

Operational Comparison of Available Bandwidth Estimation Tools

Guillaume Urvoy-Keller, Taoufik En-Najjary, and Alessandro Sorniotti
Institut Eurecom
Sophia-Antipolis, France
{urvoy,ennajjar,sorniott}@eurecom.fr

This article is an editorial note submitted to CCR. It has NOT been peer reviewed. Authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

The available bandwidth of a path directly impacts the performance of throughput sensitive applications, e.g., p2p content replication or podcasting. Several tools have been devised to estimate the available bandwidth. The vast majority of these tools follow either the Probe Rate Model (PRM) or the Probe Gap Model (PGM).

Lao et al. [6] and Liu et al. [7] have identified biases in the PGM approach that lead to consistent underestimations of the available bandwidth. Those results were obtained under the ideal assumption of stationary cross traffic.

In this note, we confirm the existence of these biases experimentally, i.e., for the case of non stationary cross traffic. To do so, we compare one representative of the PRM family, namely Pathload, and one representative of the PGM family, namely Spruce, using long term (several day long) traces collected on an example path.

We first propose a methodology to compare operational results of two available bandwidth measurement tools. Based on the sanitized data obtained using the previous methodology, we next show that the biases identified by previous works are clearly observable on the long term, even with non stationary cross traffic. We further uncover the formal link that exists between the work by Liu et al. and the one by Lao et al.

Categories and Subject Descriptors

H.4 [Network Operations]: Network Monitoring

General Terms

Measurement, Experimentation, Performance

Keywords

Available bandwidth

1. INTRODUCTION

A question of great interest to applications is how much bandwidth is available to them along an end-to-end Internet path. The high variability of the available bandwidth over a wide range of timescales makes the design of measurement algorithms very challenging.

Several tools have been proposed to estimate the available bandwidth [9]. The vast majority of these tools follow either the Probe Rate Model (PRM) or the Probe Gap Model

(PGM). In the Probe Rate Model, a tool modulates its sending rate as a function of the dispersion of packets observed at the receiver. The highest possible rate for which dispersion is minimum is used as an estimate of the available bandwidth. Tools based on the Probe Gap Model (PGM) inject pairs or trains of packets at a rate equal to the capacity of the narrow link (the link with the minimum capacity along a path, see Figure 1). Dispersion of the trains or pairs of packets at the receiver side is used to infer the rate of the cross traffic at the narrow link. The difference between the cross traffic and the capacity of the narrow link is used as an estimate of the available bandwidth of that path. This holds only if the narrow link of a path is also the tight link (the link with the minimum available bandwidth along a path, see Figure 1), which is assumed to be the case in the Probe Gap Model.




Figure 1: Example of an Internet path (the italic indicates the used bandwidth): Link 1 is the narrow link: The link with the smallest capacity (100Mbps). Link 3 is the tight link: The link with the smallest available bandwidth (10 Mbps).

A recent work by Lao et al. [7] has shed light on a bias of PGM tools that tend to underestimate the available bandwidth. This result is obtained for a two hop network and assuming a constant rate (hence stationary) fluid cross traffic. The work by Liu et al. [6] addresses the problem of available bandwidth estimation in a general setting with multi-hop paths and bursty (stochastic) cross traffic. However,

they retain the stationarity assumption of [7] concerning the cross traffic. This means that their focus is on the long term average available bandwidth. In [6], two potential sources of bias are identified for available bandwidth measurement tools. First, considering the cross traffic as a constant rate fluid can lead to underestimate the actual available bandwidth. This bias is referred to as **elastic** since it can be mitigated by using large packet sizes and large packet trains. Second, not taking into account the multi-hop nature of an Internet path, i.e., the modeling of a path as a single hop (the tight link) can also lead to underestimate the available bandwidth. This issue affects PGM tools but not PRM tools a priori. This second source of bias is **non elastic** in the sense that it cannot be mitigated by altering the measurement strategy.

In this note, we evaluate experimentally how the biases predicted in [7] and [6] affect actual measurements, given that the stationarity assumption upon which [7] and [6] are based does not hold for day long observations of a path. To do so, we picked one representative of the PRM family, namely Pathload [1] and one representative of the PGM family, namely Spruce [10]. In addition to checking the impact of the result in [7], we address the problem of the comparison of available bandwidth tools on several day long periods.

Our contributions are the following. We propose a methodology to pre-process samples obtained from available bandwidth measurement tools and to infer the correlation that exists between those measurements, be it linear or not and be it on the short or on the long term. We demonstrate that the results of Pathload and Spruce are in general highly correlated, though the type of correlation observed is non linear, which is in line with the results predicted by [6]. In addition we show that the bias predicted in [7] corresponds to the non elastic bias of [6] and that it is the elastic bias that is dominant on our example path.




Figure 2: Time series of Pathload and Spruce. Top: first trace. Middle: second trace. Bottom: last trace

2. DATASET

We collected three long traces (12, 11 and 17 days) of Pathload and Spruce measurements between the University of Oslo, Norway and the Public University of Navarra, Spain in September and October 2006. Note that Pathload and Spruce are active tools that need to be connected to the two ends of a path to measure the available bandwidth, which prevents the collection of large scale samples. This is why we concentrate here on a single path but over a long period of time. Traceroute between the two sites along with information collected from the local administrators enabled us to gain a fairly accurate view of the path that consists of 18 hops. The core of the path consists of links in the Geant network, which feature at least 2.5Gbits/s¹, while there are apparently two bottleneck links at 100 Mbit/s in between the probes and the access links of each University. The capacity of the path, as measured by Pathrate [2], over a week long period, is 100 Mbits/s, in line with the information listed above. Note that PGM tools like Spruce require knowledge of the capacity of the narrow link, in contrast to PRM tools like Pathload. The load on the path, as measured by Pathload varies between 10 and 30%, which is typical of paths between well connected sites in the Internet [3].

Measurements were performed with the two tools running alternatively, with a safe margin of 1 minute added between any two measurements. On average, it takes 27 seconds for Pathload to return a result and 11 seconds for Spruce. The longer duration for Pathload directly follows from the iterative convergence procedure used in the Probe Rate Model as opposed to the Probe Gap Model that relies on the transmission of a single (or a few) train of packets.

3. TOOL COMPARISON

Before proceeding with the actual comparison of the tools, we cleaned the data. We first removed values larger than the path capacity (a negligible number of samples). Second, we discarded samples that are too far from the core of the measurement distribution. This operation was performed for each day and each night separately, as those periods visually exhibit different statistical characteristics. For each such period, we dropped all values outside the region $[\hat{q}_{0.25} - 1.5 \times IQR, \hat{q}_{0.75} + 1.5 \times IQR]$ where $\hat{q}_{0.25}$ and $\hat{q}_{0.75}$ are the empirical 25th and 75th quartiles of the distribution and $IQR = \hat{q}_{0.75} - \hat{q}_{0.25}$. The above data cleaning process leads to a situation where some samples are missing. To enable the comparison between tools, we averaged the values using jumping time windows of 3 minutes for both time series. Third, we used wavelets to attenuate local random fluctuations of the time series. We used the Haar Wavelets to decompose the signal using 4 levels. We then discarded the detailed signal at level 4 and reconstructed back the signal. This denoising operation conserves over 99.5% of the energy of the initial signal for all three traces. The noise removed relates to the fact that available bandwidth estimation tools estimate the available bandwidth over finite periods of time and do not measure the long term average available bandwidth [6]. The three traces before the denoising process are presented in Figure 2. The resulting time series obtained after the denoising phase are presented in Figure 3.

Visual inspection of Figure 3 reveals that (i) Spruce almost consistently returns estimates smaller than the ones of

¹See <http://www.geant.net>.

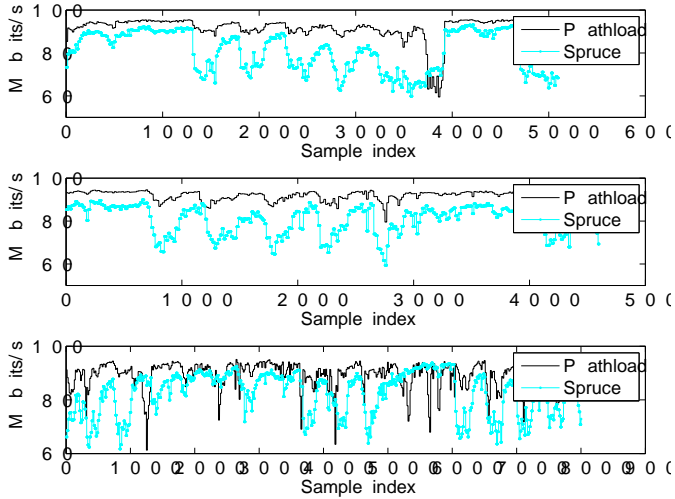



Figure 3: Denoised time series of Pathload and Spruce. Top: first trace. Middle: second trace. Bottom: last trace

Pathload, and (ii) there are periods the two tools are correlated. We used two metrics to evaluate correlation: the Pearson correlation to assess a linear relationship between the tools, and the Spearman² correlation to assess a non-linear correlation [4].


As correlation in the three traces apparently varies over time, we computed the Pearson and Spearman correlation over sliding windows of size W . We considered two values of W , namely 200 and 1000 samples³. A sample corresponding to 3 minutes, those two values roughly correspond to 10 hours and 50 hours respectively. Obviously, the smaller W , the higher the variation in the time series of coefficients. We first observed that the Pearson coefficient of correlation is consistently low (below 0.4), irrespectively of W . This clearly precludes the existence of any linear correlation between Pathload and Spruce. The use of the Spearman correlation however revealed the existence of some non linear correlation between the tools [8]. We plot in Figure 4 the Spearman coefficients of correlation over time for each of the three traces. Note that the trend observed when $W = 1000$ persists when considering $W = 200$. Our conclusion from this study is that although both tools see the same network events, as they often go up and down in a synchronized manner, they cannot be used interchangeably to analyze the dynamic of the available bandwidth in the network, as the extent of their variation is different. We can also observe that when the results are correlated, Spruce consistently returns estimates smaller than the ones of Pathload. We investigate this issue in the next section.

²The Spearman correlation, also called the Rank correlation, measures the correlation not between the initial samples, but between the ranked version of the samples. If $X = (x_1, x_2, \dots, x_n)$ and $Y = (y_1, y_2, \dots, y_n)$ are the two samples, and $Rank(X)$, $Rank(Y)$ the corresponding ranked samples, the Spearman coefficient of correlation is computed as the Pearson coefficient of correlation of $Rank(X)$ and $Rank(Y)$.


³Note that choosing smaller values of W is possible though values smaller than a few tens is questionable from a statistical point of view.



(a) First trace



(b) Second trace



(c) Third trace

Figure 4: Spearman coefficient of correlation over sliding windows of 200 and 1000 samples.

4. PATHLOAD-SPRUCE DISCREPANCY

References [7] and [6] both predict that PGM tools can underestimate the available bandwidth of path. The authors in [6] uncovered two sources of bias for available bandwidth measurement tools: one elastic bias related to the burstiness of the cross traffic and one non elastic bias related to the multi-hop nature of an Internet path.

Let us first consider the non elastic bias. Spruce is affected by non elastic bias as it relies on the simplifying assumption that an Internet path can be modeled by a single hop, which should be both the tight and the narrow link of the path.

While not stated explicitly in [7], the authors of this work focus on the non elastic bias studied in [6] and precisely identify, for a two hop network case, the cases where Spruce underestimates the actual available bandwidth, as we will show later in this section.

We adopt, in a first approximation, the two queue model of [7] to model the Norway-Spain path. We believe that this model is valid because of the characteristics of the path (see Section 2) and the fact that the core links of the path we consider are well provisioned and should with high probability

be lightly loaded. As a consequence, those links should have a negligible effect on the dispersion of the packet trains sent by an available bandwidth measurement tool. Note that the large scale study in [5] of the dispersion of packet trains at every router along a large set of paths in the Internet suggests that a significant fraction of the paths in the Internet have less than three bottlenecks (see Figure 10(a) in [5]), similarly to our example path.

We further assume that: (i) The capacity C of the two links is equal to 100 Mbits/s (see Section 2); (ii) Traffic is one hop persistent and not path persistent. The latter assumption simply means that the majority of the traffic flowing out of the first university does not reach the second one. The tight link might be either the first or the second link of the path. Using the results in [7], we obtained that the bias of Spruce, i.e., the difference between the exact value A of the available bandwidth and the value A_{Spruce} estimated by Spruce is given by Equation (1), where u_1 and u_2 are the respective utilizations of the uplink and downlink of the two universities:

$$A - A_{\text{Spruce}} = \begin{cases} Cu_1^2 & \text{if tight link=1}^{st} \text{ link} \\ Cu_1u_2 & \text{if tight link=2}^{nd} \text{ link} \end{cases} \quad (1)$$

Note that Equation (1) can equivalently be obtained using the results in [6] using Equation (33) for the Spruce estimator and Equation (7) to derive the fluid dispersion at the output of the second hop. This formally proves that the bias analyzed in [7] corresponds to the non elastic bias in [6] for the case where the two links have the same capacity (extension to the case where the two links do not have the same capacity is possible).

To find an upper bound on the utilization u_1 and u_2 , we use Pathload. Indeed, Pathload should return accurate available bandwidth estimates as (i) it is not prone to the non-elastic bias since it does not make any assumption about the path and (ii) it uses long enough packet trains (100 packets) to limit the effect of the elastic bias. In contrast, Spruce uses packet pairs to uncover the available bandwidth and as such, is extremely sensitive to the elastic bias (see [6]).

We use the 95 quantile of $\frac{C - A_{\text{Pathload}}}{C}$ as an estimate for both u_1 and u_2 . We do this so as to obtain conservative estimates of the utilizations. For our three traces, we obtain the following values : 9%, 8.8% and 11.3%. Overall, we obtain that the bias of Equation (1) (for the two cases) is upper-bounded by 0.81 Mbit/s, 0.77 Mbits/s and 1.28 Mbits/s. Those values are up to one order of magnitude smaller than the average difference between Pathload and Spruce measurements observed in our three datasets (10.6, 11.2 and 6.9 Mbits/s respectively).

The above results indicate that Spruce is mostly affected by some elastic bias. An implication of this result should be that the difference between Pathload and Spruce should increase with the intensity of the cross traffic, as we can expect that the variance of the cross traffic be positively correlated with its mean. This is indeed what we can see in Figure 3 where we observe that during the day time where the (cross) traffic is higher, the discrepancy between Spruce and Pathload measurements increases. The elastic bias should thus be a major cause of the non linearity in the correlation between Spruce and Pathload observed in Section 3.

5. CONCLUSION

In this note, we have collected available bandwidth measurements on an Internet path for long periods of time for one PRM-based tool, Pathload and one PGM-based tool, Spruce. We demonstrate that Spruce and Pathload measurements are often correlated, though the type of correlation that exists between them is non linear. We relate this non linearity to the existence of a non elastic bias in Spruce measurements that has been evidenced by [6] and [7], though under some strong stationarity assumptions concerning the cross traffic. We hope that our results will foster new research in the area of available bandwidth estimation, and especially the comparison of existing tools and the practical use of those measurements.

6. REFERENCES

- [1] C. Dovrolis and M. Jain, "End-to-End Available Bandwidth: Measurement methodology, Dynamics, and Relation with TCP Throughput", In *ACM SIGCOMM*, Pittsburgh, USA, August 2002.
- [2] C. Dovrolis, P. Ramanathan, and D. Moore, "What Do Packet Dispersion Techniques Measure?", In *IEEE INFOCOM*, Anchorage, Alaska, April 2001.
- [3] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot, "Packet-level Traffic Measurement from the Sprint IP Backbone", *IEEE Network Magazine*, November 2003.
- [4] R. V. Hogg and A. T. Craig, *Introduction to Mathematical Statistics*, Macmillan Publishing Co., Inc., New York, 1989.
- [5] N. Hu, L. E. Li, Z. M. Mao, P. Steenkiste, and J. Wang, "Locating internet bottlenecks: algorithms, measurements, and implications", In *ACM SIGCOMM*, Portland, Oregon, USA, 2004.
- [6] X. Liu, K. Ravindran, and D. Loguinov, "Multi-hop probing asymptotics in available bandwidth estimation: stochastic analysis", In *IMC*, Berkeley, CA, USA, 2005.
- [7] L. Loa, C. Dovrolis, and M. Sanadidi, "The Probe Gap Model can Underestimate the Available Bandwidth of Multihop Paths", *ACM SIGCOMM Computer Communication Review*, 36(5):29–34, 2006.
- [8] W. Martinez and A. Martinez, *Computational Statistics Handbook with MATLAB*, Chapman & Hall, 2002.
- [9] R. S. Prasad, M. Murray, C. Dovrolis, and K. C. Claffy, "Bandwidth Estimation: Metrics, Measurement Techniques, and Tools", *IEEE Network*, 17(6):27–35, November 2003.
- [10] J. Strauss, D. Katabi, and F. Kaashoek, "A measurement study of available bandwidth estimation tools", In *IMC*, Miami, Florida, USA, 2003.