

# A Measurement Study of Active Probing on Access Links

Bjørn J. Villa and Poul E. Heegaard

Department of Telematics,  
Norwegian Institute of Science and Technology, 7491 Trondheim, Norway  
{bjorn.villa,poul.heegaard}@item.ntnu.no

**Abstract.** This paper presents a measurement study of two different methods for active probing of cross traffic on access links. The categories used in the study are packet pair probing and one-way-delay probing. The first approach uses measured increase in packet spacing as indicator of cross traffic presence, while the latter uses increase in one-way-delay for probe packets as indicator. These methods have been chosen because they are fundamentally different in terms of requirements, benefits and challenges. The main novelty of this paper is the presentation and discussion of measurement results from an access network using an adaptive video service as cross-traffic. The findings clearly illustrate the potential strengths of a probing method based on one-way-delay measurements, under the condition that the required timing accuracy is achieved and delay characteristics is available for the involved network path. The benefit of using Precision Time Protocol instead of Network Time Protocol is illustrated, even in networks of limited size.

**Keywords:** Packet Pair probing, One-Way-Delay probing, Access links.

## 1 Introduction

The dynamics of Internet traffic on different levels, ranging from per session, per user and up to backbone traffic aggregates is a topic of great interest in the research community. The reasons for performing such studies are diverse and so are also the techniques and methods applied. One obvious reason for studying Internet traffic dynamics on an aggregated level is the need for knowledge about how to best design and scale the future Internet as a whole. Another reason, on a lower level, is the growing amount of adaptive services on the Internet, which is characterized by their ability to change their requirements according to varying network conditions. Such services would obviously benefit from being able to obtain accurate views of different network- and traffic metrics in real-time.

Obtaining information about Internet traffic is best done by means of passive measurements performed in a non-intrusive way on the network links of interests. However, in some cases such passive measurements are not possible due to e.g. lack of access to the relevant links or involved equipment. In these cases, the use of active measurements is an alternative to consider. Such measurements are based on injecting probe traffic into the network between end-points and then study how this specific

traffic is treated. Based on the findings for the probe traffic, one can then make some statements about the traffic conditions along the same path as the probe traffic followed. Examples of such statements could be related to e.g. packet loss or delay, described by basic mean value considerations or higher order statistical views.

In this paper we compare two active probing methods for estimating cross-traffic amount on access links by means of theoretical discussions and measurements in a controlled lab environment. The metric of interest describing the cross traffic amount is based on buffering time observations for injected probe traffic. These observations can then be analyzed over certain periods and used for different purposes. The methods used in our study are based on sequences of packet pairs and single packets, at different rates. The way buffering time observations are extracted for these two methods are quite different and will be further described later in this paper.

The measurement part of our study is done in a controlled lab environment reflecting a typical broadband access network, and using a high quality video streaming service as the cross traffic component. The reason for focusing on the access link part is that this is where we quite often encounter the bottleneck across a network path. The choice of video as cross traffic component is based on the growing popularity of this service type on the Internet.

The structure of this paper is as follows. Section 2 provides an overview of related work; Section 3 describes the active probing methods; Section 4 presents the measurement setup; Section 5 provides the results and an analysis; Section 6 provides the conclusions and an outline of future work is given in Section 7.

## 2 Related Work

There is a lot of research in the field of active probing addressing different research questions ranging from the application layer down to the physical layer. In the context of our work, i.e. estimating the amount of cross traffic (measured by increased buffering time for probe traffic) on an access link – the most related research are found in the domain of methods for available bandwidth estimations. For this purpose there are several approaches, most of which fall into either the Probe Rate Model (PRM) or Probe Gap Model (PGM) categories [1]. The PRM approach is based on the principle of self-induced congestion and by this detecting available capacity, while the PGM approach uses observed inter-arrival time (IAT) variations [2][3] for probe packets to estimate the current level of cross traffic. However, the original idea of using probe packets as basis for active measurements was suggested in [4] where back-to-back packets were sent to detect the capacity of bottlenecks.

The use of one-way-delay (OWD) observations [5][6] for probe packets through a network can also be used for estimating available bandwidth as per the PGM approach. However, as the computation of OWD is based on time information from different nodes in the network it has very strict requirements in terms of clock accuracy and synchronization [7]. As described in [8] the main protocols for distributing clock information across a network, NTP (Network Time Protocol) and PTP (Precision Time Protocol) have different capabilities in this regard. The latter is

stated to give accuracy in the order of  $\mu\text{s}$ , while the former in the order of  $\text{ms}$ . However, it should be noted that this depends a lot on the specific hardware and software used. In a recent work [9] the performance of the Linux PTP daemon was evaluated and their findings were in line with [8]. What concerns NTP there are also improvements in this provided by the NTPv4 [10] which could bring the accuracy down into the  $\mu\text{s}$  region in certain cases.

As presented in [11] there are many sources of delay components along a network path and not all of them are influenced by cross traffic. This represents as source of error for all delay based probing methods.

### 3 Active Probing

The metric of interest to be measured by the active probing is amount of cross traffic present, represented by introduced additional buffering time for the probe traffic over some time interval. In Fig. 1 a simplified model for an access link as a basic queuing system is presented. The service rate  $\lambda_{out}$  corresponds to access capacity (bits/s), the  $\lambda_{cross}$  corresponds to the uplink capacity for the access node and  $V_B$  is the configurable buffer size (bytes) for a specific access. The indicated time parameters  $t_a$  and  $t_b$  represents time between packets in a packet pair and time between packet pairs as sent, while  $t_{a^*,i}$  and  $t_{b^*,i}$  are the corresponding values when probe traffic is received on the client side. The  $t_{s,i}$  and  $t_{r,i}$  parameters are timestamps for when a packet was sent and received.




Fig. 1. Access Link Buffer Model

In the zero cross-traffic case the client side will in theory receive the probe traffic with  $t_{a^*,i} = t_a$ ,  $t_{b^*,i} = t_b$  and a constant  $t_{s,i} - t_{r,i}$ . When cross-traffic is introduced there will be time variations in the received probe traffic, caused by additional buffering time for the probe packets in the access link buffer.

When using packet pairs, the original time between the packets ( $t_a, t_b$ ) is assumed known by the receiver. Thus, any changes to this would be caused either by cross-traffic or fluctuations in processing load on involved network components. By calculating the difference between packet spacing as received and sent – a series of samples is produced. An important benefit of this method is that there is no need for accurate time synchronization of sender and receiver side. Another benefit is that an increasing amount of probe traffic does not lead to potentially an over-sampling scenario, i.e. the registration of a certain buffering time component more than once. However, the method has a weakness in the sense that cross-traffic may delay the first packet in a pair, and thereby reduce the spacing between the packet pairs [12]. One

way to handle this is to consider both packet pair spacing and the time between packet pairs by summarizing this into packet pair period samples  $t_{pp,i}$  as given in Eq.1.

$$t_{pp,i} = [t_{a^*,i} - t_a]^+ + [t_{b^*,i} - t_b]^+ \tag{1}$$

In addition, some computation is required on the period sample time series as presented in our submitted work [13] in order to carry forward a time shift component to the next  $t_{a^*,i}$  or  $t_{b^*,i}$  observation. However, for the purpose of comparing the two active probing methods in this paper we have left this computational correction out. It is further important to note that for this method to be able to capture delay components higher than just the buffer output time of a single cross-traffic packet  $T_p$ , it is a requirement that the arrival rate  $\lambda_{cross}$  towards the bottleneck is higher than the service rate  $\lambda_{out}$ . This can be seen from Eq. 2 where the maximum value for observed time between packets in a packet pair  $t_{a^*,Max}$  is expressed.

$$t_{a^*,Max} = \frac{(\lambda_{cross}t_a)}{\lambda_{out}} + T_p, \quad \text{for } t_a \leq V_B/\lambda_{out} \tag{2}$$

In a real life network the condition  $\lambda_{cross} > \lambda_{out}$  would normally apply since an access node typically is served by at least a gigabit connection and each customer connection would be in the order of tens of Mbps. One could also claim that the approach of using packet pairs has a drawback in the sense that it requires two probe packets to produce a single cross traffic sample. However, if both packet pair spacing and time between packet pairs are considered, this is no longer applicable.

When using a sequence of single probe packets it is the OWD for each packet which is used to obtain buffering time samples. The sender adds a time stamp to each packet when sent  $t_{s,i}$ , and the receiver adds his own timestamp to the packet when received  $t_{r,i}$ . If we then know the reference OWD during times of zero cross traffic  $T_{owd,0}$ , we can for each probe packet when received - compute a sample  $t_{owd,i}$  for buffering time induced by cross-traffic.

$$t_{owd,i} = [(t_{r,i} - t_{s,i}) - T_{owd,0}]^+ \tag{3}$$

This approach has the benefit of that each probe packet gives one cross-traffic sample, and all samples are independent. However, even though the samples are independent they may actually lead to a degree of oversampling if more than one probe packet is in the buffer at the same time. The reason is that each probe packet will be delayed according to the total amount of packets ahead of it in the buffer, even if there are other probe packets as well there. Investigation of the over-sampling issue is left for future work. Further on, as the method uses timestamps from different sources (sender and receiver) it requires a high degree of accuracy in time synchronization. The use of NTP or even PTP may not be accurate enough. The challenge of actually knowing the reference OWD when no cross-traffic is present is also a significant challenge.

## 4 Measurement Setup

In order to perform a comparison of the two active probing methods, an access network testbed was established (cf. Fig. 2). In order to minimize cross process impacts on the client and server side, both the probe generator and the probe receiver were put on dedicated nodes. In a real life, this may be more integrated at least on the client side – but this depends on the specific application.




Fig. 2. Access Network Testbed

As probe traffic generator and receiver the Rude/Crude tool [14] was used. This has the capability of generating IP packets according to trace files, describing both packet size and time between packets. It also provides application level time stamping which is easily available on the receiver side. The accuracy of this tool has been shown in [15] to be in the area of  $2\mu\text{s}$ . The cross-traffic used in the measurements was a video stream operating at 5Mbps based on the MS Smooth Streaming platform [16]. In our earlier work [17] the nature of this traffic is described in more detail with special focus on its burst oriented nature. Based on this, we can state that a video server of this type connected on a 100Mbps link will send bursts of data towards the client at rates close to its link speed, independent of what the average video stream bitrate is. Thus, we have that the earlier stated  $\lambda_{\text{cross}} > \lambda_{\text{out}}$  requirement for packet probing method is met. The access capacity towards the client was configured using the QoS mechanisms provided by the Cisco switch used. This configuration gives the  $\lambda_{\text{out}}$  and also the buffer size  $V_B$  available for the specific access.

For the purpose of time synchronization of probe sender and receiver both NTPv4 and PTPv2 [18] were used in the measurements. The NTP configuration was made so that the probe receiver used the probe sender as NTP server, in order to maximize timing accuracy when using this protocol. When using PTP as synchronization protocol, the same direct relationship between the probe sender and receiver was made. The probe sender in PTP master role and the probe receiver in PTP slave role.

The operating systems used on the video client and server side were MS Windows 7 Professional, while probe sender and receiver were using Linux Ubuntu 12.04.

All processing of measurement data was done post-experiment in order to keep the cross-process impact for each node as low as possible.

## 4.1 Measurement Scenarios

A range of different probe traffic patterns was used in the measurements. The parameters subject for change were the intra-packet time values  $t_a$  and  $t_b$ , while the probe packet size was fixed at 100Byte in all cases. When using packet pair probing,  $t_a$  was always smaller than  $t_b$ , thus reflecting the time between packets in a pair. In the case with a sequence of single probe packets and measurement of OWD,  $t_a$  was set equal to  $t_b$ . The sum of  $t_a$  and  $t_b$  gives the period in the probe pattern and thereby also the probing rate in pps and bps. Ideally, the probing rate should be kept as low as possible, in order to minimize the chances of self-induced congestion or other undesirable service impact.

**Table 1.** Configuration parameters

$V_B$ [MB]	Packet Pair			One-Way-Delay		
	$t_a$ [ms]	$t_b$ [ms]	pps	$t_a$ [ms]	$t_b$ [ms]	pps
256/500/1000	0.55/0.65/0.75	4.6/4.5/4.4	388	2,60	2,60	383
256/500/1000	0.55/0.65/0.75	3.7/3.6/3.5	470	2,13	2,13	468
256/500/1000	0.55/0.65/0.75	3.0/2.9/2.8	562	1,78	1,78	559
256/500/1000	0.55/0.65/0.75	2.5/2.4/2.3	654	1,53	1,53	650
256/500/1000	0.55/0.65/0.75	2.1/2.0/1.9	752	1,33	1,33	747

The parameters given in Table 1 represent the range of different scenarios included in our measurements, for which we compared the two methods of active probing. For each scenario, measurements were done for a period of 10 minutes both with and without the cross-traffic (i.e. the 5Mbps video stream). The capacity on the access link was set to 10Mbps for all scenarios, but with different buffer ( $V_B$ ) settings configured in the router.


## 5 Results

In this section, a selection of the measurement results is presented. The specific scenario for packet pair probing where  $V_B=256/t_a=0.55/t_b=4.6$ , and for OWD probing where  $V_B=256/t_a=t_b=2.6$  is presented in detail. The presentation of the results are mainly given by means of graphical summaries, and especially by estimated probability density function (PDF) plots and the corresponding cumulative distribution function (CDF). The differences in distributional properties for the received probe traffic with and without cross-traffic present are quite well presented by this. Whenever appropriate, interesting numerical indicators are also included.

The effect of time synchronization method used (NTP, PTP) is only presented for the OWD probing method. The reason for this is that the packet pair method only uses receiver side time information, and therefore is not affected by this.


### 5.1 Packet Pair Probing Results

For the packet pair probing method, the difference between the PDF for probe packets received when there is no cross-traffic present, and when the video stream is introduced is clear in terms of the reduced distribution peak for the latter case. In addition, it is interesting to note the appearance of a small peak in the distribution for  $t_{a^*,i}$  (cf. Fig 3, left side) in the low value region in the case when cross-traffic is present. The source of this effect is the occurrence of  $t_{b^*,i} > t_b$  samples, which delays the first packet in the next packet pair – as discussed in section 4. The result of this is the low value group of  $t_{a^*,i} < t_a$  observations.



**Fig. 3.** PDF view of Packet Pair Results at  $t_a=0.55\text{ms}$ ,  $t_b=4.6\text{ms}$  and  $V_B=256$

By viewing the same measurements using a CDF plot rather than PDF, it is even easier to see that the packet pair probing is able to detect the cross-traffic (cf. Fig 4). The CDF for  $t_{a^*,i}$  observations are lifted in the low region, and reduced in the high region. The similar effect is also seen in the CDF for  $t_{b^*,i}$  observations. Thus, both  $t_{a^*,i}$  and  $t_{b^*,i}$  observations detect cross-traffic, but at the same time they have a negative impact on each other (as indicated by the low region CDF lift).




**Fig. 4.** CDF view of Packet Pair Results at  $t_a=0.55\text{ms}$ ,  $t_b=4.6\text{ms}$  and  $V_B=256$

Common for both the PDF and CDF view of the packet pair method is that even in the no Cross-Traffic scenario the received probe traffic has some deviations from the original pattern. This represents a significant source of error, which could be critical depending on how the results of the probing are to be applied.

### 5.2 One-Way-Delay Probing Results

For the OWD probing method, the difference between the PDF for probe packets received with and without cross-traffic present is quite significant in terms of shape (cf. Fig 5, left side). The dominating peak is shifted upwards when cross-traffic is present, which also contributes to a higher mean value.




**Fig. 5.** PDF and CDF view of OWD results at  $t_a=t_b=2.6\text{ms}$ ,  $V_B=256$  and NTP

The effect on the CDF (cf. Fig 5, right side) illustrates the difference even better, as the two graphs follow each other up to the ~90% level and then the cross-traffic graph flattens out. This indicates that the high ~10% amount of values differ significantly in magnitude. Same as for the packet pair probing method, both the PDF and CDF view of the measurements show variations even in the no Cross-Traffic scenario.

### 5.3 Effect of NTP versus PTP on One-Way-Delay Results

The effect of using PTP as time synchronization protocol between the probe sender and receiver instead of NTP can be illustrated by a PDF and CDF plot for OWD observations as given in Fig. 6.



**Fig. 6.** PDF and CDF view of OWD results at  $t_a=t_b=2.6\text{ms}$  and no Cross-Traffic, NTP and PTP

The use of using PTP instead of NTP clearly gives a sharper peak in the PDF function and an increased function derivate in the mid region of the CDF function. These are both indications of more accurate sync. The difference in mean value for



the observations when using NTP and PTP is about 6μS, which may increase the error in each probe sample value collected.

### 5.4 Cross Traffic Amount Detection

The active probing methods generate samples for cross-traffic amount by measuring additional buffering time for the probe traffic caused by the cross-traffic. These samples are generated by either looking at the probe packet pair IAT or probe packet OWD. The difference between measured IAT or OWD values and their corresponding reference values generate a time series of buffering time samples ( $t_{pp,i}$  or  $t_{owd,i}$ ). By summarizing the time series for each method, over a 10 minute measurement period while the video stream is at 5Mbps, we get a view on each methods capability of cross-traffic amount detection (cf. Fig. 7).




Fig. 7. Comparison of probing methods

We can see that the method based on OWD measurements detects a higher cross-traffic amount for all probe rate levels, than the method based on packet pairs. It is also worth noticing that the two highest probe rate levels give distinct higher results for both methods. The amount of cross-traffic detected does not have the same linear increase as for the probing rate, but instead it seems to cross a threshold between the third and fourth probing rate level used. Although not further investigated, this threshold may indicate that the probe traffic has crossed a level where TCP mechanisms in cross-traffic are affected and thereby the nature of the traffic is changing.

For the packet pair method and the different  $t_a$  values used (cf. Table 1) there are also quite noticeable differences. The highest  $t_a$  value gives the higher cross-traffic amount detection across all probe rate levels.

For the OWD based method it is clear that the use of PTP for time synchronization instead of NTP has a significant effect, as it always gives a higher cross-traffic estimate. The graph for the PTP case also has a more logical profile than in the NTP case, as the cross-traffic volume consistently increases with increasing probing rate. The NTP graph starts with a moderate decrease, followed by a significant peak before dropping again. Repeated experiments shows that the results for the NTP case fluctuate more than in the PTP case, which is expected due to clock drift.

The presence of potential oversampling for the OWD measurements was not investigated. However, as the trend for both packet pair probing and OWD probing is quite similar across all probing rate levels we believe that this error factor did not contribute much to the measurement results.

## 6 Conclusions

The measurement based comparison of packet pair and OWD probing presented in this paper highlights strengths and weaknesses for both methods. The objective of the comparison was not to make a statement about which one is better, as this would have to be done in the context of a specific application. However, in a scenario such as the one we used in our measurements, it is clear that the OWD based method is able to detect more of the cross-traffic than the packet pair method.

Concerning the packet pair method, we believe that our findings related to dependencies between  $t_{a^*}$  and  $t_{b^*}$  observations have significance for a range of suggested probing methods based on the packet pair principle, e.g. in the area of available bandwidth estimation. To our knowledge, this specific dependency has not been documented earlier. We further believe that our findings related to sensitivity for parameter values  $t_a$  and  $t_b$  selected are of interest. Finding the optimal parameter selection for a complex traffic mix is foreseen to be quite challenging, but an approach where a range of values are used may be a beneficial approach.

The most interesting finding for the OWD method is the significant impact of using PTPv2 instead of NTPv4. Keeping in mind that our access network lab is of a much smaller size than what a real network would be, it clearly demonstrates the shortcomings of using a NTP based synchronization for purposes like this. Another interesting finding is the challenge to establish a reference point for OWD when no cross-traffic is present. The variations observed even in the small access network lab were higher than what we expected.

## 7 Future Work

To further analyze the capabilities of active probing methods it would be interesting to also make measurements using other tools than Rude/Crude, and also use different operating systems and even HW components. Reason being that one can never neglect the possibilities when doing measurements that some of the things which are being observed have underlying reasons not related to the topic investigated.

The implementation of time synchronization between probe sender and receiver should also be closer investigated. The specific implementation of NTPv4 in Ubuntu 12.04 may have flaws, or there could even be some not obvious configuration options with a positive impact on accuracy. In a scenario with at least  $\mu\text{s}$  accuracy of time synchronization between hosts on the Internet, active probing using OWD measurements becomes very attractive.

The cross-traffic used in our measurements was of a specific type, operated at a specific quality (bitrate) level. A more composed and potentially complex cross-traffic profile would also be interesting to include in a measurement study. However, the burst oriented nature of the video service used gives a very challenging traffic pattern. Thus, we do not think a more composed cross-traffic scenario will make the research question significantly harder.

## References

- [1] Prasad, R., Dovrolis, C., Murray, M., Claffy, K.: Bandwidth estimation: Metrics, measurement techniques, and tools. *IEEE Network* 17(6), 27–35 (2003)
- [2] Strauss, J., Katabi, D., Kaashoek, F.: A measurement study of available bandwidth estimation tools. In: *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement, IMC 2003*, pp. 39–44. ACM, New York (2003)
- [3] Goldoni, E., Schivi, M.: End-to-end available bandwidth estimation tools, an experimental comparison. In: Ricciato, F., Mellia, M., Biersack, E. (eds.) *TMA 2010*. LNCS, vol. 6003, pp. 171–182. Springer, Heidelberg (2010)
- [4] Keshav, S.: A control-theoretic approach to flow control. In: *Proceedings of the Conference on Communications Architecture & Protocols, SIGCOMM 1991*, pp. 3–15. ACM, USA (1991)
- [5] Jain, M., Dovrolis, C.: Pathload: A measurement tool for end-to-end available bandwidth. In: *Proceedings of Passive and Active Measurements (PAM) Workshop*, pp. 14–25 (2002)
- [6] Huang, Y.C., Lu, C.S., Wu, H.K.: Available bandwidth estimation via one-way delay jitter and queuing delay propagation model. In: *Wireless Communications and Networking Conference*, vol. 1, pp. 112–121 (2006)
- [7] Shin, M., Park, M., Oh, D., Kim, B., Lee, J.: Clock synchronization for one-way delay measurement: A survey. In: Kim, T.-h., Adeli, H., Robles, R.J., Balitanas, M. (eds.) *ACN 2011*. CCIS, vol. 199, pp. 1–10. Springer, Heidelberg (2011)
- [8] De Vito, L., Rapuano, S., Tomaciello, L.: One-way delay measurement: State of the art. *IEEE Transactions on Instrumentation and Measurement* 57(12), 2742–2750 (2008)
- [9] Kovachazy, T., Ferencz, B.: Performance evaluation of ptpd, a ieee 1588 implementation, on the x86 linux platform for typical application scenarios. In: *2012 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, pp. 2548–2552 (2012)
- [10] Mills, D., Martin, J., Burbank, J., Kasch, W.: *Network Time Protocol Version 4: Protocol and Algorithms Specification*. RFC 5905 (Proposed Standard), IETF (June 2010)
- [11] Hernandez, A., Magana, E.: One-way delay measurement and characterization. In: *Proceedings of the Third International Conference on Networking and Services, ICNS 2007*, p. 114. IEEE Computer Society, Washington, DC (2007)

- [12] Hu, N., Member, S., Steenkiste, P., Member, S.: Evaluation and characterization of available bandwidth probing techniques. *IEEE Journal on Selected Areas in Communications* 21, 879–894 (2003)
- [13] Villa, B.J., Heegaard, P.E.: Estimating available bandwidth on access links by means of stratified probing. In: *Proceedings of the 6th International Conference on Computer and Electrical Engineering Conference (ICCEE 2013)* (October 2013)
- [14] Laine, J., Saaristo, S., Prior, R.: Real-time udp data emitter (rude) and collector for rude (crude) (January 2000), <http://rude.sourceforge.net/>
- [15] Ubik, S., Smotlach, V., Saaristo, S., Laine, J.: Low-cost precise qos measurement tool. CESNET Tech report number 7/2001
- [16] Zambelli, A.: IIS smooth streaming technical overview. Tech. Rep. (March 2009), <http://www.microsoft.com/silverlight/>
- [17] Villa, B.J., Heegaard, P.E.: Group based traffic shaping for adaptive http video streaming. In: *Proceedings of the 27th IEEE International Conference on Advanced Information Networking and Applications (AINA 2013)* (March 2013)
- [18] IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems 1588-2008, IEEE Std. 1588 (2008)