

Evaluating the Potential of Bandwidth Estimators

¹Xiliang Liu, ¹Kaliappa Ravindran, and ²Dmitri Loguinov

¹The City University of New York, New York, NY 10016

xliu@gc.cuny.edu, ravi@cs.ccny.cuny.edu

²Texas A&M University, College Station, TX 77843

dmitri@cs.tamu.edu

Abstract—Many packet probing based available bandwidth estimation techniques have been proposed recently. However, their performance evaluation procedures remain in a coarse-granular fashion, which provide little information about the sources of performance inadequacy. We initiate the effort towards fine-granular performance metrics and evaluation procedures that help identify the causes of measurement errors, that can be used to retroactively optimize the design of measurement methods. As the first step to this end, we propose a metric called *single-hop potential*, to assess the capability of bandwidth estimation algorithms to reach the measurement target in an asymptotically accurate manner. We propose off-line methodology to evaluate the single-hop potential of various bandwidth estimators and show how this procedure can help improve the design of measurement techniques.

I. INTRODUCTION

End-to-end estimation of the spare capacity along a network path has recently become an important Internet measurement research area. Current measurement techniques such as TOPP [5], Pathload [2], IGI/PTR [1], Pathchirp [7], and Spruce [8] are all based on packet probing, where measurement packet streams are sent to the network path, and the inter-packet delays of the streams when arriving at the other end are used by some inference algorithm to estimate the path available bandwidth. In this abstract, we also call the inference algorithms *bandwidth estimators*.

The devise of the techniques has to take into consideration a long list of factors that may affect its measurement accuracy. Practical issues include time measurement errors, OS context switch compensation, layer two effects, etc. Algorithmically, it is strongly desirable that in real Internet environments, the estimator has the capability to reach the measurement targets, and is robust to all kinds of noises. However, most of the measurement techniques to date relied on constant-rate fluid cross-traffic to justify the design of their bandwidth estimators. Recent progress [3] starts showing that a good estimator in fluid cross-traffic can actually perform quite poorly in bursty cross-traffic. Hence, more understanding about the algorithmic performance of bandwidth estimators in realistic cross-traffic conditions is needed.

On the other hand, the current performance evaluation procedures for available bandwidth techniques is very coarse-granular. The common practice is to conduct Internet measurement and compare the results to router MRTG data or TCP throughput. The downside of coarse-granular evaluation is that it does not help much to identify the exact reasons for measurement inaccuracies, thus provides little guidance towards future

improvements. Furthermore, due to the dynamics of available bandwidth, the results of coarse-granular evaluation are less reproducible and can even become conflicting. Therefore, we argue in favor for fine-granular performance metrics and evaluation procedures. In stead of using a monolithic accuracy metric, we advocate the identification of the major contributing factors to the final accuracy and the proposal of individual metrics for each factors. As an initial effort toward this end, we first disregard all practical factors and solely focus on the performance evaluation of bandwidth estimators. We believe the isolation of algorithmic performance from practical implementation issues is both feasible and advantageous. Note that simulation-based evaluation can completely remove all effects of practical issues and purely focus the investigation on the algorithmic core of the measurement techniques. On the other hand, the solutions to practical issues are universally applicable to all measurement techniques and can be pursued independent of the design of bandwidth estimators [6].

We propose two algorithmic metrics, single-hop potential and multi-hop robustness, to evaluate the performance of bandwidth estimators. We use the first metric to assess the single-hop asymptotic accuracy; we use the second metric to assess how much this accuracy is subject to non-bottleneck link interference. We report preliminary results regarding the first metric. The evaluation of the second metric is ongoing and will be reported in the future.

In section II, we provide more details about this two metrics. In section III, we propose evaluation procedures to assess the single-hop potential of various bandwidth estimators. In section IV, we provide preliminary results and discuss how these results can help with the design of bandwidth estimators. In section V, we briefly address our future work along this direction and conclude.

II. PERFORMANCE METRICS FOR BANDWIDTH ESTIMATORS

Before presenting the performance metrics, we examine several issues in more detail. We first define available bandwidth more formally to clarify the measurement targets of bandwidth estimators.

A. Measurement Targets

The available bandwidth of a network hop is its *residual* capacity after transmitting cross-traffic. Since at any time instance, the hop is either idle or transmitting packets at its

capacity speed C , the available bandwidth $B_\delta(t)$ is the average unutilized bandwidth over the time interval $[t, t + \delta]$. Note that $B_\delta(t)$ is a dynamic metric. It varies over time t as well as a wide range of observation interval δ . However, any individual samples of $B_\delta(t)$ at a fixed time instance is unlikely to be interesting to the applications. The long-term time average of $B_\delta(t)$ (presumably in existence), denoted by $A = \overline{B_\delta(t)}$, on the other hand, is usually more relevant to applications and more amenable for measurement. Hence, It is natural to set $A = \overline{B_\delta(t)}$ as the measurement target.

Similarly, we can define the cross-traffic intensity $Y_\delta(t)$ and its long-term time average $\lambda = \overline{Y_\delta(t)}$. Assuming $\lambda \leq C$, the relationship between λ and A is $A = C - \lambda$. This shows that λ is an equally effective measurement target given C known. Note that most current estimation techniques also set A or λ as their measurement goal, sometimes in an implicit way.

B. Bandwidth Estimators

Every measurement techniques have a bandwidth estimator as its core. Some estimators have mathematical expressions ; while others can only be described using computer algorithms. Some estimators are inherently iterative, meaning that they need to first decide the ‘‘proper’’ input probing construction (usually the rate) before conducting measurements with that input construction. Others are non-iterative, their input probing constructions are predetermined. Table I shows a categorization of bandwidth estimators.

TABLE I
CATEGORIZATION OF BANDWIDTH ESTIMATORS

	Mathematically Expressed	Algorithmically Expressed
Iterative	IGI/PTR	Pathload, TOPP
Non-Iterative	Spruce	Pathchirp

Due to the nature of the measurement targets (the fact that they are time averages), one thing in common for all bandwidth estimators is that they usually have to conduct multiple runs of measurements and use the average of those results to approach A or λ .

C. Performance Metrics

Single-hop potential is a metric assessing the asymptotic performance of an available bandwidth estimator in single-hop environments. That is, as the number of measurement runs becomes large, what value will the average of those estimations approach to? An ideal estimator should be unbiased, meaning that the asymptotic average of large number of estimations should be able to reach the measurement targets A or λ .

The significance of this metric is obvious. If the estimator is inherently biased and does not have the capability to catch the measurement target in a single-hop general cross-traffic environments, it can naturally be one of the contributing factors of measurement errors. An evaluation of single-hop potential for available-bandwidth estimators helps to identify, mitigate, compensate for, or even eliminate this type of errors.

The reason why we do not propose a metric called ‘‘multi-hop potential’’ is as follows. First, there is no theoretical

foundation supporting such a metrics to date. Second, we expect much more involved evaluation procedures for a metric like ‘‘multi-hop potential’’. Hence, we propose another metric called ‘‘multi-hop robustness’’ to capture how sensitive the single-hop potential is to the cross-traffic interference from non-bottleneck links. We argue that a metric like ‘‘multi-hop potential’’ can be viewed as a combination of the these two metrics. In other words, multi-hop potential= single-hop potential+ multi-hop robustness.

The evaluation and more details of second metric will be reported in future work. In this abstract, we focus on the evaluation of the first metric only.

III. EVALUATION PROCEDURE

We devise an off-line procedure to evaluate the single-hop potential of given estimators based on cross-traffic packet arrival trace and the hop capacity C . The trace file provides information regarding the arrival time and packet size for every cross-traffic packet.

We now introduce several notations to describe our evaluation procedure. We use \vec{x} to denote the input probing packet stream construction. This includes all relevant information regarding the packet stream such as the inter-packet spacings, the packet sizes, the number of packets in the stream, etc. We use \vec{y} to denote the output probing packet stream construction after \vec{x} pass through the single-hop path. We use Υ_0^τ to denote a cross-traffic trace in the time interval $[0, \tau]$. Note that \vec{y} is a deterministic function of \vec{x} , the arrival time t to the hop of the first packet in the stream \vec{x} , cross-traffic trace Υ_0^τ , and hop capacity C . We denote this function as:

$$\vec{y} = \gamma(\vec{x}, \Upsilon_0^\tau, C, t). \quad (1)$$

The bandwidth estimator can be viewed as another function that takes \vec{x} and \vec{y} as input and generate an estimation of A associated with time instance t . We denote the estimator as:

$$A(t) = \varepsilon(\vec{x}, \vec{y}) = \varepsilon(\vec{x}, \gamma(\vec{x}, \Upsilon_0^\tau, C, t)). \quad (2)$$

The spirit of our single-hop potential evaluation procedure is to compute $\overline{A(t)}_0^\tau$, the time average of $A(t)$ in a finite time-interval $[0, \tau]$, and examine whether this metric decently approximates the measurement target A or not. A decent approximation indicates good single-hop potential of the bandwidth estimator $\varepsilon(\vec{x}, \vec{y})$.

The validity of our evaluation procedure is justified as follows. Given an unbiased estimator $\varepsilon(\vec{x}, \vec{y})$ and assuming ASTA (Arrivals See Time Average) property¹ [4], the limiting time average $\overline{A(t)}_0^\infty$ equals to A . In real testing, we can not have an infinite long cross-traffic trace. Thus, we use $\overline{A(t)}_0^\tau$ to approximate $\overline{A(t)}_0^\infty$. As long as we choose a sufficient large τ ,

¹Note that in real bandwidth measurement, only a set of samples of $A(t)$ at discrete time instances can be collected. ASTA assumption implies the equality of the asymptotic average of those samples to the long-term time average of $A(t)$. By assuming ASTA, our evaluation does not capture the measurement error introduced by non-ASTA inter-probing pattern. Using ns-2 simulation to evaluate signal-hop potential will not be able to identify whether the measurement error is due to ASTA bias or due to the problems of bandwidth estimators.

this should serve a good approximation. The way of choosing τ is to make sure that $\overline{B_\delta(t)}_0^\tau$ is close enough to A or the average cross-traffic intensity in $[0, \tau]$ is close enough to $C - A$.

Again note that the computation of $\overline{A(t)}_0^\tau$ is again approximate. We sample $A(t)_0^\tau$ at a set of equally spaced time instances and use the average of those samples to approximate $\overline{A(t)}_0^\tau$. We make sure the number of samples is large enough so that when it is doubled, the result is of little difference.

Note that the function $\gamma(\vec{x}, \Upsilon_0^T, C, t)$ is deterministic and can be easily computed algorithmically. We omit the technical details but remind readers of the fact that the evaluation procedure involves intensive computation but no event-based simulation. The estimator $\varepsilon(\vec{x}, \vec{y})$ is essentially the algorithmic core of each measurement techniques. For iterative measurement techniques such as TOPP, Pathload, and Pathchirp, we can compute the single-hop potential at a range of input packet stream constructions and examine whether the construction used by the measurement techniques has good single-hop potential or not. This method separates the estimation part from the iteration part in the measurement techniques and make it possible to apply our evaluation procedure to both iterative and non-iterative techniques.

IV. PRELIMINARY RESULTS

We already applied our evaluation procedure to several bandwidth estimators such as TOPP, Spruce, IGI/PTR using several different cross-traffic traces. Our evaluation generate quite a few interesting results and bring new knowledge on the behavior of those bandwidth estimators. We are currently working on the evaluation of Pathload and Pathchirp estimators.

In this abstract, we only report a small portion of the results we obtained so far. We use a Poisson random number generator to generate cross-traffic trace with exponentially distributed inter-packet arrivals. The average arrival rate is 500 packet per second. The packet sizes (in bytes) are decided by a random variable uniformly distributed in $[1, 1500]$. Hence, the long term cross-traffic intensity λ is 3Mbps. The hop capacity is set to be 10Mbps. Thus, the available bandwidth is $A = C - \lambda = 7$ Mbps. We generated 100-second cross-traffic traces, but we only used the traces of the first 20 second since the cross-traffic intensity in the first 20 second is already extremely close to 3Mbps. We approximate $\overline{A(t)}_0^{20}$ by computing the values of $A(t)$ at 1000 equally spaced time instances and taking the average of those 1000 samples. We find doubling the number of samples produces a difference less than 0.1% for all bandwidth estimators.

A. Spruce Estimator

Spruce sends probing packet-pairs of 1500B packet size with intra-pair gap g_I set to the hop transmission delay of the packet s/C . Each probing pair generates an available bandwidth estimate A_i based on the output intra-pair gap g_O using the following bandwidth estimator:

$$A_i = C \left(1 - \frac{g_O - g_I}{g_I} \right). \quad (3)$$

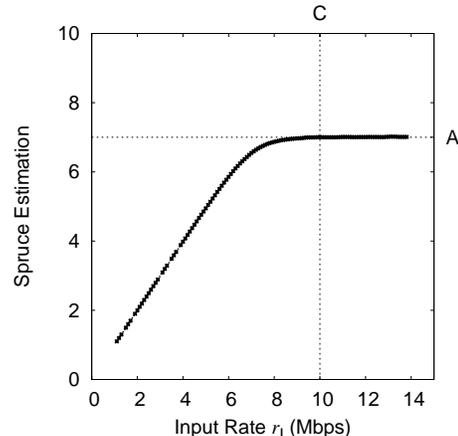


Fig. 1. Single-hop potential evaluation of Spruce Estimator.

In our evaluation, we find the average of the 1000 samples obtained using the above estimator produced unbiased estimation for $A = 7$ Mbps. We also tried to input probing packet-pairs with different input gap g_I and evaluate the single-hop potential with those packet-pair constructions. Figure 1 shows the average of the 1000 estimations with respect to the input rate $r_I = s/g_I$. It shows that when $r_I \geq 8$ Mbps, Spruce estimator is almost unbiased. When $r_I < 8$ Mbps Spruce estimator is negatively biased, producing underestimation of A . The amount of underestimation decreases linearly as r_I increases towards A . This suggests that Spruce in fact can reduce its probing rate with little compromise of the measurement accuracy.

B. PTR Estimator

PTR sends probing packet-trains to the hop and uses the output rate of the probing train as an estimator of A :

$$A_i = \frac{(n-1)s}{g_O}. \quad (4)$$

where n is the number of packets in the train, s is the packet size, and g_O is the output spacing between the first and last packet after the whole train leaves the hop.

PTR estimator is used at the point where output rate starts to fall below the input rate (called the turning point). We compute the average output gaps of 1000 samples and use (4) to estimate A at different input probing rates. Figure 2 shows the estimation with respect to the input probing rate. It shows that PTR is actually a negatively biased estimator at the turning point. We also find that, as the probing packet size or the packet train length decreases, the location of turning point becomes further and further away from available bandwidth, and the amount of bias increases. That suggests that in order to produce accurate estimation, PTR has to use long probing train and large packet size. It also reveals a fundamental tradeoff between measurement accuracy and transient congestion in PTR-based bandwidth estimation, which is not available in Spruce.

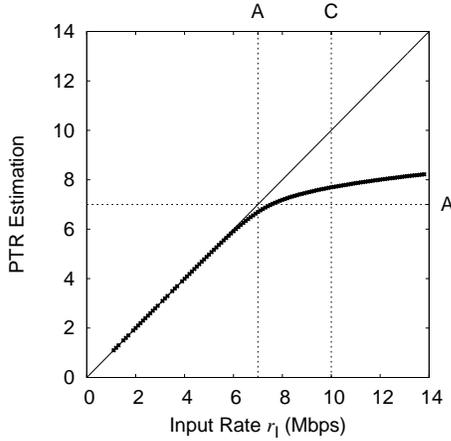


Fig. 2. Single-hop potential evaluation of PTR Estimator.

C. IGI Estimator

IGI also sends probing packet-trains to the hop and uses the following formula as an estimator of A :

$$A = C - \frac{\sum_{y_i > g_I} (y_i - \frac{s}{C})}{\sum_{i=1}^{n-1} y_i}. \quad (5)$$

where y_i $i = 1, 2, \dots, n-1$ are the $n-1$ inter-packet spacings of the output probing trains. g_I are the input inter-packet spacing. Similar to PTR estimator, IGI is also used at the turning point. We compute the IGI estimations at different input probing rate points. Figure 3 shows the estimation with respect to the input probing rate. It shows that IGI produce good estimation at the available bandwidth input rate point, not the turning point. Again note that when the input rate is small, IGI is not a converging estimator and the estimation results are unstable. In this testing, we used packet train of 64 packets as suggested by [1]. When shorter packet trains are used, the turning point will be associated with a smaller rate and the stability problem of IGI estimator becomes more apparent. This suggests that in order to produce stable and accurate results, IGI also has to use long packet train and large probing packet size. Finally, notice that when the input rate is larger than A , IGI becomes stable but produces positive biased estimation of A , which overestimates available bandwidth.

V. CONCLUSION

We proposed a metric called single-hop potential and its associated evaluation procedure to test the asymptotic behavior of various bandwidth estimators. This allows retroactively parameter tuning to optimize the bandwidth estimation techniques. We expect that it is especially useful for estimation techniques like Pathload and Pathchirp. Those techniques involves a lot of heuristic parameters. Our evaluation procedure can be used to find out the single-hop potential with various heuristic parameter settings and choose the one that provides the best single-hop potential, which is our work currently under way. On the other hand, note that the single-hop potential also depends on the cross-traffic nature. By comparing the testing results using different cross-traffic types, we might

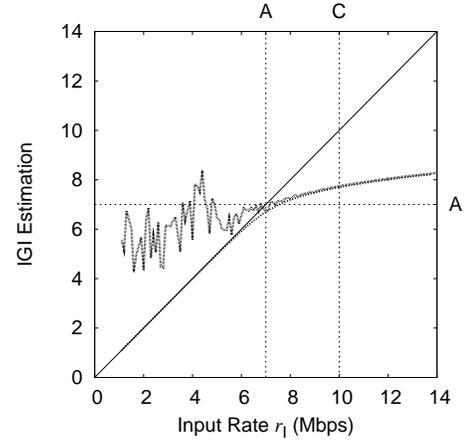


Fig. 3. Single-hop potential evaluation of IGI Estimator.

be able to confirm our conjecture that heuristic parameters in Pathload and Pathchirp do not have an optimal settings for arbitrary cross-traffic. This will potentially reveal some fundamental limitations of the current bandwidth estimation techniques.

REFERENCES

- [1] N. Hu and P. Steenkiste, "Evaluation and Characterization of Available Bandwidth Probing Techniques," *IEEE JSAC*, August 2003.
- [2] M. Jain and C. Dovrolis, "End-to-end available bandwidth: measurement methodology, dynamics, and relation with TCP throughput," *ACM SIGCOMM*, August 2002.
- [3] X. Liu, K. Ravindran, B. Liu, and D. Loguinov, "Single-Hop Probing Asymptotics in Available Bandwidth Estimation: A Sample-Path Analysis," *ACM IMC*, October 2004.
- [4] B. Melamed and D. Yao, "The ASTA property," *Advances in Queueing: Theory, Methods and Open Problems*, pages 195–224, 1995.
- [5] B. Melander, M. Bjorkman, and P. Gunningberg, "Regression-Based Available Bandwidth Measurements," *SPECTS*, July 2002.
- [6] M. Pucci, A. McIntosh, J. Alberi, and T. Raleigh, "Overcoming Precision Limitations in Adaptive Bandwidth Measurements," *NYMAN*, New York, 2003.
- [7] 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.
- [8] J. Strauss, D. Katabi, and F. Kaashoek, "A measurement study of available bandwidth estimation tools," *ACM IMC*, 2003.