# Detection and Localisation of Performance Limitations of TCP Connections on ADSL

Denis Collange, Jean-Laurent Costeux and Louis Plissonneau
France Télécom R&D, Sophia-Antipolis, FRANCE

*Abstract*— **Our research work consists of an investigation of the performance of TCP connections of a huge number of ADSL clients, and the location of its limitations. We are particularly interested in detecting the performance limitations due to the network. Then, we focus on finding out the elements responsible for limiting the performance of TCP connections. From ISPs point of view, tracking down the performance limitations is a determining factor for quality of service and troubleshouting. As for end users, they might be interested in identifying the nature of limitations.**

## I. Introduction

### A. Aim of the study

One of the major preoccupations of ISPs nowadays remains to identify the problems of TCP connections: clients encounter problems while downloading or uploading, call the hotline and want their problems to be solved. Operators need to quickly identify the type of the problem and its location (they particularly insist on searching whether the limitations are due to the applications, the TCP end-hosts or the network).

In our case, we would like to find out which are the pertinent criteria in performance limitation detection and to classify the TCP connections in order to identify the causes of these limitations. In case of a network problem, we also want to detect the network elements involved in these performance limitations.

Several studies have been published in order to identify the causes of performance problems. In [1], the authors introduce a taxonomy of performance limitations based on a specific estimation of the RTT. In [2], [3] and [4], a correlation of TCP connections according to their RTTs and loss rates is carried out using active measurements. In [5] and [6], the authors propose a new classification criteria based on inter-arrivals of the packets (the measurement methods are passive).

All these articles try to correlate flow performances in order to recognize groups sharing the same bottleneck. But in [7], the drawbacks of active measures for detecting performance problems are highlighted: the measures are perturbed and the load is increased. The other approach, namely passive measures, remains a difficult one because of storage and analysis problems. Moreover, most of the studies based on passive measures use simulations or experimental networks.

In [8], we have implemented and tested on our passive captures the tool T-RAT, proposed in [1]. We have applied this tool to numerous ADSL connections, especially P2P connections, and have mainly obtained the following causes of limitation for connections: application limited, upload limited, congestion limited, and sender window limited. But T-RAT approach is based on *flights* of packets and flights are often difficult to identify. Another limitation lies in the impossibility to locate the network elements responsible for a congestion: T-RAT only identifies the types of limitations and not the elements.

In [9], we have implemented the method proposed in [5]. This method has been tested and validated with NS simulator, but it could not cope with the large number of TCP connections we have in our real traces. Furthermore the high values of RTT in our captures perturb this method.

In this study, we try to develop a method of classification based on performance indicator (we use the delay here) in order to apply it to ADSL traffic.

### B. Measurements Settings

First of all, we shall explain in detail our experimentation protocol. A classical ADSL architechture is organized as follows: the BAS[1] collects the traffic issued from the DSLAM[2] before forwarding it through the POP[3] to the France Telecom IP backbone. Each client is connected to one DSLAM using one VP.

Our probe is located between a BAS and the IP backbone. This BAS comprises of 10 DSLAM, connecting around 4000 clients. It is noteworthy that we capture all TCP packets of the BAS without any sampling or loss (using an adapted version of tcpdump). The collected data thus represents a huge amount of traffic; the major part of the traffic comes from TCP connections (mainly http and peer-to-peer traffic).

We then calculate flow-level delays, duration and traffic volume for all TCP connections over one day. We evaluate the delay as the interval between the passage of a packet with some data on the probe and the return of the corresponding acknowledgement (the compression of ACKs thus reduces the number of delay measures). We focus on "local delay" which represents the round trip time spent in the part of the network between the probe and the client.

## II. Analysis of Round Trip Times

In order to detect the limitations and to identify the elements involved, we investigate the performance obtained on the BAS and try to find certain criteria (threshold for example) to sort out the delays.

[1]Broadband Access Server
[2]Digital Subscriber Line Access Multiplexer
[3]Point-Of-Presence

### A. Delays of Connections on the BAS and on the DSLAMs

We measure all the connections of the BAS and evaluate the corresponding delays as explained in I-B. We observe that the mean delay by hour for all TCP connections is not correlated to the time of the day. We note instead some very high values of delay. To determine the cause of these values (are they due to congestion and where is the congestion ?), we first try to locate them on the network. Thus for each DSLAM, we follow the evolution over time of the mean of the median delay by connection. This method is effective for detecting and identifying problems on a DSLAM, if most of the clients on this element encounter bad delays. Indeed, some DSLAMs have problems at some periods of the day only, but this first result gives very restrained informations.

Our goal now is to find a threshold on the delays. We should determine this theshold by correlating the delays and the load. As long as the delays are correlated to the traffic, there is no problem of overload. When the delays are over the threshold, they are no longer correlated to the load, which means that there is an overload (load over or equal to 1, buffers stay full) in the part of the network between the BAS and the client, or that there is another limitation, probably a delay due to the application.

But, in this first step, we do not observe any correlation between the load and the delays as the delays are extremely application dependant: delays for P2P applications are much higher than those for interactive applications. Thus, the observation of the performance on the global traffic does not give an accurate information on the state of the network.

### B. Performance of Interactive Applications

In order to obtain accurate informations about the performance of the network and potential limitations, we filter the delays corresponding to HTTP connections according to the TCP port number. We have chosen HTTP because it is an interactive application. In that case, the delay is a performance criterion perceived by the end user. Furthermore, this application is frequently used by users and gives permanent informations on every DSLAM.

We try to correlate HTTP delays with the traffic volume and the number of connections. We first have to exclude extreme values. Then we have found that a threshold for off-peak hours and another one for loaded hours is a simple but efficient approach to detect network problems. In fact, as long as the delay is correlated to the load, the network encounter no problem of overload. Above these thresholds, we assume that there are congestions in the network. But the coefficient of correlation seems to be low between the delays and the load (at best, we observe time correlation between load of the DSLAM and the delays of the DSLAM).

In fact, even if the corresponding connections have been filtered from our measures, P2P applications still have a strong influence on the HTTP delays (buffers are filled with P2P packets), particularly during the night. This explains why the correlation is so difficult to obtain.

### C. Factorial Analysis and Classification

The last step of our method to distinguish faulty elements by measuring the performance consists of using a classification of performance; for example in [5], the authors use inter-arrivals. Here we classify the connections according to their delays and relate the obtained connection classes to the DSLAM, location, ISP[4] and application.

We thus conduct a factorial analysis to obtain a different representation of all the statistics we can obtain on delay (mean, minimum, maximum, deciles...). We obtain two new axes representing delay up and delay down. These new axes are made distinct with the help of a Varimax Rotation. We thus obtain an excellent representation of the performance values, with non correlated and quasi normal variables.

Then, we run a K-Means classification on these new variables and distinguish 8 classes. We draw up the characteristics of each class according to the delay, the traffic volumes and the duration of connections. Among these classes, a very large one with *normal* mean delays (below 300 ms) appears. But 7 classes with unusual performances are of greatest interest. A detailed analysis allows us to explain the causes of limitations of each class. We then present the classes according to the DSLAM, location, ISP and application.

As for the application, the ISP and the location of the destination are two interesting criteria for differentiating the performances of the connections. This is done by relating the delay informations with the traffic volume information.

## III. Future Work

We are currently applying two alternative methods proposed in [11] to our measures. On the one hand, we analyze the time series extracted from packet traces, and on the other hand, we compute a statistical classification of TCP connections according to their limitation causes. Our goal is to compare these methods and to obtain compatible classifications.

## References

[1] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker, “On the Characteristics and Origins of Internet Flow Rates,” in *ACM Sigcomm*, Pittsburgh, PA, August 2002.

[2] D. Rubinstein, D. Towsley, and J. Kurose, “Detecting shared congestion of flows via end-to-end measurements,” in *ACM Sigmetrics*, June 2000.

[3] M. Kim, T. Kim, Y. Shin, S. Lam, and E. Powers, “A wavelet-based approach to detect shared congestion,” in *ACM Sigcomm*, August 2004.

[4] K. Harfoush, A. Bestavros, and J. Byers, “Robust identification of shared losses using end-to-end unicast probes,” in *8th International Conference on Network Protocols (ICNP)*, November 2000.

[5] D. Katabi and C. Blake, “Inferring congestion sharing and path characteristics from packet inter-arrival times,” MIT/LCS Technical Report, 2001.

[6] D. Katabi, I. Bazzi, and X. Yang, “A passive approach for detecting shared bottlenecks,” in *IEEE International Conference on Computer Communications and Networks (ICCCN)*, Arizona, 2001.

[7] P. Huang, Feldmann, and W. Willinger, “A non-intrusive, wavelet-based approach to detecting network performance problems,” in *ACM Sigcomm Internet Measurement Workshop*, August 2001.

[8] F. Hernanz, “Analyse de flots de trafic sur un réseau pour la détection de problèmes,” Technical report, France Telecom, February 2002.

[9] R. B. Abdeljelil, “Détection et identification des problèmes de congestion dans les réseaux d'entreprise,” Technical report, France Telecom, September 2004.

[4]Note that the geographical and ISP informations are retrieved using MaxMind [10]

[10] www.maxmind.com.
[11] M. Siekkinen, G. Urvoy-Keller, E. Biersack, and T. En-Najjary, "Disambiguating network effects from edge effects in TCP connections," Technical report, Institut Eurecom, May 2005.