NAPlr: Multi-Gigabit Internet Access Profiling in Web Browser

Péter Orosz, Tamás Marosits, Tamás Skopkó
Dept. of Telecommunications and Artificial Intelligence
Budapest University of Technology and Economics
Budapest, Hungary
orosz@tmit.bme.hu, marosits@tmit.bme.hu, skopko@tmit.bme.hu

***Abstract*—Advancements in internet service provision have enabled subscribers to access multi-gigabit throughput for both wired and mobile connections. Despite underutilization, users demand compliance with subscription parameters, necessitating precise measurement of multi-gigabit connections. This paper investigates the challenges associated with profiling multi-gigabit internet access using web browsers, addressing limitations introduced by browser capabilities, client hardware, and SOHO router performance. We present a novel measurement system designed to accurately evaluate key service quality parameters, including throughput, latency, and packet loss, for multi-gigabit connections. The system architecture comprises user endpoint devices and measurement servers strategically positioned in high-bandwidth, network-neutral locations to ensure accurate and scalable measurements. Empirical results demonstrate the system's efficacy in measuring multi-gigabit connections, even within the constraints of client and network hardware limitations. This research establishes a robust framework for reliable and comprehensive assessment of high-speed internet access services, contributing to improved service quality and user satisfaction.**

I. Introduction

Recent advancements by Internet Service Providers have enabled subscribers of wired internet access to opt for packages offering multi-gigabit throughput for several years. Similarly, mobile internet access now surpasses 1 Gbps. Although the majority of users do not utilize this throughput to an extent where its precise value is critical, the same observation applies to 1 Gbps connections [1] [2]. Despite this underutilization, customers demand adherence to the parameters specified in their subscription contracts, necessitating accurate measurement of multi-gigabit connection characteristics.

Among quality of service (QoS) parameters, throughput typically receives the most emphasis. However, given the diverse range of modern applications, parameters such as latency, its variation, and packet loss are also critically important for evaluating service quality [3].

Examining the characteristics of multi-gigabit access networks presents several challenges. These challenges often pertain to features of the connection that are not apparent during routine use. Starting from the client side, the initial challenge is the browser used for measurement. Browsers can limit their own speed significantly below the network's maximum throughput. For example, Chrome was observed to use approximately 2 Gbps on a 10 Gbps link, whereas Firefox and Edge could achieve around 4 Gbps.

The next limiting factor is the client itself, where factors such as computing power, memory, and network interface card characteristics influence the measurable speed.

Finally, the last limiting elements are the SOHO routers, as their LAN-side ports are typically capable of a maximum speed of 1 Gbps. Consequently, even a high-performance client cannot fully utilize the throughput of a multi-gigabit internet connection. In this demo paper, we introduce our advanced Internet access profiler, namely Network Access Profiler (NAPlr). Its main novelty is the aggregated measurement capability incorporating multiple clients to form a virtual client that performs measurement towards the measurement server. We demonstrate that a 10 Gbps Internet access link can be profiled with high accuracy using three measurement clients concurrently, each equipped with an Intel i7 CPU, 8 GB of RAM, and a 10 GbE Intel NIC. A 10GbE layer-3 switch will provide the gateway function between the local network and the internet access WAN link. The measurement will be performed in Firefox 115. In this setup, the main performance constraining factor is the web browser and its JavaScript engine.

II. Related Work

While there are many different tools available for bandwidth estimation (even available as command line tools), in this section, we focus on solutions offering measurement in a web browser. All of these tools work by saturating the user's network link.

Presumably, the most widely used Internet access profiling tool is SpeedTest by Ookla [4]. Most users seem to use it via web browser but Ookla also offers mobile app. Moreover, it is also available as command line client. It uses TCP for data transmission and fallbacks to legacy HTTP if needed. ISPs can deploy their own measurement servers to provide measurement points for users and the client application chooses an optimal, low latency and high bandwidth measurement endpoint automatically. The data is collected on the server but providers can access the collected data.

RTR-NetzTest is a very similar SpeedTest alternative but it is only available through browser [5]. The measurement server source code is also available. Measurement is based on TCP/WebSockets and during the development, a compact

978-3-903176-66-9 ©2024 IFIP

architecture with minimal customization was in mind. Its measurement method is well documented.

Measurement Lab has an open source measurement software NDT7 [6]. It uses only one TCP session to saturate the link and that is stated to be enough event for 3 Gbps on decent hardware. It collects the data by its own server module so the data remains at the server maintainer. MacMillan et al. published a comparison against SpeedTest [7]. It also offers a JavaScript library for building a web interface.

BME SmartCom Lab develops szelessav.net measurement system for years [8]. It is not an open source solution but publicly available for testing and offers some customization of the measurement methods. It also stores detailed measurement data in database for later analysis. Current development focus is on scaling and to make it capable for assessment of multi Gbps user links by utilizing multiple end devices.

III. Architecture and Methodology

Internet access profiling with NAPlr consists of metering various metrics (i.e., downlink and uplink aggregated TCP goodput, packet loss, RTT, and RTT variance) throughout the measurement requested from the user side. The user has to use a web application in a browser. The app is fetched from the HTTP server of the measurement system as well as its initial settings. Scheduling and metering is done on the server side and the measurement traffic is exchanged using multiple WebSocket connections adaptively between the server and the client's browser [9]. This architecture is overviewed in Fig. 1. Bandwidth measurements occupy significantly more bandwidth than fetching the web app thus the group of WebSocket connections is marked with thicker line.

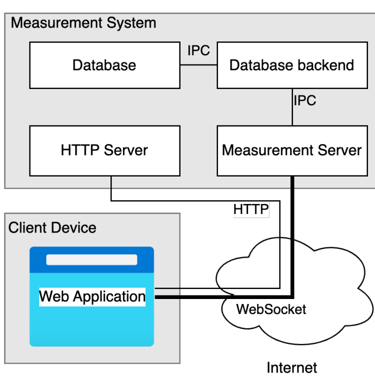

Fig. 1. Architectural main building blocks

One assessment itself consists of the following sessions: RTT on idle connection; download bandwidth measurement and upload bandwidth measurement. After completing all the sessions the results will be summarized and displayed to the user. During the bandwidth measurements the system will try to saturate the available data path capacity and analyses the traffic of the involved WebSockets.

Although it is not intended to vary or modify the measurement methods, some client and server side parameters can be modified to fit the current environment or application.

A. Conditions

- To achieve the specified measurement accuracy and to serve multiple users in parallel or to measure multi-Gigabit accesses, measurement servers must be scaled and configured appropriately. Load-balanced configuration of measurement servers is recommended.
- The gateway device at the user end must be able to serve his/her Internet access in both directions.

B. Parts of the measuring system

- User endpoint device: desktop or mobile device, but can also be a dedicated fixed wired or mobile device for automated measurements. If a single device is not capable to fully utilize the bandwidth, multiple devices can be involved into a parallelized measurement.
- Measurement server: schedules measurements requested from the endpoint devices; receives or transfers the data stream depending on the measurement phase; measures and evaluates the associated QoS metrics; stores measurement data.

C. Syncronization

If more than one user devices are involved in the measurement, they have to be synchronised. One of the simpliest ways is to use a PIN. The first client device that requests a new measurement will be the master instance. The user will receive a PIN from the measurement server and will have to enter it on all the other client (slave) instances. Once all slave devices have connected with the PIN, the measurement will start at the same time on all client devices. In case of automated measurements, pre-generated PINs are used.

D. Data storage

All sampled measurement data is stored including detailed time series patterns. In the case of multi-client measurement, the samples of the individual clients are stored in separate records, as well as the aggregated data. This allows detailed reconstruction of the measurements and can be subject to subsequent big data analysis.

IV. Demonstration

Since the primary aim is to assess the IP path capacity between the user's Internet access and the measurement server, it is recommended to place the measurement servers at a high bandwidth, network neutral location, e.g. in an Internet Exchange [10]. A typical measurement setup is shown in Fig. 2. This is also our demonstration setup. User-end devices are usually (although not technically necessarily) located on the user's LAN network.

During the demonstration, we illustrate that the developed measurement system is effective for evaluating service quality parameters of multi-gigabit connections. The resource-efficient nature of the system dictates that the number of required clients is determined by the bottleneck capacities along the path between each client and the measurement server. The cumulative capacity must match the throughput of the access

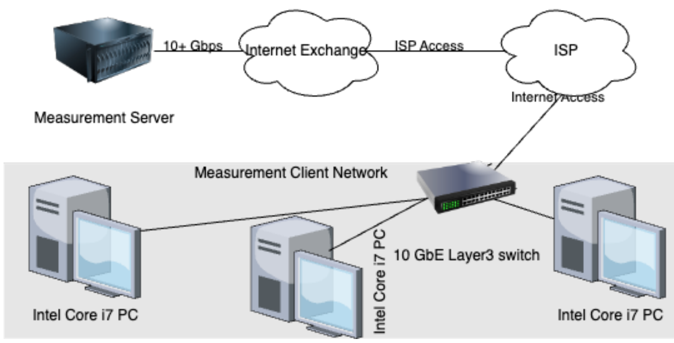

Fig. 2. The demo setup

network to ensure the clients' traffic fully saturates the connection. However, achieving this can be physically impossible; for instance, saturating a 10 Gbps WAN link via a router modem with four 1 Gbps LAN ports, even with WiFi connectivity, is infeasible.

There are no specific requirements for clients; the operating system, browser, and medium for connecting to the router can all vary. Clients participating in the same measurement session are linked by a PIN assigned to the primary client by the measurement system. This PIN must be entered by the other clients before commencing the measurement, and all clients must approve the start of the measurement. Unlike single-client measurements, the server initiates the measurement upon receiving approval from all connected clients.

The measurement system assesses a broad spectrum of service quality parameters, although some parameters are challenging or impractical to interpret in a multi-client environment. For example, latency measurements on unloaded and loaded networks are inherently restricted to the primary client. Clients conduct throughput measurements synchronously in both download and upload directions, with the aggregate values forming a sample of the measurement outcome. These samples naturally exhibit variability, similar to single-client measurements, and no universally applicable method exists to mitigate this fluctuation. The final measurement result is derived from the sample series.

Ideally, service quality measurements on a multi-client, multi-gigabit access network should yield results comparable to those obtained from a single-client measurement, with throughput measurement errors being minimal (a few percent, typically). However, this does not necessarily hold true for latency and packet loss measurements on a saturated network. For instance, measuring a 2 Gbps link with a single client equipped with a 10 Gbps NIC versus two clients with 1 Gbps NICs each presents different saturation scenarios: the former saturates the connection but not the client's NIC, while the latter saturates both the connection and the clients' NICs during latency measurements on a fully loaded network.

Throughout the measurement process, users receive continuous updates on download and upload rate results for both single-client and multi-client cases. In multi-client scenarios, the presented sample values represent aggregated results from all participating clients. Upon completion, each client, regardless of role, reports the aggregated result for the entire connection, which is made available for download in various formats. The database results are accessible to analysts from the Hungarian National Media and Infocommunications Authority, and registered users can retrieve them at any time.

V. Conclusions and Future Work

This demo peper introduced and demonstrated the efficacy of the Network Access Profiler (NAPlr) system for accurately measuring service quality parameters of multi-gigabit internet connections within web browsers. The system successfully addresses the challenges inherent in such measurements, particularly the limitations imposed by client hardware, browser capabilities, and SOHO router configurations. Through a synchronized multi-client approach, NAPlr enables comprehensive assessments of throughput well beyond 10Gbps, latency, and packet loss, offering a robust solution for profiling high-speed internet access in both wired and wireless environments. The system's resource-efficient design ensures scalability and adaptability across diverse network conditions, making it a valuable tool for both end-users and network analysts. Future work may explore the extension of NAPlr to accommodate emerging networking technologies and further refine the system's capabilities in capturing more fine-details of quality of service parameters.

References

[1] Cisco, "Cisco Annual Internet Report Complete Forecast Update, 2018–2023," 2020. [Online]. Available: https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2020/pdf/PSO-2556.pdf

[2] R. Sanchez-Arias, L. G. Jaimes, S. Taj, and M. S. Habib, "Understanding the State of Broadband Connectivity: An Analysis of Speedtests and Emerging Technologies," *IEEE Access*, vol. 11, pp. 101 580–101 603, 2023, conference Name: IEEE Access. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10244030

[3] E. Saiz, E. Ibarrola, E. Atxutegi, and F. Liberal, "A unified framework of internet access speed measurements," in *2015 ITU Kaleidoscope: Trust in the Information Society (K-2015)*, 2015, pp. 1–7. [Online]. Available: https://ieeexplore.ieee.org/document/7383644

[4] "Speedtest® Methodology | Ookla®." [Online]. Available: https://www.ookla.com/resources/guides/speedtest-methodology

[5] "RTR-NetTest." [Online]. Available: https://www.netztest.at/

[6] "NDT (Network Diagnostic Tool) - M-Lab." [Online]. Available: https://www.measurementlab.net/tests/ndt/

[7] K. MacMillan, T. Mangla, J. Saxon, N. P. Marwell, and N. Feamster, "A Comparative Analysis of Ookla Speedtest and Measurement Labs Network Diagnostic Test (NDT7)," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 7, no. 1, pp. 1–26, Feb. 2023. [Online]. Available: https://dl.acm.org/doi/10.1145/3579448

[8] "Szélessáv.net." [Online]. Available: https://szelessav.net/

[9] A. Melnikov and I. Fette, "The WebSocket Protocol," Internet Engineering Task Force, Request for Comments RFC 6455, 2011, num Pages: 71. [Online]. Available: https://datatracker.ietf.org/doc/rfc6455

[10] P. Orosz, T. Skopkó, and T. Marosits, "Application-Aware Analysis of Network Neutrality: A Scalable Real-Time Method," *Infocommunications journal*, vol. 15, no. 1, pp. 77–86, 2023.