

Anomaly root cause diagnosis from active and passive measurement analysis

Ziad TLAISS *[†]

Orange Labs Networks, Lannion[†]
IMT Atlantique, Brest*

Abstract—Diagnosis demands a deep analysis of data to identify the root cause of an anomaly and still mostly relies on human experts. The increase of Internet traffic combined with the arrival of the encrypted protocol QUIC which invalidates many troubleshooting methods, urges to automate this process. To this effect, both domain familiarity and analysis skills are required. In this work we present our methods and strategies to detect the root cause of anomalies from active and passive network measurement and we share our plan towards an automatic root cause diagnosis. We focus on four root causes : transmission, congestion, application limited and packet delay variation, and present the building blocks of classification methods.

Index Terms—troubleshooting, active probe, passive probe, TCP, QUIC, Congestion Control algorithm, time series data

I. INTRODUCTION

Troubleshooting is essential to the operators to offer good network services to their clients. Telecommunication operators have made substantial investment to ensure high QoS (Quality of Service) and QoE (Quality of Experience) to their customers, this has been made possible thanks to relentless tracking of network degradation. In this article we focus on troubleshooting that rely on transport layer and more precisely on congestion control algorithm behaviour. Operators use active and passive probes to perform network measurement. This troubleshooting method is considered as one of the best solutions to detect root causes that directly impact the end user. In this purpose we should distinguish two types of network measurement : active and passive measurement.

Active measurement consists in observing test traffic generated by active probes, sometimes called robots. These probes will generate packets that are sent over the network to perform different measures. The main role of active probes is to act as an end-user client connected to the Internet through an access network. In passive measurement, probes are deployed in the network with the objective of tracking and collecting real users data.

The collected data from active and passive measurement is presented as a time series data. For confidential matters all the data presented in this work will be extracted with active probes only. However the analysis and diagnosis methods of the data collected with passive probes still remain the same while using TCP. Sequence number, acknowledgment, receiving window, selective acknowledgment, round trip time and bytes in flight represent the different metrics captured by probes and extracted or calculated from TCP packet headers.

II. METRICS DESCRIPTION

Active and passive probes timestamp each captured packet which leads to a time series. In this section we describe the most important metrics that we collect and analyse to detect the root cause of an anomaly.

- Sequence (SEQ): The sequence number identifies the first data byte in a packet. [1].
- Acknowledgment (ACK) : The acknowledgment number is the sequence number of the next byte the receiver expects to receive.
- Receiving window (RWIN) : The receiving window identifies the number of bytes that the receiver can accept. It is sent to the source by the receiver in every ACK.
- Selective acknowledgment (SACK) : With selective acknowledgments, the data receiver can inform the source about all segments that have arrived successfully.
- Bytes in flight (BIF) : Bytes in flight are the number of bytes sent by the source but not yet acknowledged. Contrarily to the RWIN, SEQ and ACK, BIF is not included in the TCP header. To measure the BIF, we calculate the bytes difference between the sequence number and the last received ACK. If the congestion control algorithm has the sack option then we take this information into consideration for better BIF estimation.
- Round-Trip-Time (RTT) : Round trip time is the delay between the emission of a packet and the reception of the corresponding acknowledgment. Like the BIF, the RTT is not included in the TCP header. To measure the RTT we compute the time difference between the observation of the sequence number and its acknowledgment.

Identifying the root cause of an anomaly is essential for an efficient troubleshooting. This requires a deep knowledge of the congestion control algorithm's (CCA) behaviour; this behaviour can be caught by observing the time series of CCA states (e.g. : Slow-Start, Congestion Avoidance, Fast Retransmit, Fast Recovery [3]). CCA states are easily readable in the sender stack, typically a server in an active measurement scenario. On the contrary, internal CCA states are out of reach for a mid-point observer, typical of a passive measurement scenario. In that case, human experts typically derive CCA state time series from the analysis of the previously defined metrics (BIF, SEQ, etc.). For example, detecting Slow-Start state exit time and determining if it's an early exit (or not) is the first necessary step to investigate an anomaly. This

detection is crucial because in this state the CCA estimates the bottleneck capacity, and overestimating or underestimating the bottleneck capacity will lead to a bad connection throughput.

III. TIME SERIES ANALYSIS FOR ROOT CAUSE IDENTIFICATION

In this section we analyse time series captured by active probes so as to classify them into root causes. We focus on four of the most frequent root causes in Internet networks. In each case we expose our data and explain the reason underlying our diagnosis. The presented data are collected from RAN and fully wired network with various TCP versions.

A. Transmission loss

Transmission losses occur when packets are dropped without any congestion in a network. It mostly affects Radio Access Networks (RAN). Radars, weather and RAN interference can be the origin of this issue. We can identify it by analysing packet sequence numbers and their acknowledgements. Transmission loss appears as individual and isolated packet losses as shown on Fig. 1 (where we have three individual losses at 13.61s, 13.82s and 14.03s). A loss can be identified by a sequence number decrease due to packet retransmission. This loss is considered isolated as we have only one packet lost from a burst of packets. It's harmful when it is repeated many times and more precisely at the beginning of the connection. Indeed it causes the CCA to exit the slow-start state early, because it considers these losses as a sign of a congestion, thus underestimating the bottleneck.




Fig. 1. transmission loss detection using seq & ack evolution

B. Network congestion

Network congestion can occur when the emitted traffic throughput is greater than the network capacity. It may affect any network, wireless or wireline and mostly depends on link capacity and buffer sizing. Contrarily to transmission loss, to detect congestion we look for a burst of packet losses as shown on Fig. 2 (from 1.88s to 1.94s). These lost packets will be accompanied by a significant increase in RTT as shown on Fig. 3 at 1.90s where RTT value starts to rise from 200ms to 625ms at 2.30s.

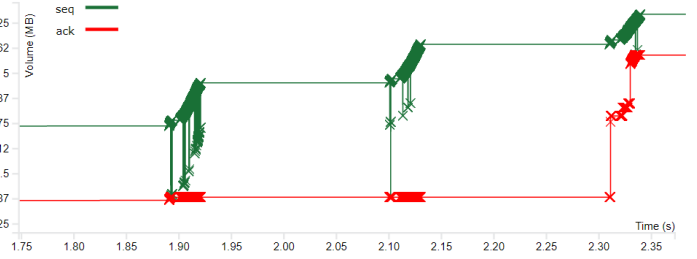


Fig. 2. congestion detection using seq & ack evolution




Fig. 3. RTT evolution

C. Application limited

Beside network issues, the performance of a TCP transfer can be affected by the endpoints' ability to send or receive fast enough. This problem only depends on the server and terminal. Servers are typically sized not to be overwhelmed by normal traffic, so the "slow sender" case is rare. On the contrary, clients are typically battery-powered devices with limited resources, and are fairly often the limiting factor. We thus focus on this case here.

This "slow receiver" situation is defined by the RWIN limiting the BIF, as shown on Fig 4. Here we can notice the BIF's exponential growth in the slow-start phase, ending at 1.4s by reaching the RWIN's limitation. This indicates that the receiver's buffer cannot grow beyond that size; accordingly, the source's sending rate cannot increase any further without overwhelming the receiver.




Fig. 4. application limited detection using rwin & bif evolution

D. Packet delay variation impact

Packet delay variation, also called jitter, is a network issue that normally affects wireless users. Jitter causes bad performance due to early Slow Start exit when the CCA

misinterprets it as a sign of congestion. To verify if jitter is the root cause of bad performance we first compute the slow-start exit time. After that, we correlate with other time series (e.g. SEQ, ACK, RTT) looking for packet loss or RTT increase. If no loss is detected and high RTT variance is observed, then we can conclude that CCA has exited slow-start due to jitter.




Fig. 5. bif evolution




Fig. 6. bif evolution

In Fig 6 the CCA exits slow start at 50000 bytes (0.7s), while the true bottleneck value we can derive at the end of transmission is around 1.6×10^6 bytes as shown on Fig 5 which lead to a bad connection throughput.

IV. TOWARDS AN AUTOMATIC TROUBLESHOOTING!

Finding the root cause of degradation within the network is time consuming for human experts. With the skyrocketing growth of Internet traffic, this task becomes intractable without automation. The faster we can detect the heart of the problem in networks, the quicker we can solve it before impacting the experience of customers.

A. Slow-start exit automatic detection

Slow-start exit can be detected by searching the exponential part in BIF time series. For this purpose, we distinguish two steps : the first one with traffic with headers in clear (TCP), the second one with deeply encrypted traffic (QUIC).

1) *Detection with readable headers:* Our work mainly focuses on slow-start automated detection in TCP where BIF series is at hand. We first tried to identify the Slow-Start phase by fitting the BIF series to an exponential, but the results were disappointing with bursty traffic which typically exhibits an on-off pattern. We then designed a promising innovative smoothing method which will be published soon.

2) *Detection with encrypted headers:* Strictly applying the previous slow-start detection method on passive measurements with encrypted traffic is impossible. Indeed, only packets timestamp and length are reachable in QUIC, thus BIF and RTT are out of reach.

To overcome this difficulty, we propose to use machine learning algorithms trained on clear-headers samples: thus, we first classify TCP captures, then turn them into QUIC-like captures, erasing all header information. We then feed them into the ML algorithm, using their previously determined class as supervision. We then expect the algorithm to extract the proper regularities, assuming it is possible.

B. Labelled dataset

We gathered more than 400 captures from active probes in Orange networks and manually labelled them into root causes using human expertise. This dataset relies on probes in a broad range of conditions: 4G, wireline, various TCP versions, various RTT.

This dataset will be used both to evaluate analytical methods and to train machine-learning algorithms.

The next step is to "QUIC-ize" this dataset so as to mimic encrypted traffic, so as to verify the consistency of our results;

V. CONCLUSION

This thesis contributes to the research and industry in the field of network monitoring and troubleshooting. Automatic troubleshooting is not simple, especially with the arrival of QUIC that invalidates passive troubleshooting methods due to packets headers encryption. In this thesis, typical degradation root causes are analysed, descriptors are defined and new classification methods are proposed. These methods - coupled with new learning solutions such as Machine Learning - seem promising; they potentially allow a high automation of processes while being compatible with deeply encrypted flows such as QUIC.

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to my supervisors Mr. Alexandre Ferrieux and Ms. Isabelle Hamchaoui in Orange Labs and Ms. Sandrine Vaton and Ms. Isabel Amigo in IMT Atlantique for their advice on this research and on my PhD project.

REFERENCES

- [1] Information Sciences Institute University of Southern California, "TRANSMISSION CONTROL PROTOCOL," RFC 793, September 1981.
- [2] M. Mathis, J. Mahdavi, S. Floyd, A. Romanow, "TCP Selective Acknowledgment Options," RFC 2018, October 1996.
- [3] W. Stevens, "TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms," RFC 2001, January 1997.