

A First Look at Starlink Performance

François Michel*

UCLouvain

francois.michel@uclouvain.be

Danilo Giordano

Politecnico di Torino

danilo.giordano@polito.it

Martino Trevisan

University of Trieste

martino.trevisan@dia.units.it

Olivier Bonaventure

UCLouvain

olivier.bonaventure@uclouvain.be

ABSTRACT

With new Low Earth Orbit satellite constellations such as Starlink, satellite-based Internet access is becoming an alternative to traditional fixed and wireless technologies with comparable throughputs and latencies. In this paper, we investigate the user-perceived performance of Starlink. Our measurements show that latency remains low and does not vary significantly under idle or lightly loaded links. Compared to another commercial Internet access using a geostationary satellite, Starlink achieves higher TCP throughput and provides faster web browsing. To avoid interference from performance enhancing proxies commonly used in satellite networks, we also use QUIC to assess performance under load and packet loss. Our results indicate that delay and packet loss increase slightly under load for both upload and download.

CCS CONCEPTS

• **Networks** → **Network measurement**; *Wireless access networks*.

KEYWORDS

Starlink, Satellite Communications, Low Earth Orbit, Network Performance, Measurements

ACM Reference Format:

François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. 2022. A First Look at Starlink Performance. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22)*, October 25–27, 2022, Nice, France. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3517745.3561416>

1 INTRODUCTION

Internet access technologies and Internet protocols are constantly evolving. Broadband technologies such as xDSL and cable modems are prevalent today, but they are being replaced by optical fibers. In densely populated areas, such as cities, fiber deployment can be profitable, while in rural or mountainous areas, however, it can be much more expensive. For this reason, network operators have been working on other Internet access technologies since considerable time. Some propose Fixed Wireless Access (FWA) technologies [15, 17]. Others are deploying hybrid networks that

combine cellular and xDSL [22, 28]. Given the opportunities offered by these rural areas, several companies nowadays offer satellite-based Internet access solutions.

Classical Satellite Communications (SatCom) use geostationary satellites with an orbit of 22 236 miles. A single satellite can cover a large portion of the Earth at the price of a latency of several hundreds of milliseconds due to their high elevation [18, 30]. Such communication technology may provide connectivity to thousands customers with connections easily reaching a speed up to 100 MB, with the drawback of a minimum latency of about 600 ms [37].

A new approach is to use a constellation of Low Earth Orbit (LEO) satellites to dramatically reduce communications latency. The first large-scale deployment of this kind is the Starlink constellation, currently operating more than two thousand satellites. The commercial service started in beta version in October 2020 in the United States and from 2021 in European countries. It promises Internet access with latency on the order of 20 ms and bandwidth speeds between 100 and 200 Mbps [12]. Being this a newborn service, its operation and performance have not been fully investigated yet. The only comparable work has been proposed by Kassem *et al.* [34], which shows how Starlink performance changes from different vantage points. We here focus on how the performance of a single Starlink vantage point changes when accessing globally distributed resources, under high and heavy network loads, with the TCP and QUIC transport protocols.

For many years, TCP has been the dominant protocol for Internet services [32, 41]. SatCom operators therefore widely adopt TCP Performance Enhancing Proxies [23] (PEP) to mitigate the impact of increased latency on TCP performance. Recently, the Internet Engineering Task Force (IETF) has standardized the QUIC protocol [26, 31]. In short, QUIC combines the features of TCP and TLS [39] into a single protocol above UDP. QUIC is already widely deployed by major cloud providers and it drives a growing share of Internet traffic [40, 42, 44]. In contrast with TCP, QUIC cannot be optimized by using PEPs in satellite networks since QUIC packets are encrypted and authenticated. Given the current growth of QUIC traffic, it is important to evaluate new access networks using both QUIC and TCP.

In this paper, we benchmark the Starlink service and compare it to traditional SatCom networks. We measure the performance in terms of throughput for QUIC and TCP, latency, and packet loss, and find that Starlink delivers on its performance promises and enables the use of demanding services such as high-definition video streaming or cloud gaming. We also find that Quality of Experience (QoE) for Web browsing with Starlink is far better than with traditional SatCom and comes close to wired access. To enable

*François Michel is FNRS Research Fellow

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

IMC '22, October 25–27, 2022, Nice, France

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9259-4/22/10...\$15.00

<https://doi.org/10.1145/3517745.3561416>

the research community to perform experiments with an emulated Starlink connection, and compare it with other connection technologies (included 3G, 4G, and classical geostationary SatCom with data from [37]), we have created a data-driven model for the ERRANT network emulator tool [43] and make it available at <https://github.com/SmartData-Polito/errant>.

2 TESTBED AND MEASUREMENTS

For our experimental campaign, we use three off-the-shelf PCs equipped with 8 cores and 16 GB of memory running Ubuntu 20.04 and Linux kernel version 5.0.4. The two first PCs are located in the UCLouvain campus in Louvain-la-Neuve, Belgium. The first PC (PC-Starlink) is connected to the Internet via Starlink with a regular subscription. The second PC (PC-Wired) is connected to the UCLouvain campus network via a 1 Gbit/s Ethernet adapter. The third PC (PC-SatCom) is connected to the Internet via a traditional SatCom equipment for which we have purchased a regular plan offering up to 100 Mbit/s in downlink and 10 Mbit/s in uplink. The SatCom operator is a reseller and relies on a major European provider that uses geostationary satellites to provide Internet access. Our user equipment consists of a dish antenna and a modem that connects the PC to the network. For each setup, the TCP receive window is the kernel default, i.e. 131072 bytes by default with a maximum of 6291456 bytes through automatic buffer tuning. The congestion control is Cubic. We use the three PCs to run the experiments that we describe in detail below and summarize in Table 1.

QUIC measurements. Some of the performance metrics of this article are gathered using the QUIC protocol [26]. We assess the network performance with two kinds of transfers: (i) *bulk* HTTP/3 (H3) [16] 100MB transfers and (ii) *light* QUIC transfers with regularly sent messages, similar to a real-time video traffic. The latter sends 25 variable length messages per second during 2 minutes. Each message has a size in the 5-25kB range. The average bitrate of this transfer is 3 Mbit/s, far below both downlink and uplink capacities announced by Starlink. The QUIC client runs on PC-Starlink while the server is located in the UCLouvain university campus. Half of the experiments are transfers from the server to the client (download) and the other half are from the client to the server (upload). Using QUIC instead of TCP ensures that we measure the end-to-end latency as it forbids the use of middleboxes and proxies interfering with the traffic at the transport layer as it can be done for TCP with PEPs. The way QUIC identifies and retransmits packets also allows us to exactly point every lost packet and disambiguating original packets from retransmissions. The QUIC H3 server is able to provide more than 400Mbps of QUIC traffic to other endpoints connected to wired networks outside our campus. The QUIC implementation used is `quiche` [1] compiled in *release* mode from commit `ba87786`. Its initial `max_data` and `max_stream_data` transport parameters are set to 10MB and the receive window varies through automatic buffer tuning. The congestion control used is Cubic.

Latency. We measure the latency of Starlink by probing a set of 11 anchors using `ping`. Our set of anchors includes 7 servers used inside the RIPE Atlas project [8]. The servers are located in Europe (Amsterdam $\times 2$, Nuremberg $\times 2$), North America (New

Table 1: Overview of the datasets.

Measure	Network	Duration	Target
Latency	Starlink	5 Months	11 Anchors
Throughput	Starlink	4 Months	Ookla Servers
	SatCom	2 Weeks	
Web Browsing	Starlink	4 Months	120
	SatCom	2 Weeks	Websites
QUIC H3	Starlink	5 Months	Our server
QUIC messages	Starlink	5 Months	Our server

York, Fremont) and Asia (Singapore). We also include 4 nodes of the RIPE Atlas project hosted by volunteers in the same country as our Starlink connection (Belgium). Every five minutes, we measure the latency towards the anchors running 3 pings. We also measure the link latency under light and heavy network load by studying the evolution of the Round-Trip Time (RTT) measured by QUIC with our messages and H3 transfers.

Packet loss. Starlink provides a new kind of wireless network access. In general, packet losses come from two causes: congestion or medium imperfection (e.g., Wi-Fi interferences). We study the packet losses under light and heavy network load using our QUIC setup with both bulk H3 transfers and messages variants.

Throughput. We measure Starlink download and upload throughput using the command line version of the Ookla SpeedTest service [11]. The application selects the closest test server and probes download and upload capacity by opening several parallel TCP connections. We perform a speed test every half an hour using PC-Wired from 20 December to April 7 2022. We compare Starlink with SatCom using PC-SatCom, on which we run identical measurements, scheduling them at the same pace. Finally, we also measure Starlink throughput using our QUIC H3 setup.

Web Browsing. We measure the performance of Starlink for Web browsing by running on PC-Starlink automatic visits to websites and collect metrics that can be used as proxy for users' perceived QoE. We rely on BrowserTime, a tool performing automated visits to websites [2]. We rely on the rank provided by SimilarWeb [13], an online ranking service out of which we pick the top-120 website for Belgium. Among the statistics collected with BrowserTime, we focus on two metrics that have been shown to be correlated with users' QoE [19]: (i) *onLoad*: the time when the browser fires the `onLoad` event – i.e., when all elements of the page have been downloaded and parsed; (ii) *SpeedIndex*: proposed by Google [10], it represents the time at which the visible parts of the page are displayed. It is computed by recording the video of the browser screen and tracking the visual progress of the page during rendering. Every half an hour, we test 30 websites chosen at random and ensure they do not overlap with the speed test experiments. We collect data from December 20 to April 7 2022. We compare the browsing experience offered by Starlink with the SatCom link by running the same experiments on PC-SatCom and collect the resulting metrics.

3 RESULTS

In this section, we report our results and findings. We first discuss the measured latency and then focus on packet loss and throughput, comparing Starlink with traditional SatCom. Finally, we discuss

Figure 1: Distribution of the RTT to the anchors. The top axis reports the distribution minimum. Notice the logarithmic \sim axis.

Figure 2: RTT towards the European anchors.

QoE-related metrics for web browsing and the presence of middle-boxes.

3.1 Latency

We begin our analysis by looking at the RTT. We first measure the latency without load on the link, which is the best latency Starlink subscribers could achieve. We then perform QUIC downloads and uploads, thus generating bandwidth pressure and study how the RTT evolves under load.

Latency during inactivity : Figure 1 shows the distribution of the measured latencies towards our set of anchors. The axis (in logarithmic scale) represents the distribution of RTT measured by ping in the form of a boxplot: boxes range from the 25th to the 75th percentile, while whiskers range from the 5th to the 95th. The black central stroke represents the median, while on the upper axis we indicate the absolute minimum of the distribution. The 4 left-most boxes are the four local anchors. In the median case, the RTT is in $\approx 40-52$ ms and exceeded 70ms in less than 5% of cases. The minimum observed RTT for these anchors is 24-28ms. Similar considerations apply to the two Dutch anchors. The lowest RTT we observe is for the two German probes, where Starlink reaches in only 42ms in median. The lowest RTT we observe is 20.5ms, confirming Starlink's 20ms latency promise. We observe that these values allow high QoE for voice calls and are compatible with latency-sensitive services such as cloud gaming. Indeed, GeForce Now, one of the leading platforms, mandates a latency below 80ms [7]. To reach the most distant anchor points in the U.S. (San Francisco) and Asia (Singapore), the RTT is necessarily much higher, but not more than the distance between the endpoints

(a) H3 bulk traffic. (b) Messages traffic.
Figure 3: Measured per-packet RTT distribution.

would suggest. San Francisco and Singapore are reached in a median of 184 and 270ms respectively. Using traceroute, we verified the path taken by packets towards San Francisco and Singapore and the exit nodes from the Starlink network were the same as for the European anchors (i.e. one exit in the Netherlands and the other in Germany). This suggests that inter-satellite links (ISL) are not currently enabled, although ISL-capable satellites have been launched [5] and ISL activation is planned by the end of 2022 [4].

To investigate how latency evolves over time, we depict in Figure 2 various percentiles and the minimum values, focusing on European anchors. The axis spans the few months of measurements, and we compute our statistics using 6 hours bins. The picture is fairly flat, indicating stable performance and no particular changes in Starlink infrastructure over this period. The RTT to the European anchors remains constant around 50ms in median and ranges from 40ms (25th percentile) to 60ms (75th). The minimum measured latency is on the order of 20ms. Interestingly, we observe that the distribution takes on slightly smaller values of a few milliseconds from February 11 onwards - see the small step in the middle of the figure. We suspect that this improvement is related to new satellites joining the constellation in early 2022, although we have no direct evidence [9]. Moreover, we observe an increase in RTT during the last week of April and the first week of May. Since, at that time, we did not run different experiments, we speculate that in this period Starlink was more loaded or going through reorganization, but we cannot confirm this. Finally, we observe that distribution of RTT is rather flat over the hours of the day. The median RTT is around 50ms and a Mood's test suggests the samples are drawn from distributions with the same median. Similar considerations hold for throughput measurements as well, and this can hint low utilization of the infrastructure as most operator links are impacted by diurnal patterns.

Latency under load : We now study the latency evolution under link pressure. We perform HTTP/3 downloads and uploads towards our server and study the evolution of the RTT during the file transfer. Figure 3 shows the distribution of the RTT for every acknowledged packet during the experiments. We compute the downloads curve by running an additional one-week experiments session with packets captured on the server as they were too few RTTs samples coming from the client capture for download transfers. Each curve contains more than 2 millions RTT samples. We note a median 95th and 99th percentiles RTT of 95 (resp. 104), 175 (resp 237) and 210 (resp 310) ms for downloads (resp. uploads). We

Table 2: QUIC packet loss ratios

H3 #	H3 "	Messages#	Messages'
1.56%	1.96%	0.40%	0.45%

can see that the RTT increases more for uploads than download. This difference may be explained by the larger available bandwidth for downloads allowing emptying the router queues faster than for uploads, having thus a smaller impact on queuing delay for equally-sized queues.

We finally study the RTT evolution with the QUIC messages transfer. Compared to the H3 traffic, the RTT stays mostly under 100ms, similarly to the values we obtain for ping European anchors. The downloads (resp. uploads) have 50 (resp. 66) ms median RTT, 71 (resp 87) ms 95th percentile and 87 (resp. 143) ms 99th percentile RTT. The larger RTT for uploads relates quite not implementing packet pacing. The largest messages (25 kB) are thus stacked in the network's buffers making the RTT increase slightly.

Take Away: The minimum latency of Starlink is in the order of 20ms for close destinations, as publicly advertised. Under traffic load, it may increase to a few hundreds of milliseconds.

3.2 Characterizing packet losses

Packet losses can be caused by congestion or imperfections on the medium. For download, we determine losses by looking at QUIC received packet numbers on the client. As in QUIC retransmitted data have different packet numbers from the original data and QUIC does not introduce packet number gaps, every missing packet number means the packet has been lost. For uploads, we determine the received packets by looking at the ACK frames returned by the server.

Packet losses during HTTP/3 transfers : We first study the packet losses encountered during HTTP/3 bulk transfers. In this case, losses can be due to both congestion and medium imperfections. The first two columns of Table 2 show the packet losses recorded during the H3 transfers. We can see that uploads suffered from more loss events than downloads. Nearly 2% of the packets were lost during uploads while a bit more than 1.5% were lost during downloads. Figure 4a shows the measured distribution of the loss bursts lengths during H3 transfers. The loss burst length is the number of consecutively lost packets for each loss event. As we can see, the majority of loss events during uploads concerned only one packet at a time, while more than 75% of loss events during downloads concerned several consecutive packets. We also look at the duration of a loss event. Indeed, some wireless technologies such as 802.11 implement retransmission mechanisms that may delay the arrival of subsequent packets, resulting in small silent periods during the transfer [6]. As packets are captured on the client, we can compute the duration of loss events during downloads. We identified 244 008 loss events. The median loss event duration is 49 microseconds. The 75th and 90th percentiles are respectively 58 and 113 microseconds. The 95th and 99th percentiles are 1.5 and 7.5 milliseconds. We also identified a small number of longer loss periods lasting more than 1 second identifying a possible loss of connectivity.

(a) H3 transfers (b) Messaging transfers

Figure 4: Measured loss bursts distribution. Note that Table 2 shows that packet losses are far less common for messages transfers. Thus, loss burst greater frequency is only apparent.

Packet losses during low bitrate transfers : We now focus on the low bitrate messaging use-case. The two last columns of Table 2 show the packet loss ratios measured during those transfers. Conversely to the H3 experiments, the computed loss ratio is only slightly smaller for downloads than for uploads. The loss ratio is also significantly lower compared to H3. Given the low bitrate of the messaging use-case and the overall low RTT previously measured, we can expect that fewer packet losses were caused by congestion here. Note however that from the transport viewpoint, there is no known way to distinguish between congestion and medium-induced losses. This is why loss-based congestion control algorithms such as Cubic [24] interpret every loss event as a congestion signal.

Figure 4b shows the loss bursts lengths distribution for the messaging experiments. While packet losses are a lot less frequent compared to the H3 experiments (see Figure 2), the loss bursts are in general longer when occurring. We conjecture that most of the loss events occurring during the H3 experiments are due to congestion they are more frequent and only concern a few packets, while the loss events encountered during the messages transfers may be related to the medium, sometimes being even comparable to small network outages with some loss bursts of more than 100 packets (also present for H3 transfers). Concerning the loss events duration, most events for message transfers were shorter than 1ms. However, we noted 95th and 99th percentiles of 104 and 127 ms which are larger than the percentiles for H3 downloads (note that the loss events for message transfers are a lot more rare than H3 loss events and that H3 transfers probably mostly encounter congestion-induced losses). Similarly to H3, we also detected small network outages with loss events lasting more than 1 second.

Finally, we checked that those losses were neither caused by our network nor our server by running downloads for both H3 and messages transfers from a machine in Amsterdam (i.e., close to an exit point of the Starlink network) towards our H3 server. For H3 (resp. messages) downloads, over more than 5.8 M (resp. 2.8 M) packets sent by our QUIC server, only 10 (resp. 8) were lost, making loss events nearly absent outside Starlink.

Take Away: The loss events occurring when the link is loaded are more frequent and only affect a few consecutive packets. Without link pressure, the loss events are more rare, concern overall more consecutive packets and last longer.

(a) Download Throughput. (b) Upload Throughput.
Figure 5: Measured throughput distribution.

3.3 Throughput

Figure 5 shows the throughput distribution for three experiments: Ookla Speedtest on Starlink, H3 bulk download on Starlink and Ookla Speedtest on the regular SatCom access. We first discuss the Ookla speedtests and then the H3 results.

Speedtest results: By looking at Figure 5a, we can see that Starlink’s download throughput ranges between 100 and 250 Mbit/s. The median value is 178 Mbit/s, while the maximum is 386 Mbit/s. This maximum is surprisingly high given the company’s public statements, i.e., download speeds between 100 Mbit/s and 200 Mbit/s. We note that they enable the use of bandwidth-intensive services, such as High-Definition video streaming. Netflix’s 4K videos require a download bandwidth of 15 Mbit/s [6], while Disney+ recommends 25 Mbit/s [3].

The upload throughput, in Figure 5b, is significantly lower, reaching a median of 17 Mbit/s. Fewer than 5% of the cases exceed 30 Mbit/s and the highest observed rate is 64 Mbit/s. For both metrics, we cannot find a seasonality in the measurements. Looking at the different hours of the day, the median throughput varies by less than 10% with no apparent day-night cycle. Furthermore, we have not observed any increasing or decreasing trend in the measurements over our three months of experiments, and the distributions assume approximately the same average values and variability.

Comparing with traditional SatCom, we find that Starlink provides higher throughput in both scenarios. Considering download, with a median value of 178 Mbit/s, Starlink is more than twice as fast as SatCom (82 Mbit/s). The situation is similar for upload: the traditional SatCom connection inherently offers lower upload throughput (4.5 Mbit/s in median), as it is limited to a bitrate of 10 Mbit/s.

We can briefly compare these values with mobile networks looking at recent related work. Safa et al. [29, 43] conducted a large-scale measurement campaign in 2018 involving 4 European MNOs in 2 countries. For download, they found that in the best case (4G with good signal quality), mobile networks provide a median throughput of 29.5 Mbit/s. For upload, the authors found a median bitrate of 14 Mbit/s, comparable to Starlink’s 17 Mbit/s. However, keep in mind that these throughput measurements [29, 43] are already 4 years old at the time of writing and thus possibly outdated.

HTTP/3 transfers : We now measure throughput using HTTP/3 with our server located in Belgium, the same country as the Starlink access. We report the measured throughput distribution for the

(a) onLoad (b) SpeedIndex
Figure 6: Web browsing performance.

download and upload of 100MB of data in Figure 5. We ran two experiment sessions, one until the 25th of April and one starting from the 25th of April. We observed a difference of download throughput during the two sessions but the upload throughput stayed the same. All the parameters are the same for the two sessions but we observed an increase of download capacity for QUIC. Figure 5 thus shows the results for the second session as they represent the most up-to-date results for Starlink. The download bitrate sits mostly between 100 and 150 Mbit/s which is in line with what is announced by Starlink but lower than the best results obtained with the Ookla TCP speedtests and lower than what our QUIC server can deliver to other wired endpoints. We also excluded the possibility of an incorrect receive window tuning of quiche by running additional experiments with a 150MB receive window, leading to similar results. The difference in download throughput may be due to the fact that regular speedtests use at least four concurrent TCP connections while the QUIC download uses one single connection, reacting more strongly to losses [1, 35]. It is also possible that Ookla speedtests are prioritized by the operator, similarly to what happens for conventional operators. The measured upload throughput is similar for the two sessions and is in-line with the Ookla speedtest results: they have the same median, although the QUIC results are more stable.

Take Away: In download and upload, Starlink outperforms traditional SatCom. The measured throughput with QUIC is lower compared to TCP speed tests for downloads but similar and more stable than TCP speed tests for uploads.

3.4 Browsing Performance

We now quantify the Starlink performance for Web browsing. We compare the user experience of Starlink users against other access technologies. We resort to the onLoad and Speed Index metrics that have been shown to correlate with it [9]. We continuously visit a set of 120 popular websites in our country, using PC-Starlink, PC-SatCom and PC-Wired.

In Figure 6, we show the ECDF of QoE-related metrics. Starting from onLoad (Figure 6a), we find that it generally ranges from a few to 15-20 seconds, depending on the website and conditions. Starlink (solid red line), overall, provides a median onLoad of 2.12s and an interquartile range (IQR) between 1.60s and 2.78s. Experiments with SatCom equipment (blue dashed line) show that onLoad is

¹While not present on the graph, all packet captures for the first session will be provided in the artifacts of the article.

substantially larger, 1091s on median. The distribution ranges from 836s (25^c percentile) to 1359s (75^c percentile). It is likely that this performance is due to the high latency of the SatCom connections, which affects the operation of TCP and HTTP. Note that rendering a web page requires opening multiple connections to different servers to retrieve all page objects. In our dataset, a single visit results in 15 connections on average. On SatCom, opening a connection (including the TLS handshake) takes an average of 680ms, while Starlink requires only 167ms. Finally, the green dashed line reports the baseline performance of a well-functioning wired network. The median on-load is 124s, still considerably lower than the other two cases. Although we do not run experiment on mobile networks, we mention that Rajiullah et al. [38] use a large testbed of mobile nodes to visit a number of popular websites. They measure load time on the order of 20s, thus moderately higher than what we measure on Starlink.

Similar considerations apply to the SpeedIndex (Figure 6b). Starlink shows a median performance of 182s, outperforming SatCom with a 819s median SpeedIndex. Starlink performance is closer to the Wired setup, with median of 170s.

Take Away: For Web browsing, Starlink outperforms SatCom and has close performance to regular wired access. Looking at QoE-related metrics, Starlink is 80% faster than traditional SatCom.

3.5 Middleboxes and traffic discrimination

SatCom solutions often deploy PEPs to alleviate the problems due to the high link latency. Some operators also apply Traffic Discrimination (TD) to control the bandwidth used by applications on their network. In this section, we analyze the presence of middleboxes and TD on the Starlink network.

PEPs and middleboxes: We first use traceroute and Tracebox [20] to detect PEPs and middleboxes. Traceroute shows us the presence of two levels of NAT at the two first nodes: the Starlink access point (92.168.1.1) and a carrier-grade NAT node (100.64.0.1) at the exit of the satellite link. Tracebox does not show the presence of any PEP: the TCP handshake is correctly performed in the destination network. Only the TCP and UDP checksums are altered by the NATs.

Traffic discrimination: We employ Wehe [33], a state-of-the-art tool to detect Traffic Discrimination (TD). It replays packet traces of 22 popular services including video streaming (e.g., Netflix, YouTube) and video call (e.g., Zoom, Skype). It then replays the same traces with randomized bytes to prevent the operator from correlating this traffic to the original service. In case of Starlink, we launched ten times the complete Wehe tests but could not find any TD policy in place, at least for these popular services.

4 DISCUSSION AND FUTURE STEPS

This study presents an initial characterization of Starlink from the perspective of a single site in Western Europe. Our TCP and QUIC measurements show that Starlink delivers on its promised low latency and high throughput. It enables the use of latency-sensitive services that struggle with traditional SatCom. Interestingly, early simulations of LEO constellations (see Hypatia [27], among others) predicted similarly low values for RTT, especially in this first

phase of low utilization. However, we emphasize the presence of (moderate) packet loss even at low network utilization.

Given the limited time between the commercial launch of Starlink and this study, our ping latency measurements are still in the early stages in terms of temporal and spatial scale. As we aim at tracking latency evolution over time, the number of anchors we probe is limited and does not allow us to provide a complete picture of latency for a comprehensive set of targets worldwide. Inter-satellite links do not seem to be enabled, but Starlink plans to deploy them by the end of 2022. At that time, coverage will be significantly expanded, and it will be possible to study how packets are routed through the sky and how performance varies around the globe.

Our QUIC measurements reveal additional details about the RTTs and packet losses under load. During HTTP/3 bulk transfers, RTTs increase more than when applications exchange messages at a low rate. Thanks to QUIC's precise acknowledgments, our measurements show that packet losses are more frequent during bulk transfers and provide some characterization of the loss patterns.

At the application level, we have studied QoE for web browsing and found it to be radically better than traditional SatCom. Note, however, that we only studied a limited number of websites because we wanted to visit them hourly. We did not account for differences in experience that could be due to different browsers, different devices, or other factors. Also, we only visited landing pages, while a more realistic campaign should include internal pages. Finally, further measurements should assess QoE for a wider range of services, since in many cases (e.g. video calls) there are well-established performance indicators to study. However, we believe that the results in this article give a good insight to researchers of what they can expect in terms of throughput, latency and packet losses when developing solutions that may be used together with Starlink network accesses.

A ARTEFACTS

All the data gathered to compute the results discussed in this paper are publicly available. This includes pings, traceroute, Tracebox, speed test and BrowserTime results as well as more than 530 Gigabytes of QUIC packet captures along with their encryption keys. The data can be found online at <https://smartdata.polito.it/a-look-at-starlink-performance-open-data/>.

B ETHICAL CONSIDERATIONS

Our work does not employ data coming from individuals as it is based uniquely on active measurements. During our experiments, we took care to avoid harming the destination servers, anchors and crawled webpages. We run ping measurements towards anchors every five minutes and perform speed test approximately every hour. We believe such a workload cannot harm the proper operation of the targets. Regarding web measurements, we contacted each website approximately once per hour. Considering that the target of our analysis were some of the most popular websites in Western countries, our belief is not to have caused an overload on the servers or any undesirable side effect.

