Institut Eurécom
Department of Corporate Communications
2229, route des Crètes
B.P. 193
06904 Sophia-Antipolis
FRANCE

Research Report RR-07-186

On The Operational Comparaison of Available Bandwidth Estimation Tools

Alessandro Sorniotti, Taoufik En-Najjary, Guillaume Urvoy-Keller

Tel : (+33) 4 93 00 81 00
Fax : (+33) 4 93 00 82 00
Email : {Sorniott, ennajjar,urvoy}@eurecom.fr

[1]Institut Eurécom's research is partially supported by its industrial members: BMW Group Research & Technology - BMW Group Company, Bouygues Télécom, Cisco Systems, France Télécom, Hitachi, SFR, Sharp, ST Microelectronics, Swisscom, Thales.

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 but little is known about their accuracy and the type of information that can be obtained with them. This paper is an initial effort to assess the type of information that can be obtained from long term experiments with existing available bandwidth measurement tools. We propose a methodology to compare available bandwidth measurement tools with one another. We apply this methodology to three tools: Pathload, Spruce and Pathneck. We obtain insights on the degree of correlation between those tools, on the feasibility of single end measurements of the available bandwidth and on the type of events that can be observed with those tools.

1 Introduction

The number of computers connected to the Internet with an high speed access, be it DSL, cable or fiber, continues to grow at a fast pace [2]. The impact of the resulting traffic increase on the regional, national and international backbones is however not clear. Especially, it seems that the over-provisioning of backbones is such that the average load is small to negligible, though some spikes can regularly be observed as consequences of DDoS, worm outbreak or flash crowd events [12].

A natural question in this context is the utilization of end-to-end Internet paths, which relates to the daily experience of end users. More precisely, the absolute value of the available bandwidth (defined as the minimum unused bandwidth of the links of a path) as well its fluctuations over time are of high interest. A large scale and long term study of the available bandwidth in the Internet would allow to assess the spare capacity currently available in the network.

Several tools have been proposed to estimate the available bandwidth. However, they in general require to be connected to the two ends of a path, which constitutes a severe limitation to large scale measurements.In this paper, we take the stance of performing long term experiments with existing available bandwidth tools to compare them and to understand the information that be obtained from such long term experiments.

Little comparative studies have been made so far on those tools. Noticeable exceptions are [10, 11], which focus on either specific settings such as high speed networks [10], or on short time comparisons of some tools [11]. When performing the comparison, we would like to determine if the tools return the same (or very similar) results, or alternatively, if the results are not similar, we would like to know if the time series are correlated? The former case is the ideal one where tools can be used interchangeably. In the latter case, if the correlation is strong enough, tools can be used interchangeably if the focus is on the variation of the time series rather than on absolute values.

There exist two main families of available bandwidth estimation tools: the Probe Rate Model (PRM) family and the Probe Gap Model (PGM) family. We consider one tool based on the Probe Rate Model, namely Pathload [3] and one tool based on the Probe Gap Model, namely Spruce [11]. In addition, we consider Pathneck[6], a tool initially devised to detect bottlenecks in the Internet but that offers also an upper bound on the available bandwidth.

Contributions of this paper are the following. We propose a methodology to compare long term samples from available bandwidth tools and apply it to Pathload and Spruce. We demonstrate that Pathload and Spruce are in general highly correlated. In addition, we show that Spruce consistently return values smaller than the ones of Pathload, and we attempt to explain this observation using the analytical model in [9]. As for Pathneck, we demonstrate that it does not provide a meaningfull upper bound on the available bandwidth and we further highlight that there is almost no correlation between Pathneck and Pathload measurements.

The remaining of this paper is organized as follows. In Section 2, we discuss the related work in available bandwidth estimation. In Section 3, we present our dataset and the techniques used to clean the data. In Section 4, we introduce the notion of rank correlation that we use extensively to compare the tools. In Section 5, we compare Spruce and Pathload. In Section 6, we investigate the feasibility of single end measurements using Pathneck. Conclusions and future work are presented in Section 7.

2 Available Bandwidth Measurement Tools

A question of utmost 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* in a wide range of timescales makes the design of measurement algorithms very challenging. The first tool that tackled this problem was *cprobe* [1]. *Cprobe* estimates the available bandwidth based on the dispersion of long packet trains at the receiver. The underlying assumption in this work is that the dispersion of long packet trains is inversely proportional to the available bandwidth. Dovrolis et al., however, have shown that this is not the case [4]. The dispersion of long packet trains does not measure the available bandwidth in a path. Instead, it measures the *Average Dispersion Rate* (ADR), which is an upper bound for the path's available bandwidth and a lower bound for the path's capacity.

Next generations of available bandwidth estimation tools can be classified into two families: the PRM (Probe Rate Model) family and the PGM (Probe Gap Model) family.

Pathload [7] follows the Probe Rate Model. This means that it 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.

Spruce follows the Probe Gap Model. Tools based on the Probe Gap Model 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). 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 if the narrow link of a path is also the tight link (the link with the minimum available bandwidth along a path), which is assumed to be the case in the Probe Gap Model.

Pathneck does not belong to the PGM nor to the PRM family. Pathneck sends long trains of UDP packets with carefully chosen TTL values for the packets at the edges of the train. Pathneck then relies on the feedback received from intermediate routers, where the packets at the edges of the train expire, to estimate the dispersion of the train at those intermediate routers. From the above discussion on cprobe, one can note that Pathneck does not directly estimate the available bandwidth but the ADR of a path, i.e. an upper bound of the available bandwidth.

Tools	Starting day	Duration
Norway-Spain		
Pathload vs. Spruce	09/02/2005 09/17/2005 10/25/2005	12 days (period 1) 11 days (period 2) 17 days (period 3)
Pathneck vs. Pathrate	02/01/2006	4 days
Pathneck vs. Pathload	02/05/2006	13 days
Norway-Taiwan		
Pathneck vs. Pathload	03/01/2006	14 days

Table 1: Dataset description

3 Dataset

3.1 Traces

We collected traces over two paths. One path in Europe between the University of Oslo, Norway and the Public University of Navarra, Spain; and one intercontinental path between the University of Oslo and the National Cheng Kung University in Taiwan. Those two paths are similar in terms of number of hops (approximately 18 hops) but differ in terms of round trip time. The RTT is equal to 70 ms on the Norway-Spain path and 400 ms on the Norway-Taiwan path, where one hop accounts for almost 200 ms of RTT. Application of Pathrate [4] (an active measurement tool that aims at estimating the capacity of a path) to the two paths during a full week each time, returned consistent (98% of the samples) values of 100 Mbits/s for the capacity of each path.

Measurements were collected for several week-long periods, with two tools running in parallel each time. Details on each trace are provided in Table 1. On average, it takes approximately 27 seconds for Pathload to return a result, 11 seconds for Spruce and 3.7 seconds for Pathneck. The longer duration for Pathload results directly from the longer estimation time required by a tool using the Probe Rate Model as compared to a tool using the Probe Gap Model, as the former uses an iterative procedure to converge while the latter relies on the transmission of a single (or a few) train of packets. To minimize the load in the network and the correlation between consecutive samples, we added one minute of delay between each measurement of the two tools. Each sample is identified by a timestamp, calculated as the minute in which the measurement process started.

3.2 Data Cleaning and Filtering

Cleaning the data aims at removing measurement outliers. This task was performed in two stages. First, we removed values larger than the path capacity (equal to 100 Mbits/s as estimated by Pathrate). Second, we discarded samples that are too far from the core of the 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 drop all values outside the region

$[\hat{q}_{0.25} - 1.5 \times I\hat{Q}R, \hat{q}_{0.75} + 1.5 \times I\hat{Q}R]$ where $\hat{q}_{0.25}$ and $\hat{q}_{0.75}$ are the empirical 25th and 75th quantiles of the distribution and $I\hat{Q}R = \hat{q}_{0.75} - \hat{q}_{0.25}$.

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.

At this point of the process, we have removed from the initial time series values that violates physical constraints or statistical outliers, i.e., any measurement point which lies too far from the core of the distribution. We next use 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 of the traces we collected.

4 Rank Correlation

A visual inspection of some traces of our dataset suggests that some of the tools, especially Pathload and Spruce (see Section 5) are correlated. A natural way to check for correlation is to first use a scatterplot representation of the data. A scatterplot by itself does not constitute a proof and this is why it is often complemented by the (Pearson) coefficient of correlation of the two time series. The latter provides meaningful results especially if a linear relationship exists between the two random variables. Otherwise, the usefulness of the coefficient of correlation depends on the nature of the relation.

In the case of our dataset, we observed that while tools are apparently correlated, i.e., the corresponding time series goes up and down in a synchronized manner, variances of the two random processes were different, which precludes a linear correlation. The net result, in our case, is that coefficient of correlations values were low and fail to fully capture what we visually deemed as correlation. Yet, we need an objective metric to assess the correlation that exists between the tools we compare. We use the Spearman correlation [5]. The idea behind the Spearman correlation, also called the Rank correlation, is to measure the correlation not between the initial samples, but between the ranked version of the samples. If $X = (x_1, x_2, \ldots x_n)$ and $Y = (y_1, y_2, \ldots y_n)$ 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 $Rank(Y)$.

The Spearman coefficient of correlation has been observed to be more efficient than the classical (Pearson) coefficient of correlation to detect some non-linear correlations. It turned out to correctly capture the type of correlation that exists in our dataset.

We illustrate in Figure 1 both the impact of the wavelet denoising and of ranking the data prior to plotting the scatterplot or computing the coefficient of correlation. On the top row of Figure 1, we plot the original time series along with the

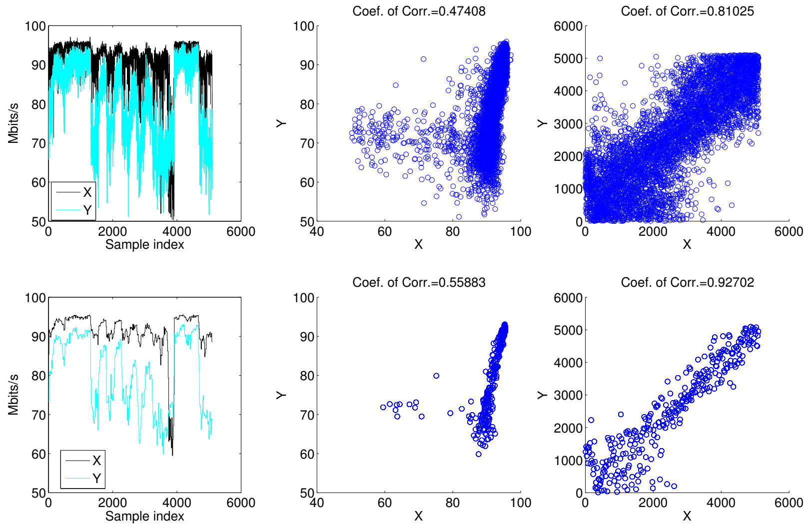

Figure 1: Top row: original time series. Bottom row: wavelets are used (over 99.5% of the energy is kept)

scatterplot of the original time series and the scatterplot of the ranked values. The corresponding coefficients of correlation are displayed as title in the figures. On the bottom row of Figure 1, we provide the same plots but once the dataset has been processed using wavelets. We observe that the coefficient of correlation as well as the scatterplot both benefit from wavelet denoising and ranking the data, though ranking has (hopefully) much more impact than denoising.

5 Pathload vs. Spruce

In this section, we focus on Pathload and Spruce. We first compare the two tools using the three traces collected on the Norway-Spain path, see Table 1. We next investigate whether the offset observed between the two tools can be explained based on the results of [9]. Last, we investigate the type of information that can be extracted from an operational use of those tools.

5.1 Tools Comparisons

We have collected three long traces of Pathload and Spruce over the Spain-Norway path, see Table 1. Denoised time series for the three traces are presented in Figure 2. Figure 3 (upper row) represents the scatterplots, for each trace, of the ranked data along with the corresponding Spearman coefficient of correlation in the title. The scatterplots and coefficients of correlation reveal that in general, Pathload and Spruce are highly correlated. Yet, correlation is much more pronounced for the first two traces (Spearman coefficient above 0.9) than for the last one (Spearman coefficient of 0.4). For the latter case, a visual inspection of the time series of the two tools (bottom curve of Figure 2) reveals that there exist some periods, e.g., interval [5000,6000] where the two tools clearly disagree while they are apparently correlated in the interval [6000,8000]. There is thus a need to measure correlation not only over the whole trace but over smaller periods of times.

To quantify the variation of correlation over time, we computed the Spearman coefficient of correlation over sliding windows of size W. The choice of W naturally impacts the observed correlation. We plot in the bottom row of Figure 3, the Spearman coefficients of correlation over time for the three traces respectively and for two values of W, namely 200 and 1000 samples[1]. A sample corresponding to 3 minutes, those two values roughly correspond to 10 hours and 50 hours. Obviously, the smaller W, the higher the variation in the time series of Spearman coefficients. Still, the trend observed when $W = 1000$ persists when considering $W = 200$.

The Spearman coefficient of correlation computed over sliding window allow to detect local correlations. Consider for instance the last 3000 samples in the third trace. Visually (see Figure 2), in the interval [5000,6000], Spruce and Pathload are not (or only weakly) correlated, while in the interval interval [6000,8000], they are strongly correlated. Those visual observations are supported by the right-bottom curve of Figure 3, where considering the case $W = 1000$, we observe a dip in the curve around value 4000 [2] followed by a constant increase toward high coefficient of correlations around index 5000.

The overall conclusions to draw from this long term comparison of Spruce and Pathload is that the two tools are often significantly correlated. The methodology we have introduced, based on the Spearman coefficients of correlation computed over time allows us to investigate precisely this correlation. Another striking observation is that Spruce always consistently outputs values smaller than the ones of Pathload. We investigate this issue in the next section.

5.2 A simple analytical model of Pathload-Spruce discrepancy

In [9], the authors prove, for a two-node network, that tools that follow the Probe Gap Model, like Spruce, tend to underestimate the available bandwidth. At first sight, those results are in line with the ones of Figure 2. Indeed, if we assume that Pathload, which is a PRM-like tool, is not affected by the bias pinpointed in [9] and thus returns accurate results, then Spruce consistently underestimates the available bandwidth. We quantify this underestimation as the mean values of the

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

[2]4000 is the index in the time series of sliding window - an offset of $W = 1000$ must be added to find the corresponding values in the initial time series.

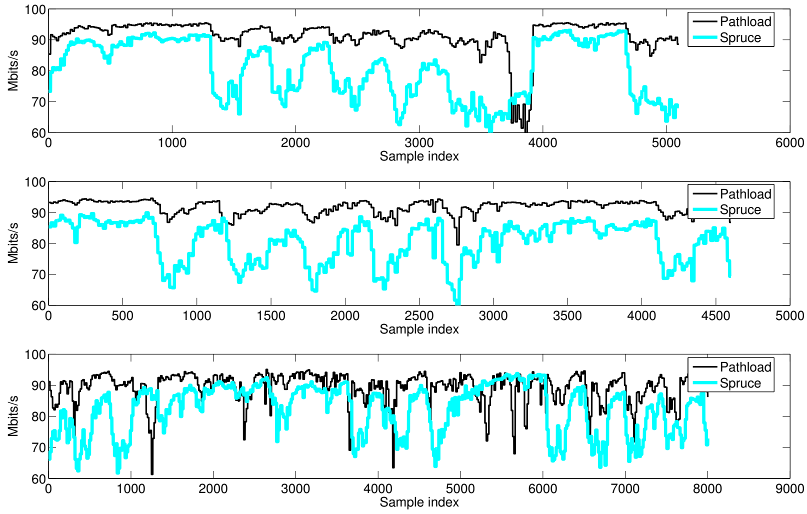

Figure 2: Denoised time series of Pathload and Spruce

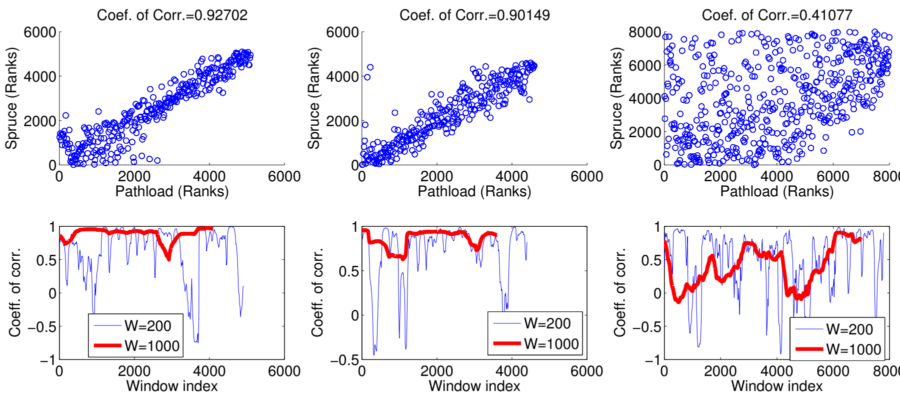

Figure 3: Top: Scatterplot of ranked data. Down: Spearman coefficient of correlation over sliding windows of $200, 1000$ samples.

difference of the two times series for the tree traces. We obtain values of 10.6, 11.2 and 6.9 Mbits/s.

The analytical model used in [9] is a two queue model. The authors account for either one hop persistent or path persistent cross traffic. In addition, the bottleneck[3] might be either the first or the second link of the model.

We decided to adopt, in a first approximation, the two queue model of [9] to model the Norway-Spain path, where the first queue represents the uplink of the source institution and the second queue represents the downlink of the destination institution. We assume that:

1. The capacity C of the two links is equal to 100 Mbits/s (Pathrate measured a path capacity of 100 Mbits/s in both directions);

2. The bottleneck is the second link, i.e. the downlink of the destination institution.

3. Traffic is one hop persistent and not path persistent, which is the general case in the Internet.

Based on the above assumptions, and using the equations of Section 3.2.2 in [9], 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. We do not provide the complete derivation of the result due to space limitation.

$$A - A_{\text{Spruce}} = C u_1 u_2 \tag{1}$$

According to our second assumption $u_1 \leq u_2$, and thus the bias of Equation 1 is upper bounded by Cu_2^2. We assess the utilization u_2 using Pathload measurements as $\frac{C - mean(A_{\text{Pathload}})}{C}$. For our three traces, we obtain the following estimations of u_2 : 8.8%, 8.2% and 10.8%. Overall, we obtain that the bias of Equation (1) is upper-bounded by 0.76 Mbit/s, 0.67Mbits/s and 1.15 Mbits/s. Those values are significantly smaller than the ones observed in our dataset (10.6, 11.2 and 6.9 Mbits/s respectively).

There are several potential sources of errors in our reasoning. A first option is that the utilization as estimated by Pathload are underestimated (i.e., Pathload overestimates the available bandwidth), as the bias of Equation (1) increases with increasing utilizations. However, even if we estimate utilizations based on Spruce measurements, the upper bound on the bias of Spruce is below 3 Mbit/s, i.e., two to three times smaller than what we observed in our 3 traces. Another option is that the two-queue model is too simplistic to model an Internet path. A last option is that the analysis in [9], while highlighting a flaw in the PGM tools (as evidenced

[3] As in the Probe Gap model, the narrow link is also the tight link, we simply refer to this link as the bottleneck

in our dataset), does not uncover the whole story. We leave for future work a more in depth study of those issues.

5.3 Path Analysis

Numerous questions can be raised based on available bandwidth measurements collected for a given path. We focus on one such questions: What kind of periodic events can we observe based on available bandwidth measurements?

Observation of Figure 2 reveals a number of periodic events. Especially, intervals [1500, 3500] in trace 1, [800, 2800] in trace 2 and [6000, 8000] in trace 3 clearly exhibit periodic patterns with a rough period of 500 samples each. 500 samples corresponding to 25 hours, we do indeed observe a day effect over a total durations of 2000 samples, i.e. approximately 4 consecutive days. We checked that those days correspond to week days. We also observe week-end effects as those days correspond to periods where the available bandwidth is close to the capacity of the path, i.e. 100 Mbits/s. For instance, the first 800 samples the second trace were collected during a week-end.

Those preliminary observations indicate that available bandwidth measurement tools can be extremly useful to assess the dynamic of Internet paths. Our long term objective being large scale available bandwidth measurements, we investigate the performance of Pathneck in the next section.

6 Pathneck

We focus in this section on the feasibility of single end available bandwidth estimation using Pathneck that was claimed to provide an upper bound of the available bandwidth [6]. Pathneck does not rely on the cooperation of the destination. Recently, a new tool called *abget* was proposed in [8] where the available bandwidth is estimated provided that the remote entity runs a TCP -like (usually a web) server. We leave for future work a more in depth study of *abget*.

6.1 Asymptotic Dispersion Rate

As discussed in Section 2, Pathneck estimates the ADR of a path. Pathrate [4] also estimates the ADR as an intermediate step to estimate the capacity of a path. In this section, we compare the ADR estimates of Pathneck and Pathrate. Note that, at first sight, Pathrate should be more accurate than Pathneck as the former is a two-end tool and uses an adaptive technique to assess the ADR (using increasing lengths for the packet trains) while the latter sends a fixed size train of packets and relies on ICMP messages for its estimation.

Figure 4 presents both the denoised time series (using the techniques describes in Section 3) and the Spearman coefficient of correlation over time. The global Spearman correlation is weak (0.26). We first observe a clear offset of about 20 Mbits/s between the two time series. As Pathneck should return an upper bound on the available bandwidth, this offset suggests a major flaw in Pathneck as it is highly unlikely that the link utilization is consistently close to 30% during 4 consecutive days.

In addition, we observe no significant correlation between Pathrate and Pathneck measurements as the Spearman coefficient of correlation over time oscillates between positive and negative values and is (almost) always below 0.5 in absolute value.

The main conclusion from this section is that Pathneck is apparently not able to accurately estimate the ADR.

In the next section, we compare Pathneck to Pathload.

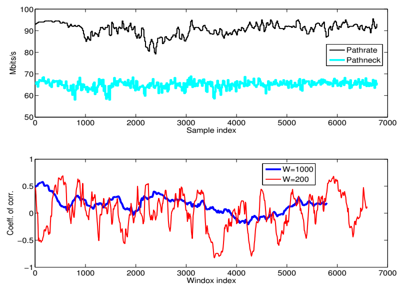

Figure 4: Top: Denoised time series. Down: Spearman coefficient of correlation over sliding windows of $200, 1000$ samples.

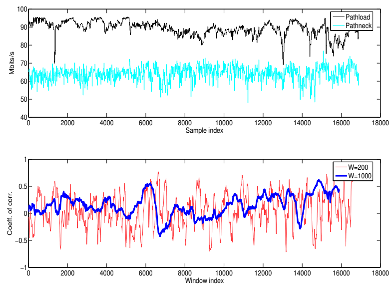

Figure 5: Top: Denoised time series. Down: Spearman coefficient of correlation over sliding windows of $200, 1000$ samples.

6.2 Pathneck vs. Pathload

6.2.1 The Norway-Spain path

We compared Pathneck to Pathload on the Norway-Spain path using a 13 day long trace. Results, see Figure 5, are similar to the ones obtained when we compared Pathneck to Pathrate: (i) there is a significant offset of around 30 Mbits/s between the two tools and (ii) little correlation exists between the corresponding time series. Those results strongly suggest that the results returned by Pathneck are correlated neither with the ADR nor with the available bandwidth of a path. The clear need for scalable measurement of the available bandwidth is thus apparently not fulfilled with Pathneck.

6.2.2 The Norway-Taiwan path

We considered the Norway-Taiwan path as an example of intercontinental path, i.e., a path for which the two ends are in far apart time zones. Our initial idea was that the variations observed for this path should not exhibit any clear periodic patterns as the ones observed on the Norway-Spain path. We compared Pathneck to Pathload for this path. We plot in Figure 6 the resulting time series for the two tools and for a period of 14 days. We first observe that the offset of about 30 Mbits/s between Pathload and Pathneck observed on the Norway-Spain path persists on the Norway-Taiwan path.

However, the most striking result is that there is almost no fluctuation of the available bandwidth, whatever the tool is. Curves are so flat (once they are cleaned) that computing any correlation coefficient is meaningless. As Pathrate measured a capacity of 100 Mbits/s for this path and Pathload reports an available bandwidth consistently close to 100 Mbits/s, our conclusion is that the utilization of the path is consistently low. This result is striking as one might expect such a long path not to be as well provisioned as a relatively small intercontinental path as is the case of our European path. A possible explanation is that the Internet traffic between Asia and Europe in general and on this specific path in particular, is in general weak, as compared to a European path or to a inter-continental path between Europe and the US.

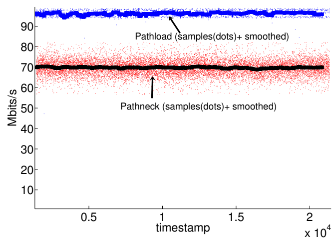

Figure 6: Pathneck vs. Pathload on the Norway-Taiwan path

7 Conclusion

In this paper, we took the stance of collecting available bandwidth measurements for a small set of paths but over long periods of time. Several tools have been devised to estimate the available bandwidth. We consider three such tools, namely Spruce, Pathload and Pathneck. Spruce and Pathload require to be connected to the two ends of a path while Pathneck, that bases its estimation on ICMP messages, requires no support from a remote party. We proposed a technique to clean the data and compare the results of two tools. Our conclusions are manifold. First, Spruce and Pathload often behave in a synchronized manner despite a relatively constant offset. An analytical model based on the findings of [9] partly explains this observed discrepancy. Second, available bandwidth measurement tools can be very helpful to uncover periodic patterns and study the long term behavior of a path. Third, we uncovered that Pathneck tries to estimate the ADR and not directly the available bandwidth of a path. In addition, we demonstrated that this tool is in general unable to provide accurate estimates of neither the ADR nor the available bandwidth of a path.

While we are aware that with this study, we have only scratched the surface of the problem, we expect that those preliminary 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. A crucial point is surely the design of an accurate single end measurement tool to allow a large scale study of Internet paths.

References

[1] R. L. Carter and M. Crovella, "Measuring Bottleneck Link Speed in Packet-Switched Networks", *Perform. Eval*, 27/28(4):297–318, 1996.

[2] K. Cho, K. Fukuda, H. Esaki, and A. Kato, "The Impact and Implications of the Growth in Residential User-to-User Traffic", In *SIGCOMM '06*, ACM Press, 2006.

[3] C. Dovrolis and M. Jain, "End-to-End Available Bandwidth: Measurement methodology, Dynamics, and Relation with TCP Throughput", In *ACM SIGCOMM*, Pittsbugrh, USA, August 2002.

[4] C. Dovrolis, P. Ramanathan, and D. Moore, "What Do Packet Dispersion Techniques Measure?", In *Proceedings of IEEE INFOCOM*, Anchorage, Alaska, April 2001.

[5] R. V. Hogg and A. F. Craig, *Introduction to Mathematical Statistics*, Prentice-Hall, 5th edition, 1995.

[6] N. Hu, L. E. Li, Z. M. Mao, P. Steenkiste, and J. Wang, "Locating internet bottlenecks: algorithms, measurements, and implications", In *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, pp. 41–54, New York, NY, USA, 2004, ACM Press.

[7] M. Jain and C. Dovrolis, "End-to-end available bandwidth: measurement methodology, dynamics, and relation with TCP throughput", *IEEE/ACM Transactions on Networking*, 11(4):537–549, 2003.

[8] M. Jain and C. Dovrolis, "Available bandwidth measurement as simple as running wget", In *Passive and Active Measurements*, 2006.

[9] 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.

[10] A. Shriram, M. Murray, Y. Hyun, N. Brownlee, A. Broido, M. Fomenkov, and K. C. Claffy, "Comparison of Public End-to-End Bandwidth Estimation Tools on High-Speed Links.", In *PAM*, pp. 306–320, 2005.

[11] J. Strauss, D. Katabi, and F. Kaashoek, "A measurement study of available bandwidth estimation tools", In *IMC '03: Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, pp. 39–44, New York, NY, USA, 2003, ACM Press.

[12] H. Wang, H. Xie, L. Qiu, Y. R. Yang, Y. Zhang, and A. Greenberg, "COPE: Traffic Engineering in Dynamic Networks", In *SIGCOMM '06*, ACM Press, 2006.