Rate-Based and Gap-Based Available Bandwidth Estimation Techniques in Cross-Traffic Context

Wayman Tan[1], Marat Zhanikeev[1], and Yoshiaki Tanaka[1 2]

[1] Global Information and Telecommunication Institute, Waseda University
1-3-10 Nishi-Waseda, Shinjuku-ku, Tokyo, 169-0051 Japan
[2] Advanced Research Institute for Science and Engineering, Waseda University
17 Kikuicho, Shinjuku-ku, Tokyo, 162-0044 Japan
waymantan@akane.waseda.jp, maratishe@aoni.waseda.jp, ytanaka@waseda.jp

Abstract. Recent years have seen an increasing interest in end-to-end available bandwidth estimation. A number of estimation techniques and tools have been developed during the last few years. All of them can be classified into two models: rate-based and gap-based, according to the underlying approaches. The difference in characteristics of both models is studied in this paper. *PathChirp* and *IGI* are selected to represent each model and they are evaluated on low speed paths which are typical in today's local access of the Internet. Finally a hybrid method adopting both rate-based and gap-based approaches is proposed. The hybrid method is compared with *pathChirp* and *IGI*. Simulation proves that the hybrid method yields more accurate results and reduces overhead traffic.

1 Introduction

An end-to-end network path consists of a sequence of store-and-forward links that transfer packets from the source host to the destination host at each end of the path. Two commonly used throughput-related metrics of a path are the *end-to-end capacity* and the *end-to-end available bandwidth*.

In an individual link, the *capacity* is defined as the maximum IP layer transmission rate at that link. Extending the definition to a network path, the *end-to-end capacity* of a path is the maximum IP layer rate that the path can transfer from the source to the destination. It is determined by the smallest link capacity in the path.

The *available bandwidth* of a link is the unused portion of its capacity. Since a link is either idle or transmitting a packet at full capacity at any given instant of time, the instantaneous utilization of the link can only be either 0 or 1. Thus reasonable definition of available bandwidth should be the average unused capacity over some time interval. The available bandwidth A over a time period $(t - \tau, t)$ is

$$A(t-\tau,t)=\frac{1}{\tau}\int_{t-\tau}^{t} C(1-u(x))dx\,, \qquad (1)$$

where C is the capacity of the link and $u(x)$ is the instantaneous utilization of the link at time x. Similarly, the *end-to-end available bandwidth* is determined by the smallest link available bandwidth in the path over certain interval.

End-to-end available bandwidth is usually estimated with active probing techniques which send probe traffic from the source host to the destination host of the path. Active probing is free of privileged access requirement and feasible for the end users. Various active probing techniques have been proposed today. Generally, they can be classified into two models according to the underlying approaches:

The *rate-based model* (RM) uses the sending rate of the probe traffic at the sender (probing rate) to infer end-to-end available bandwidth. The RM techniques usually send probe packets from an initial low rate and increase the rate gradually. They search for *the turning point* at which the arrival rate of probe packets at the receiver strata to be lower than the probing rate. Such turning point is believed to mirror the end-to-end available bandwidth over the probing period. RM techniques include *pathChirp* [1], *Pathload* [2] and *PTR* [3].

The *gap-based model* (GM) compares the time gap of successive probe packets between the sender and the receiver to calculate the available bandwidth. Given two successive probe packets fall in the same queuing period at the single bottleneck of the path, and they are sent with a time gap Δ_{in} and reach the receiver with a time gap Δ_{out}, then Δ_{out} should consist of two segments: Δ_{in} and the transmission time of the cross-traffic between them. The GM techniques use the difference between Δ_{out} and Δ_{in} and end-to-end capacity to calculate the rate of cross-traffic and then end-to-end available bandwidth. The calculation of GM tools is usually based on a set of congested probe packets. *IGI* [3] and *Spruce* [4] are examples of tools using GM.

2 Comparison between Rate-Based and Gap-Based Available Bandwidth Estimation Tools on Typical Internet Paths

2. 1 Representative Tools of Two Models

In this section two representative tools of each class are selected to be compared. In the RM class *pathChirp* [1] is famous for its *chirp* structure – a series of exponentially spaced packets within a single probe train. The chirp structure covers a wide range of probing rate and thus works with high efficiency. The probing process of *pathChirp* is fast and with significantly light overhead traffic. *IGI* [3] is a tool in the GM class. *IGI* sends an even spaced train with carefully selected initial gap and increases the gap linearly for subsequent trains. *IGI* estimates available bandwidth when all the gaps of train at the sender equal to those at the receiver. *PathChirp* and *IGI* are selected to represent their class in our study.

2. 2 Test Environment and Methodology

Some evaluation of the existing end-to-end available bandwidth estimation tools has been done on high speed links (with capacity of 1000 Mbps) environment [5]. However, high speed end-to-end path does not yet prevail in today's Internet; most paths are with a capacity around 10 Mbps, usually constrained by the edge of the network. For practical purpose, we evaluate and compare the tools on paths with capacity of 10 Mbps. The test is performed based on simulations with OPNET Modeler. Simulation environment is crucial for this research as we need complete control over cross-traffic load on network paths and the ability to monitor link utilization with high level precision and granularity, which is not possible in real network environment. The network paths under observation are from 5 to 7 hops long with a single bottleneck in the middle of the path whose capacity is 10 Mbps. Admittedly, this topology is primitive and some special artifacts such as multiple bottlenecks are beyond the scope of our consideration, but the simple topology is favorable for discovering the fundamental characteristics of the tools.

Most available bandwidth estimation tools work well in stable cross-traffic environment. When they face highly bursty cross-traffic, however, the performance is very different for each tool. Therefore the performance in bursty cross-traffic environment is critical for evaluation. Except the idle path and light traffic scenarios, most of our test scenarios simulate highly bursty traffic environment. The

cross-traffic is generated in the form of various popular Internet applications such as HTTP web browsing, FTP file transferring, Email, video conferencing and so on.

Both *pathChirp* and *IGI* have tunable parameters. Generally the default values of those parameters are used in our test. The initial gap of *IGI* is exceptional, because some test scenarios with high utilization are performed and the default initial gap is too large. For pathChirp, the probe packet size is 1000 bytes, the exponential spread factor is 1.3 and the length of the chirp is 15; for *IGI*, the probe packet size is 500 bytes, the initial gap is 0.5 ms and the length of a train is 32. This setting remains unchanged for all the comparison scenarios. Two tools run subsequently on the same path with the same cross-traffic condition repeated by the simulator.

2.3 Comparison Results

The first test scenario is performed on an idle path. The results are shown in Fig. 1. Both tools give an estimate around 10 Mbps, which is the end-to-end capacity of the path. While the estimate of *IGI* is very close to 10 Mbps, it is interesting to see that the estimate of *pathChirp* goes up to 11 Mbps. When *pathChirp* is evaluated on high speed links [5] the same phenomenon is also seen. Please note that the actual available bandwidth (avail-bw for short in the figures) is not 10 Mbps in both *pathChirp* and *IGI* cases. It is because the probe traffic occupies a little bandwidth of the path.

In the second scenario, there is light cross-traffic with average utilization at about 15%. The results of both tools, which are also shown in Fig. 1, are exactly the same as those on an idle path. It could be understood, because 15% of 10 Mbps is not more than the probe traffic rate of both tools. It means the cross-traffic is extremely light. Actually in a path with capacity of 10 Mbps both tools can not sense cross-traffic when the average utilization is lower than 15%.

Fig. 2 shows the results when the average utilization is 30%. Due to the bursty cross-traffic, the instantaneous utilization is heavily fluctuating. The accuracy of both tools drops in this scenario. However, the estimate results of both tools show different characteristics. The estimate of *pathChirp* exhibits highly dynamic feature which reflects the fluctuation of the cross-traffic. A number of estimate results correctly follow the change of cross-traffic although some go opposite direction. But *pathChirp* is most of time over-reacting: even it follows the change of cross-traffic but the estimate value is either too high when the available bandwidth increases, or too low when it decreases. Some estimate results are higher than 10 Mbps which is obviously over the border. On the other hand, *IGI* barely

corresponds to the change tendency of the cross-traffic. The range of the estimate results is limited. In this scenario, they are almost always higher than actual available bandwidth and fall into the area between actual available bandwidth and end-to-end capacity.

There is more cross-traffic in the next scenario and the average utilization is 50%. Fig. 3 shows the results in this scenario. Again *pathChirp* shows its strong ability to follow the change of cross-traffic but is over-reacting. The performance of *IGI* is also unchanged.

The final scenario is with heavy cross-traffic and the average utilization jumps up to 70%. Fig.4 shows the results in this scenario. Unsurprisingly, the particular characteristics of both tools exist as the same. But, *pathChirp* gives less estimates lower than the actual available bandwidth compared with previous scenarios. And, there are occasionally extremely high estimate results up to 30 Mbps from *pathChirp*.

In the above test scenarios, *pathChirp* and *IGI* exhibit different characteristics: *pathChirp* is good at following changes of cross-traffic but often over-reacts, while *IGI* does not catch changes of cross-traffic accordingly but offers a stable estimate. The different characteristics probably result from the underlying approaches of two models. The RM tools perform estimation based on a single turning point. For *pathChirp,* the turning point is a single packet pair inside the chirp. Therefore, it is fast to discover the change of cross-traffic but easy to deviate from the actual value. The *GM* tools estimate available bandwidth based on a set of congested probe packets. The estimate is smoother but can not follow changes in cross-traffic promptly.

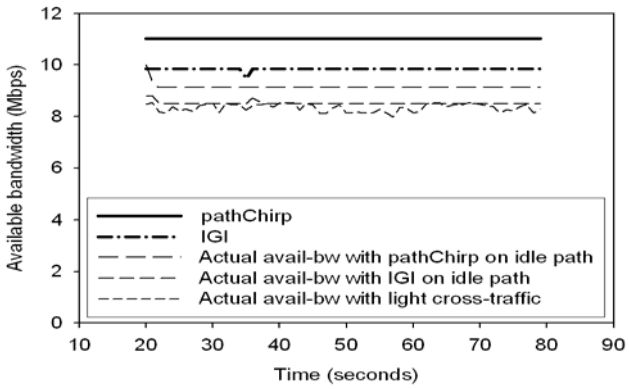

Fig. 1. Comparison between *pathChirp* and *IGI* without cross-traffic and with light cross-traffic.

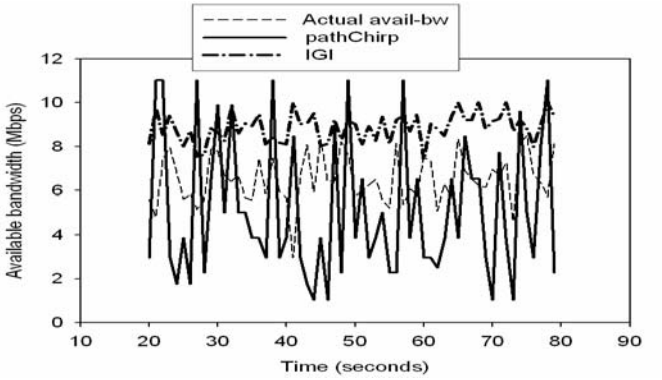

Fig. 2. Comparison between *pathChirp* and *IGI* with average utilization at 30%.

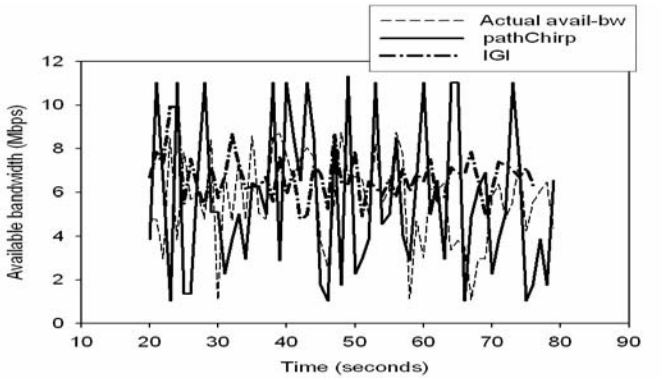

Fig. 3. Comparison between *pathChirp* and *IGI* with average utilization at 50%.

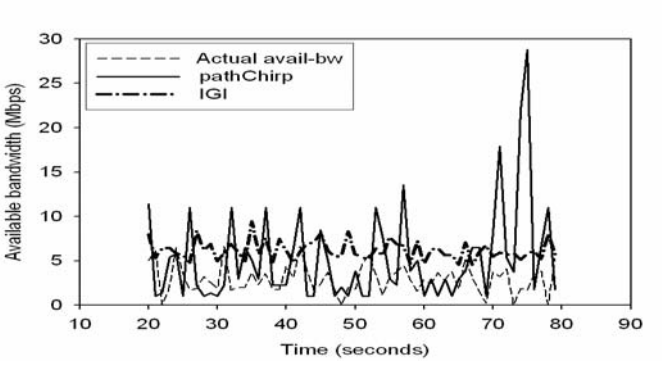

Fig. 4. Comparison between *pathChirp* and *IGI* with average utilization at 70%.

3 Hybrid Method Adopting Rate-Based and Gap-Based Approaches

3.1 Introduction to Hybrid Method

In our experiment, we explore the different characteristics in performance of rate-based and gap-based available bandwidth estimation tools. In fact, we believe that their performance characteristics are complementary. Therefore, a tool that combines both models should have the potential to improve the accuracy of the estimate. A hybrid method adopting both *pathChirp* and *IGI* logic is proposed based on this assumption.

There are two phases in the method: the chirp processing and the *IGI* processing. The first phase is the chirp processing. It uses the same logic of *pathChirp* with certain enhancement: The initial rate of the chirp is set to a small portion of end-to-end capacity. It is set to cover a wide range of possible results of available bandwidth. The spread factor is no longer user specified. It is automatically set to ensure the highest probing rate of the chirp can be slightly higher than the end-to-end capacity. Self-induced congestion is guaranteed in this way. In this phase, the sender sends a chirp to the receiver. And after it receives the chirp, the receiver performs estimate with the logic of *pathChirp*.

The second phase is the *IGI* processing with the logic of *IGI*. The initial gap of the probe train is set according to the feedback of the chirp processing, instead of using default value. This is a significant enhancement in that it can effectively reduce the number of probe trains for the *IGI* processing. The initial gap is set based on the probing rate of the turning point in the chirp processing. Since the probing rate at the turning point is expected to reflect the available bandwidth, the initial gap based on this probing rate can converge faster. Similarly the sender sends probe trains and the receiver receives them and performs calculation.

After two probing phases are finished, two available bandwidth estimate results are generated. The *IGI* estimate then acts as an anchor for the chirp estimate to prevent the chirp estimate deviation. In addition, the chirp estimate is cut to be equal to end-to-end capacity when it is larger than it. A coefficient α is used to control the weight of the estimate of two phases in the final estimate. The final estimate E is given by

$$E = \alpha E_{\text{chirp}} + (1 - \alpha) E_{\text{igi}} , \qquad (2)$$

where $0 \le \alpha \le 1$, E_{chirp} is the chirp estimate and E_{igi} is the *IGI* estimate.

3.2 Evaluation of Hybrid Method

We evaluate the hybrid method by comparing it with *pathChirp* and *IGI*. The comparison is performed in the same environment and with the same methodology as previous tests.

First, the hybrid method is compared with *pathChirp* and *IGI* when the average utilization of the path is 30%. The results are shown in Fig. 5. The coefficient α is 0.5 in this case. The hybrid method yields much closer results to the actual available bandwidth than the other two tools. It can not only follow changes of cross-traffic but also effectively restrain the estimates from overreaction. In addition, the hybrid method reduces probe traffic in the second phase compared with the original *IGI*. The average number of trains of *IGI* process in the hybrid method is 1.7 while the number is 6.2 for original *IGI*.

Fig. 6 shows the comparison when the average utilization is 50%. The coefficient α is also 0.5. Again the hybrid method performs the best. The average number of *IGI* trains is 3.5 and that of original *IGI* is 9.1.

The final comparison is performed when the average utilization is 70%. The results are shown in Fig. 7. In this scenario, the coefficient α is 0.7. The hybrid method is still the best among the three. The average number of *IGI* trains is 3.7 and that of original *IGI* is 7.6.

We set the coefficient α with different values to see how it affects the final estimate. The error rate is used for verification. The error rate is the difference between the estimate and the actual available bandwidth value in proportion to the total capacity. A single error rate value is calculated from a number of estimate

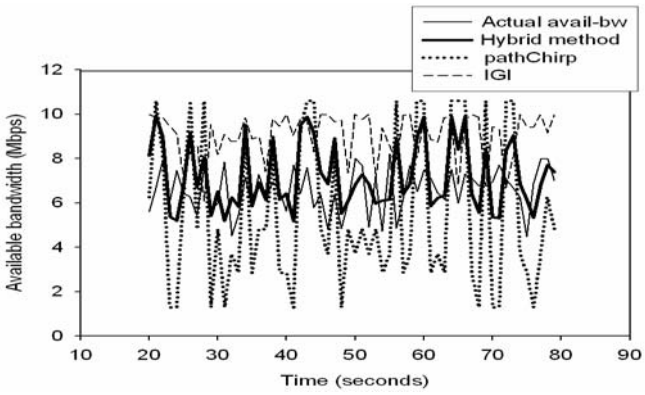

Fig. 5. Comparison among hybrid method, *pathChirp*, and *IGI* with average utilization at 30%.

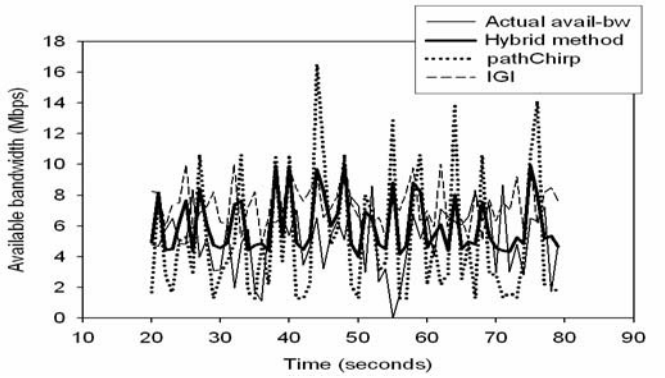

Fig. 6. Comparison among hybrid method, *pathChirp*, and *IGI* with average utilization at 50%.

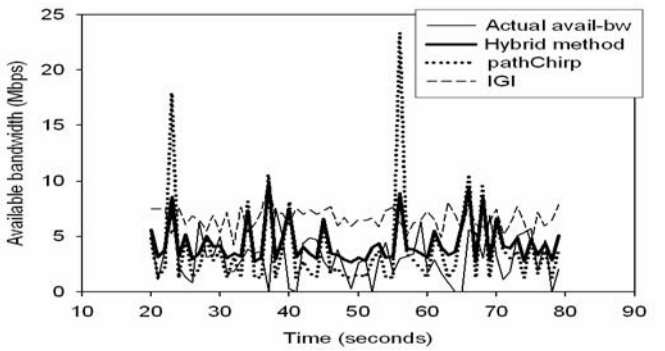

Fig. 7. Comparison among hybrid method, *pathChirp*, and *IGI* with average utilization at 70%.

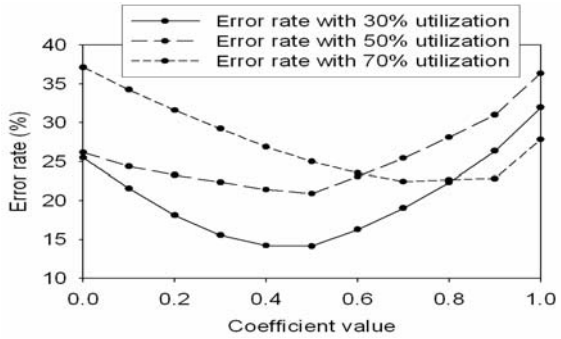

Fig. 8. Error rate vs. coefficient value with different utilization.

samples of consecutive probes. The results are shown in Fig. 8. When the load is not heavy (e.g., the average utilization is under 50%), the value of 0.5 yields the best

result. When the load is very heavy, the value of α should also increase because the *IGI* estimates are generally high and they should be restricted to a light-weight. Generally speaking, a network path with normal usage should not see very heavy utilization, so 0.5 is appropriate for the coefficient value. When heavy utilization happens, however, we can refer to the estimates of IGI as it is stable. For example, we can assume the utilization is heavy when the average estimate of IGI is no more than 6 Mbps in our test. This assumption is based on our empirical results.

4 Conclusion

In this paper, the differences in performance between *rate-based model* and *gap-based model* of end-to-end available bandwidth estimation techniques are studied. *PathChirp* and *IGI* are selected to represent each model. The evaluation is performed on low speed paths which are typical in today's Internet. A hybrid method adopting both *pathChirp* and *IGI* logic is proposed. It takes advantage of the complementary properties of the two models Simulation shows that the proposed hybrid method significantly improves the estimation accuracy and reduces overhead traffic compared with original tools. Our future work is to evaluate the hybrid method in real world environment.

References

1. Ribeiro, V. J., Riedi, R. H., Baraniuk, R. G., Navratil, J., and Cottrell, L.: pathChirp: Efficient Available Bandwidth Estimation for Network Paths. Proc. Passive and Active Measurement Workshop, Tech. Report no. SLAC-PUB-9732 (2003)
2. Jain, M., and Dovrolis, C.: End-to-End Available Bandwidth: Measurement Methodology, Dynamics, and Relation with TCP Throughput. IEEE/ACM Trans. Networking, Vol.11, No.4 (2003) 537-549
3. Hu, N. and Steenkiste, P.: Evaluation and Characterization of Available Bandwidth Probing Techniques. IEEE J. Select. Areas Commun., Vol.21, No.6 (2003) 879-894
4. Strauss, J., Katabi, D., and Kaashoek, F.: A Measurement Study of Available Bandwidth Estimation Tools. Proc. 3rd ACM SIGCOMM Conf. Internet Measurement (2003) 39-44
5. Shriram, A., Murray, M., Hyun, Y., Brownlee, N., Broido, A., Fomenkov, M., and Claffy, K.: Comparison of Public End-to-End Bandwidth Estimation Tools on High-Speed Links. Proc. Passive and Active Network Measurement, 6th International Workshop, Boston, MA, USA (2005)