of paths that we have measured to date. We plot the ratio of the QuickProbe estimate after just two round trips ($A_{QP,2}$) to the estimate given by Pathload (A_{PL}) as a function of A_{PL}.

Five points in the graph are for different commercial Internet paths that have bottleneck link capacities (C) ranging from 10 Mb/s to 80 Mb/s, and available bandwidths ranging from just $0.05C$ to $0.3C$; hence the commercial Internet cross traffic is using 70% to 95% of the bottleneck link capacity. For those paths, $A_{QP,2}/A_{PL}$ is in the range of 0.72 to 0.99.

Three additional paths have capacities ranging from 20 Mb/s to 80 Mb/s and cross traffic that uses 55% to 60% of C; in those cases $A_{QP,2}/A_{PL}$ is between 0.75 and 0.8.

The remaining eight paths have cross traffic that utilizes 5% to 53% of C with C ranging from 20 - 100 Mb/s. These lower cross traffic loads might also fluctuate and thus also be difficult to estimate accurately, but $A_{QP,2}/A_{PL}$ is above 0.9 in all of these cases.

We find that the $A_{QP,2}$ estimates are between 0.7 - 1.0 times the Pathload estimate in *nearly every* experiment that we have run (over hundreds of paths), and are often between 0.9 - 1.0 times the Pathload value. In a very few cases, three round trips are needed to obtain an estimate close to A_{PL}. Note that QuickProbe estimate is always less than or equal to Pathload estimate, indicating that QuickProbe obtains conservative overestimates. In some cases QuickProbe will underestimate the available bandwidth due to the granularity of the binary search during two round trips but the estimate can be improved in further round trips. In other cases the underestimate may be due to hitting an unusual cross traffic burst with the 9 probe packet train as compared with the average cross traffic load that is observed by Pathload's 100-packet trains.

Figure 2 shows the send and receive rates over a path from Bristol, UK to Wisconsin when $A_{QP,2} = 0.2C$. In Figure 2(a) the packets are sent at rate $A_{QP,1} = 6$ Mb/s, for one round trip after 9 probe packets are sent to obtain $A_{QP,2}$. In Figure 2(b) the packets are sent at rate $A_{QP,2} = 20$ Mb/s for 10 seconds after the second round trip. We observe that the 20 Mb/s rate is sustainable during the 10 seconds following the $A_{QP,2}$ estimate - that is, there was no packet loss, the average queuing delay for packets sent during each round trip time was less than the time to transmit 15 maximum size packets on the bottleneck link, and the receive rate in each round trip does not exceed 10% more than the send rate. There are small fluctuations in the receive rate that indicate the presence of fluctuations in the cross-traffic, but the conservative estimate $A_{QP,2}$ is sustainable (without packet loss) for the entire 10 seconds. We note that after any available bandwidth estimation method, the sender needs to monitor packet loss and changes in queuing delay and react to congestion that may be detected by these measures. In a small fraction of our experiments, we observed such changes just 3-5 round trips after obtaining the QuickProbe available bandwidth measure. However, in most experiments the estimated available bandwidth was sustainable over several to tens of seconds.

4. CONCLUSION

We have developed and evaluated QuickProbe, a new rapid available bandwidth estimation technique that significantly extends the state of art by producing accurate estimates in one or two round trip times for a diversity of Internet paths with widely varying network loads.

5. REFERENCES

[1] N. Hu and P. Steenkiste. Evaluation and characterization of available bandwidth probing techniques. *IEEE journal on selected areas in communications*, 21(6):879–894, August 2003.

[2] M. Jain and C. Dovrolis. Pathload: A measurement tool for end-to-end available bandwidth. In *Passive and Active Measurement Workshop*, March 2002.

[3] T. Oetiker. MRTG: Multi Router Traffic Grapher. http://people.ee.ethz.ch/ oetiker/webtools/mrtg/, Feb 2005.

[4] J. Strauss, D. Katabi, and F. Kaashoek. A measurement study of available bandwidth estimation tools. In *3rd ACM SIGCOMM conference on internet measurement*, 2003.

[5] S. Ubik and A. Kral. End-to-end bandwidth estimation tools. Technical Report 25, CESNET, 2003.

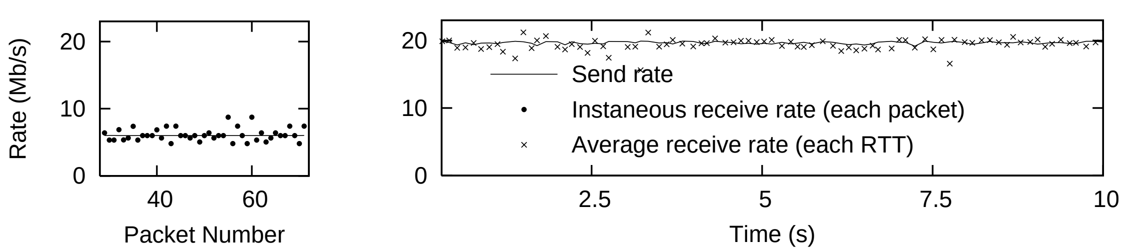

(a) Sending rate $= A_{QP,1} = 6$ Mb/s

(b) Sending rate $= A_{QP,2} = 20$ Mb/s

2: Receive rates when sending at $A_{QP,i}$ (commercial Internet path from Bristol, UK to Wisconsin, C=100 Mb/s, RTT=111 ms, loss rate=0)